The Vested Parimutuel

Paying prediction markets for information when it is worth the most

Hunch Research · Public whitepaper · August 2026 · playhunch.xyz/vpm-whitepaper

> ***Every dollar that arrives belongs to whoever was already on the other side.***

In sixty seconds

The problem. Anyone can create a prediction market now; almost nobody can *make* one. Order books need a market maker. AMMs need someone to absorb the loss on an expiring claim. The parimutuel works from the first seeded dollar with neither, but it pays the same multiple to someone who was right three seconds before the whistle as to someone who was right an hour out. Venues cope by halting trading before the event.

The rules. Two sentences, no curves.

1. **Vesting.** When a stake arrives, it vests immediately and irrevocably to the positions already standing on the other side.
2. **Matching.** A stake is accepted only to the extent the opposing side has capacity to cover it. The rest is refused at entry and returned, as a partial fill.

What follows. These are consequences of the rules, not features added to them: a minimum if-you-win payout set the instant you enter that can only rise (your losing branch still pays zero; every floor in this paper is conditional on your outcome realizing); no revision of an accrued claim by anyone arriving later; payouts that sum to the accepted pool identically; staking on the obvious winner at the buzzer returning exactly your stake, which removes the reason to lock a market before it resolves; and a market creator who seeds every outcome recovering at least their stake in every branch.

The dial. λ, the vesting fraction, is the share of each stake that vests. λ = 0 is the classic parimutuel; λ = 1 the pure mechanism; the interior is measured, not asserted (§13), and §4.5 states exactly what each point buys and voids.

The cold start. No structure lets its first arrival actually bet: a first order rests unmatched, an AMM's curve is someone's funded expected loss, a classic pool escrows one-sided money. Here the job is explicit and priced — a market opens only on an all-outcome seed, floored in every branch rather than spent (§1.3) — which is what turns venue-scale seeding into a revolving float instead of a subsidy budget.

Verification. A zero-dependency simulator, a conformance suite with published test vectors, and an O(1) reference settlement algorithm ship with this paper and are served alongside it (see the Artifacts note at the end). CC BY 4.0 and MIT.

Failure modes. Several plausible variants of this mechanism do not work, and the ways they fail are only visible under adversarial input: uncapped vesting hands the first dust position a toll on all later flow, a metered cold-start bounty leaks to buzzer entrants, and an unreserved seeding vintage voids the creation floor. Appendix A gives each one with its failing case, all pinned in the shipped code.

Abstract

We present the **Vested Parimutuel**, a settlement rule for permissionless prediction markets. A stake vests, at the moment it arrives, pro-rata to the positions already standing on the opposing outcomes; it is accepted only to the extent those positions have capacity to cover it. From these two rules a set of properties follows as accounting identities rather than as engineered guarantees: conservation of the accepted pool, a monotone win-branch floor fixed at entry, invariance of accrued claims under all later flow, neutrality of late entry (which removes the incentive that forces parimutuel venues to close entry before an event), and an exact floor for a creator who seeds every outcome. We give the settlement rule in constant time per entry and per claim, which is what makes it implementable on-chain, together with the exact residue bound introduced by fixed-point arithmetic and the block-vintage batching rule that removes the intra-block ordering game in both vesting and acceptance. We derive the mechanism's central pricing result: a stake's vesting yield is `((1−q)/q)·ln(Π_T/Π_t)`, so an entrant is priced at the pool ratio *prevailing when they enter*, discounted by the pool's remaining log growth; we also measure what the formula does not price, namely that the win-conditional yield is roughly half the unconditional yield in our crowd model, a winner's curse any user of the break-even rule must correct for. A parameter λ interpolates to the classic mechanism. We are explicit about what is not established: the mechanism has no equilibrium theorem, its late-stage pool ratio is not a probability, it makes hedging near resolution unattractive, and large early capital compresses everyone else's return. Simulation, conformance tests with fixture vectors, and every failing case we found in our own design are published with the paper.

---

1. Introduction

A prediction market is two things at once: a forecasting instrument and a financial venue. The venue needs a market-making rule. The instrument needs that rule to reward information, or the forecast is noise.

In a deep market, being early already pays: you buy at 20¢ and sell at 80¢, and that is what a price is for. The problem this paper addresses is narrower and it is a problem of the long tail: in the ten-thousandth market there is no book to sell into, no market maker willing to quote, and no operator able to subsidize. There, pools are the only structure that works from the first dollar, and a pool pays a flat multiple regardless of when the risk was taken.

One bridge in this paper's argument should be stated as an assumption, not smuggled in through the title. What the mechanism pays for, as an identity, is *time priority of capital at risk*: a position is rewarded for when it stood, not for what its owner knew, and a noise trader at vintage 1 is paid exactly what an informed trader at vintage 1 is paid (P5). The claim that this rewards *information* rests on a hypothesis about the world: that private information about an event diffuses over the market's life, so bearing counterparty risk early correlates with holding information while it is scarce. That hypothesis is an incentive argument, not a theorem; §13.2 measures its footprint in a simulated crowd and reports the places where it fails, including the fact that the pool's own late-phase forecast is *worse* than an unlocked classic pool's. The two known failure cases are already in this paper: information that arrives late (§8 is the mechanism's answer, and its weakest part) and the seed itself, which is paid best and carries no information at all (§9, §14).

Each classical structure fails the permissionless setting in a known way. **Order books** cannot cold-start: an empty book has no price, and professional market makers will not quote regional, niche, or machine-made claims; Polymarket's order books are the deployed existence proof of how far books do reach, and their economics still stop well short of the long tail. **CFMMs** trade continuously but were designed for persistent tokens; on an expiring event contract the pool rebalances into the losing side as information arrives, and the liquidity provider absorbs that. Paradigm's pm-AMM (White & Diamandis, 2024) redesigns the curve specifically for expiring claims and reshapes that loss profile; it still requires a funded liquidity provider per market, which is the constraint that binds here. **Scoring-rule market makers** (LMSR and its relatives) solve pricing elegantly and their worst-case loss is *bounded* at `b·ln n` (`b` the market maker's liquidity parameter), but that bound is a subsidy someone must fund per market. §2.2 takes up who can fund it, because the honest comparison is with a creator-funded subsidy, not only an operator-funded one. **The parimutuel** needs no market maker and carries no operator risk, but late capital entering the winning side takes value from everyone who was right earlier, so real venues stop accepting entries before the event and give up continuous trading.

The Parimutuel Market Maker (Melee Markets, 2026) frames five requirements (R1 continuous trading, R2 first-dollar cold start, R3 no subsidy, R4 entry-time price integrity, R5 profitable passive bootstrapping) and we adopt all five without modification. They are the right bar. We add two.

R6. Verifiability. A permissionless market whose pricing rule cannot be inspected has reintroduced the trusted operator through the back door. A mechanism that asks participants to accept its solvency on the strength of the designer's own testing is offering a promise rather than a guarantee, however carefully that testing was done. The rule should be publishable in full and every participant's payout independently recomputable.

R7. A structural origin for the first dollar. Floors make early capital safer; they do not summon it. A venue hosting millions of markets needs a participant class for which discovering, pricing, and seeding a brand-new market is cheap and systematic.

1.1 Contributions

1. A settlement rule stated in two sentences (§4), with its properties proved as short consequences and their hypotheses stated exactly (§5), including the block-vintage batching rule that makes same-block entry order-free in both vesting and acceptance (§4.4).
2. A constant-time formulation (§6). The rule reads as a loop over every standing opposing position per stake, $O(m^2)$ per market in the log length; it collapses to one running accumulator per outcome, updated in $O(1)$. A book of a few thousand positions already puts a single naive entry in the tens of millions of gas, growing linearly with the book forever; the $O(1)$ form is the difference between implementable on-chain and not.
3. The vesting-yield result (§7): `y(t) = ((1−q)/q)·ln(Π_T/Π_t)`, with a published verification script (discretization error 0.005%-0.02% at 10,000 steps, converging as O(1/steps)), the corollary that when a pool grows e-fold after your entry your break-even belief is exactly the pool ratio you entered at, and the measured corrections the closed form needs in practice (§7.1): composition drift and a winner's curse that roughly halves the win-conditional yield in our crowd model.
4. The λ family, stated with its composition rule and measured across its range rather than at two points (§4.5, §13).
5. An operational conformance path: a settler-injection harness with 106 published fixture vectors, versioned and content-hashed, so a third-party implementation has something concrete to pass (§5.2, Artifacts).

6. A published account of the designs we rejected and the adversarial cases that rule them out (Appendix A).

1.2 What we do not claim

The early-entry-reward gradient is not new; Pennock's Dynamic Pari-Mutuel Market has it (§2.1), and we claim no novelty for it. Nor do we claim an equilibrium result: §9 and Appendix B are incentive arguments, conjectures, and simulation evidence, and we mark which is which throughout. Nor, finally, do we claim that payment tracks information position by position: by P5 it cannot, and §1 states the diffusion hypothesis the information framing actually rests on.

1.3 Cold start: who seeds, and what it costs

The mechanism in this paper does not open a market until someone has staked every outcome (§4.4), and the first objection this rule meets is always the same one: a venue will not seed thousands of markets, a creator posting a question as an opinion will not stake it, and surely a market can simply begin with its first bettor. The objection deserves an answer at the front of the paper rather than distributed across §2, §4 and §9, because the requirement it points at is not incidental to this design. It is mandatory — the properties of §5 are bought with it, and P11 voids any market that evades it — and the mechanism's claim to the long tail stands or falls on the argument of this section: that the seed is the cheapest form the first counterparty's funding problem has ever taken.

Start with the fact the objection assumes away: **no market structure has a first bettor.** A bet requires a counterparty, so every structure meets its first arrival with something that is not yet a bet. An order book holds the first arrival as an unmatched quote, binding nobody until the other side takes it; the first mover on a fresh book has placed an order, not a bet. An AMM meets it with a curve that someone funded in advance. A classic parimutuel accepts one-sided money outright — and if opposing money never arrives, hands it back at resolution at 1×, less takeout: an escrow with a fee. The refusal-at-entry of §4.1 is the same fact every venue's first participant already lives with, stated immediately and at no cost instead of discovered at settlement. The question that separates the structures is therefore never *whether* someone must fund the first counterparty; it is who does, at what expected cost, and holding what guarantee.

Priced that way, the comparison of §2.2 and §2.4 compresses to a sentence per structure. An order-book venue funds the first counterparty by attracting a professional market maker per market, and reaches past its head markets only by paying for it; liquidity incentives are spent money, and the long tail's empty books are what unspent looks like. An AMM or scoring-rule venue funds it as a per-market subsidy with a negative expected value someone chooses to bear: bounded at `b·ln n` for the LMSR, the rebalance-into-the-loser loss for a CFMM on an expiring claim. This mechanism funds it with a seed whose worst case is nominal recovery in every branch (P6, asymmetric seeds included), whose real costs are the ones §2.2 states — fees, carry on locked capital, resolution risk — none of which is an expected loss to informed flow, and which doubles as the first rung of the resolution bond (§12). Stated with §2.2's care: the per-market capital does not disappear here any more than anywhere else. It stops being spent.

The floor is what changes the arithmetic of scale, and it is why "a platform will not seed every market" has the economics backwards. A subsidy budget is consumed; a floored seed is a **revolving float**. A venue seeding its own markets programmatically parks the seed for the market's life, recovers it in every branch, redeploys it, and holds the first standing position on every book in between, collecting a vested share of all subsequent flow. The arithmetic is small: a symmetric seed of `S` per leg opens `S(κ−n+1)` of first-vintage headroom per book (§4.3), so $50 a side at $\kappa = 9$ admits a $400 first entry, and ten thousand binary markets seeded at $5 a side hold $100,000 of float in total against zero mechanism-level expected loss. §9's crowding measurement then makes small seeds the correct policy, not merely the cheap one: a seed three times the organic pool compresses ordinary early winners from 1.70× toward their floor, so the seeding policy that costs least and the one the mechanism prefers are the same policy.

Nor must the seed come from the venue, or from whoever wrote the question. In the mechanism's terms the *creator* is defined by one act, posting vintage 0, not by authorship, and whose capital takes that seat is venue policy. A venue can carry unseeded *listings* — a market posted as a question, an opinion, a challenge — open for trading the moment any party posts the all-outcome seed and, with it, takes the §12 bond seat. And the seed need not be balanced. §4.4's joint acceptance admits any legs satisfying $a_0 \le \kappa \cdot \min_{w \ne 0} a_w$, so at $\kappa = 9$ a $450/$50 straddle is a valid opening seed for a binary market: a directional first bet at nine to one that leaves on the table exactly what the second bettor needs to exist, carries the floor no ordinary entry gets, and — having consumed the whole of the headroom its own side draws on — waits for opposing flow like any first position anywhere. The demand that "there has to be a first bettor" is met literally. There is one; the mechanism's requirement is only that the first bet be the one that makes the second bet possible, and it pays that seat for the wait.

For a venue already running classic parimutuel pools, this section is the adoption case in miniature, because such a venue has already paid for cold start — pools work from the first dollar — and pays instead at the other end of the market's life. Entry closes minutes or hours before the event, since under the classic rule late money free-rides on early money (§13.1 measures the buzzer entrant's free-ride at +25.3%), and every closed minute is handle the venue does not take. Adopting this rule deletes the lock window as an identity rather than a policy — the same buzzer strategy prices at measured negative expected value here — and arrives as a settlement-rule change, not a venue rebuild: $\lambda = 0$ *is* the rule such a venue already runs, so the dial of §4.5 is a migration path taken at the venue's own pace, κ is published policy like a fee table, and §15 gives the staging. A venue that today guarantees or tops up thin pools bears an expected cost for the guarantee; the same commitment made as a vintage-0 seed is floored. What is given up is stated where it is measured, not hidden: live late pool odds (§7, §13.2) and the secondary-layer obligation of §8.

R7 asked for a structural origin for the first dollar; the seed is that origin, and it is not detachable. P1's conservation proof leans on every book being non-empty for the life of the market, which the creation rule alone guarantees in the unbounded-κ regime (§5), and a market that somehow opened unseeded would not be a degraded market but a dead one — every book's capacity zero, no entry from anyone acceptable at any size — which is why P11 voids it at creation rather than let it stand looking tradable. An implementation should therefore treat *every market opens seeded* as part of the settlement rule, with the same standing as conservation: not as deployment guidance, but because without it there is no market to conserve.

---

2. Related Work

2.1 Dynamic parimutuel markets. Pennock (2004) is the closest prior art and deserves a direct comparison rather than a citation. The DPM prices shares off the current pool ratio, so earlier buyers receive more shares per dollar; its payoff per share is non-decreasing in same-side purchases, and it supports redemption before resolution. In our terms, the DPM already delivers a monotone floor, protection from same-side dilution, an approximate late-entry neutrality (a late buyer of a 0.90 outcome pays about 0.90 to receive 1.00), and native exit. **Properties P2 and P3 below should be understood as re-deriva-**

tions of Pennock's insight under a different rule, not as new results. What differs here: the assignment is *to identified counterparties at a moment*, not to a share price, so the payout is a pure accounting identity with no price function to choose; the late-entry neutrality is exact rather than approximate; and the capacity rule (§4.3) has no DPM analogue. The exactness cuts both ways: because P3 makes accrued claims irrevocable, the last entrant must be paid exactly 1× at $\lambda = 1$, so no member of the family with $\lambda > 0$ can reproduce the DPM's near-fair *current-ratio* pricing for late entrants, and $\lambda = 0$ reproduces it only by discarding vesting altogether, which is the free-ride the mechanism exists to remove; a venue that needs a live late price and native exit more than exact claim invariance should prefer the DPM's point in the design space, and §14 says so. Pennock & Sami (2007) survey the family; Agrawal et al. (2011) unify parimutuel call auctions and cost-function market makers as convex programs, and an open question we do not resolve is whether the mechanism here sits inside that framework or outside it.

2.2 Scoring-rule market makers. Hanson (2003, 2007) introduced the LMSR; Chen & Pennock (2007) and Abernethy, Chen & Wortman Vaughan (2013) give the utility and axiomatic characterizations. We do not claim these mechanisms are unsound: their loss is bounded and deliberate. The one-line dismissal ("a bounded per-market subsidy is still a per-market subsidy") is too quick, because the subsidy need not come from the venue: Manifold runs creator-funded AMM liquidity at scale, which puts the scoring-rule family in the same funding position as our creator seed, per-market capital at risk supplied by whoever wants the market to exist. The honest comparison is then an incidence comparison. A creator-funded LMSR subsidy has worst case `−b·ln n` and an expected loss to informed flow that someone chooses to bear; what it buys is continuous two-sided quotes, native exit, and hedging. The VPM seed has a nominal floor in every branch (P6), and what it gives up is exactly those three things; its costs are fees, carry on locked capital, and resolution risk. Both are per-market capital. The claim this paper actually makes is narrower than the slogan: the VPM converts an *expected loss* into a *floored position compensated by later entrants*, which we argue is the allocation that scales to the long tail, not that the per-market cost disappears.

2.3 Parimutuel microstructure. Thaler & Ziemba (1988) is the source of the favorite-longshot bias, which we analyse in §14 rather than merely cite; Ali (1977), Ottaviani & Sørensen (2008, 2010) and Snowberg & Wolfers (2010) develop it. Ottaviani & Sørensen's work on the timing of parimutuel bets is the direct antecedent of §9 and Appendix B. Plott, Wit & Yang (2003) give the experimental treatment of last-mover free-riding: the problem P4 addresses. Lange & Economides (2005) describe a deployed parimutuel with limit orders, which is a counterexample to any blanket claim that parimutuels offer no exit.

2.4 Automated market makers and time-weighting. Angeris & Chitra (2020) for CFMM price behaviour; White & Diamandis (2024) for the pm-AMM, the prediction-market-specific AMM discussed in §1; White, Robinson & Adams (2021) for time-weighted execution, the closest crypto-native antecedent to spreading a claim across arrival time. Manifold's Maniswap is a live at-scale answer to cold-starting user-created markets: a CPMM whose per-market liquidity is funded by the creator, so it meets cold start by paying for it, which is the same trade our seed makes with a different loss profile (§2.2). Azuro and Overtime are deployed on-chain pooled designs whose liquidity providers likewise fund the quote. None of these deployed answers is subsidy-free; that is the empirical shape of the constraint R7 names.

2.5 Resolution. §12 does not propose a new oracle. Augur (Peterson & Krug, 2015), UMA's Optimistic Oracle, and Kleros (Lesaege, Ast & George, 2019) are the deployed designs for staked resolution, escalation, and dispute juries; our contribution there is operational rather than mechanical.

2.6 The Parimutuel Market Maker. Melee Markets (2026) meets the five requirements with co-adaptive outcome price curves. The litepaper states of its own design: *"The PMM's production curve family,* *parameter schedule, and rebalancing implementation are proprietary and are not disclosed; every property described in this paper is observable mechanism behavior"* (p. 2), and of its solvency: *"This floor-solvency property has been formally verified against the settlement logic"* (p. 6). Our disagreement is exactly and only with the first of those: a non-public curve is not necessarily an unsound one, and we make no claim that it is. Our objection is that a permissionless venue's soundness should not require the participant to take the venue's word for it. On cold start the honest comparison is symmetric: both mechanisms require someone to commit capital before organic flow. Their answer is an optional presale phase clearing at a constant price (p. 5); ours is a mandatory creator seed on every outcome. The seed is a funded phase by another name, and we do not pretend otherwise; what we claim for it is that it is floored (P6), that it doubles as the resolution bond §12 wants, that it needs no separate clearing rule, and that it works with a single participant. Their own limitations section notes that without a presale, *"the earliest entrants receive the entire share of counterparty-liquidity rewards"* (p. 10). That is a concentration effect converging with what we find in §9, and we cite it as independent corroboration rather than as a point against them.

---

3. Model and Notation

A market is an event, a resolution criterion, and outcomes O with $n = |O| \geq 2$, mutually exclusive and exhaustive. $\omega \in O$ is the realized outcome. Trading is admitted on $[0, T]$. The trade log $E = (e\_1 ... e\_m)$ is processed in log order; $e\_k = (\tau\_k, o\_k, c\_k)$ is a stake of c_k integer units on outcome o_k at time τ_k. Entries sharing a block share a **vintage** ν_k; §4.4 gives the batching rule.

symbol	meaning
s_i	principal of position i **actually accepted** (≤ the amount offered)
$v\_i(t)$	vested claims of i at t, contingent on $o\_i = \omega$
$F\_i(t)$	win-branch floor, $s\_i + v\_i(t)$
$P\_w(t)$	total accepted principal on outcome w at t (the book)
κ	capacity coefficient: each position grants its book $\kappa \cdot s\_i$ of matching capacity
C_w, V_w	book w's cumulative granted capacity ($\kappa \cdot P\_w$) and total vested-in
λ	vesting fraction; $\lambda = 0$ classic, $\lambda = 1$ pure
$\Pi(t)$	accepted pool at t
q	a side's share of the accepted pool, P_w/Π; in §7 it is *your* side's share
Π_i	payout to i at resolution

The unsubscripted Π always means the accepted pool; a position's payout is always written Π_i. Three conventions matter and are part of the specification, not implementation detail. First, a stake is assigned **in full in each of the $n-1$ opposing branches**, not divided among them. Because branches are mutually exclusive, at most one assignment is ever realized, so this is not money creation; per-branch conservation holds exactly, and it also means $V\_w = \Pi - P\_w$ when every entry is fully accepted. Second, integer allocation uses a deterministic rule (largest-remainder in the reference settler, floor division in the constant-time on-chain form of §6), and the two differ by a bounded residue that §6 quantifies. Third, every market carries a designated **residue owner**, an address fixed before the market opens, like κ, λ and the fee schedule (§6). Finally, a time subscript abbreviates a time argument ($\Pi_t = \Pi(t)$; on the book fixed by context, $P_\tau = P_w(\tau)$ and $q_T = q(T)$); position indices, as in Π_i and s_i, are never times.

---

4. The Mechanism

4.1 The two rules

> ***Rule 1. Flow vesting.*** *When a stake of c arrives on outcome o at time t, then for every other outcome $w \neq o$, that stake is assigned, contingent on w winning, pro-rata by principal to the positions on w existing at t. The assignment is immediate and irrevocable, conditional only on the market resolving (§12 specifies the void path).*
>
> ***Rule 2. Capacity matching.*** *Each position grants its own book $\kappa \cdot s_i$ of matching capacity. A stake is accepted only up to the capacity remaining in every opposing book; the unmatchable remainder is refused at entry and returned in the same transaction.*

A position on o entered at τ with accepted principal s is paid, if o wins,

```
Π_i = s + Σ over every stake c accepted on any other outcome after τ:  c · s / P_o(t_c)
```

Rule 1 alone would be unsound. With no cap on what a book may absorb, the first position on an empty book receives *all* opposing inflow until a second joins, so return on capital diverges as the position shrinks. Measured in the design phase on the uncapped settler, a seeding straddle returned +155% at $50 a leg and +82,037% at one cent a leg (the shipped suite pins the same divergence as P5's 200,001× single-probe comparison), so the profit-maximizing strategy would be to post dust on every new market and toll the organic flow while supplying no liquidity at all. Rule 2 is what makes Rule 1 a mechanism rather than a toll booth, and its placement matters: **capacity constrains acceptance at the book level, never the distribution**, which stays pure pro-rata. Pro-rata makes return on capital *scale-invariant*: a dust position earns exactly the multiple any capital entering at that moment earns, never more, so there is nothing for dust to dominate, and the acceptance cap bounds that common multiple by $1 + \kappa(1 + \ln g)$, g the book's subsequent principal growth (§5, P5); it is *not* a flat $1 + \kappa$ except on a static book. Keeping the distribution uncapped is also what keeps §6's constant-time form exact: a per-position cap would have broken the accumulator (Appendix A.4). Rule 2 is what a betting exchange already does. It is a matching constraint, with the partial fill as its familiar consequence.

4.2 A worked example

Binary market, $\kappa = 9$, fees zero, $\lambda = 1$. The creator seeds $25 on each outcome in a reserved vintage 0 whose legs match each other (§4.4). Five ordinary stakes follow. This exact log, both branches, ships as the first pair of conformance vectors, so every number below is machine-checked.

#	time	who	outcome	stake
0	0	creator	YES / NO	$25 / $25
1	10	A	YES	$100
2	30	B	NO	$200
3	60	C	YES	$300
4	90	D	NO	$200
5	99	E	YES	$200

Accepted pool $1,050. Settlement, both branches, against the classic parimutuel on the same log:

If YES realizes (classic pays every YES holder a flat 1.680×):

position	staked	vested	payout	multiple	classic
creator (YES leg)	$25	$76.76	$101.76	**4.070×**	1.680×
A (t=10)	$100	$207.06	$307.06	**3.071×**	1.680×
C (t=60)	$300	$141.18	$441.18	**1.471×**	1.680×
E (t=99)	$200	$0	$200.00	**1.000×**	1.680×

If NO realizes (classic pays a flat 2.471×):

position	staked	vested	payout	multiple	classic
creator (NO leg)	$25	$170.09	$195.09	**7.804×**	2.471×
B (t=30)	$200	$360.79	$560.79	**2.804×**	2.471×
D (t=90)	$200	$94.12	$294.12	**1.471×**	2.471×

(The classic NO multiple is 1050/425 = 2.47059; "flat" holds up to integer allocation, which shifts individual classic payouts by a cent, so the creator's $25 leg settles at 2.470× and B's at 2.471× under largest-remainder cents.)

Read six things off this table.

1. **Conservation.** Payouts total $1,050 exactly in both branches.
2. **The whole point.** The classic column is flat: 1.680× to A, who carried the risk for 89% of the market's life, and 1.680× to E, who arrived after the last opposing dollar. The vested column pays 3.071× and 1.000× for the same two positions, out of the same pool.
3. **Late-entry neutrality.** E enters after the last NO inflow and is paid exactly its principal. Under classic rules E takes $136 of profit, every cent of it from A, C and the creator.
4. **The floor.** The creator staked $50 across both legs and recovers $101.76 or $195.09, above break-even in both branches, because at vintage 0 the legs are each other's counterparty.
5. **Monotonicity.** A's claim of $207.06 was accrued before E arrived, and E's arrival does not touch it. Under classic rules A's multiple falls from 2.000× (immediately before E: pool $850 on a YES book of $425) to 1.680×: E takes 16% of A's expected payout having borne no risk.
6. **What a later entrant gets.** C is not punished for being third; it is paid 1.471× on $300. It simply shares only the flow that arrives after it.

4.3 Capacity, and why the bounty is gone

The obvious alternative is to hold unmatchable flow in an "unallocated bucket" and meter it out to later entrants at up to `ρ×` their stake, decaying with market age. We rejected it (Appendix A.2): it is strictly dominated by the dust position above; its residue at resolution has to be refunded, which leaves a losing stake not fully at risk; and its decay is indexed to clock time while the risk it prices is informational, so a market that settles informationally early reinstates most of the free-ride. Rule 2 needs none of it. There is no bucket, no `ρ`, no decay schedule, and no residual refund. In their place is a single published constant `κ`, a maximum-odds cap, the same object every betting exchange already publishes. Acceptance headroom on a book is `C_w − V_w`; with stakes assigned in full to every opposing branch this is the same test as `(κ+1)·P_w > Π`, and §6 computes it from two running scalars.

The cost is larger than the binary case suggests, and the blocking economics deserve an honest statement. **In a binary market** the favoured side receives partial fills once the underdog book is saturated. That is an honest liquidity signal. It is *not* free to weaponize for a trader who does not want the position: to exhaust the headroom a buyer needs, you must stake their side yourself, at full risk (conformance case P10, where the squatter eats the loss when the block side loses). But for a trader who wants the exposure *anyway*, the block is free at the margin: by scale-invariance the blocking capital earns the same multiple whether or not it crowds the rival out (conformance case P10b), and seizing the remaining headroom converts shared future vesting into an exclusive claim on it. The harm is redistributive among same-side entrants, and informational: the blocked rival's information never enters the pool. Under a binding κ, entry is therefore a race for headroom between same-side traders, which is a concentration channel §9 treats alongside crowding-out; the block-vintage rationing rule of §4.4 removes the *within-block* race, and per-account acceptance caps are the venue-level mitigation for the rest, at the cost of sybil pressure. **In an n-way market the constraint couples across branches**, because acceptance takes the minimum headroom over *every* opposing book: entry on any outcome requires positive headroom in each opposing book `w`, so the *thinnest* book gates entry on all outcomes opposing it. That is both a liveness failure, since a three-outcome market with one unloved branch refuses 97.5% of gross volume at κ = 9 (the pinned skewed-book construction, conformance case P9b), and a griefing lever, since a $450 stake on a side outcome blocked a $2,000 informed entry entirely (P9).

The sharper statement is that for `n ≥ 3` the degraded state is **absorbing, and its trigger has a closed form** (conformance case P9c). A symmetric seed `S` leaves every book exactly `S(κ − n + 1)` of headroom. One stake of that size on a single outcome consumes the headroom of the `n − 1` books opposing it, and thereafter *every* outcome has a saturated book among its own opponents, so **no entry on any outcome at any size is ever accepted again**: capacity grows only with a book's principal, and principal can no longer grow. The market is not throttled, it is dead, and it cannot be revived from inside the mechanism, because the stake that would replenish a starved book is itself refused. At the paper's own parameters (κ = 9, n = 3, a $50 seed per leg) the trigger is $350, the attacker's own stake is capped at exactly that, and the freeze is reachable at any point in a market's life rather than only off the seed: after six ordinary entries the same construction costs $770 and is equally permanent. `n = 2` is the sole exception and it self-heals, because the single opposing book is precisely the one a recovering entry needs; that asymmetry, not a difference of degree, is what makes finite κ a binary instrument. The balanced n = 3 arm in §13 reports 1.03% refusal, and that number is *not representative* of skewed books; P9 pins the skewed case. The rule for venues: **finite κ is a binary-market instrument.** Markets with three or more outcomes, ladders with far-out-of-the-money bands especially, should run κ large or unbounded; the mechanism's soundness does not depend on the cap. What holds it up without κ is the pair that Rule 2 brought with it and that stays: the mandatory all-outcome seed (no empty books can exist) and pro-rata scale-invariance (dust earns what any same-moment capital earns, so P5's dominance analysis needs no cap). With κ unbounded, thin-book multiples are simply long odds, exactly as in a classic parimutuel; what remains bounded is what matters, namely nobody's accrued claim and nobody's dominance.

4.4 Creation, entry, exit

Creation. A market opens when its creator posts a seed on every outcome in a reserved **vintage 0** that no other entry may share. The legs are counterparties to one another, which is what floors the seed (P6) and what makes the floor robust: there is no one to interpose. The seed obeys Rule 2 among its own legs, and because each leg's capacity comes only from the other legs, vintage-0 acceptance is a joint condition: the accepted legs are the greatest amounts `a_o ≤ offered_o` satisfying `a_o ≤ κ·min_{w≠o} a_w`, computed as a fixed point (the reference settler iterates it), so an asymmetric seed is partially refused rather than allowed to violate the leverage bound, and any symmetric seed requires `κ ≥ 1` to be accepted at all. A seed that leaves any outcome unbacked, including by integer allocation flooring a dust leg to zero, **voids the market at creation** (conformance case P11); this creation rule, not Rule 2 alone, is what guarantees every book is non-empty for the life of the market, and P1's proof leans on it. This also gives §12 the two-sided bond its resolution story wants, with the bond's limits stated there.

Entry. A buy is accepted up to capacity, creates a position, and vests to all opposing books. All entries in one block share a vintage, and the batching rule is normative, not an implementation choice: **(i)** every entry in a vintage vests to the books *as they stood at the start of the vintage*, so same-vintage entries never vest to each other; **(ii)** capacity granted by a vintage's entries becomes usable only from the next vintage; **(iii)** when the joint same-vintage demand `D_w` against an opposing book exceeds its headroom `H_w = C_w − V_w`, every entry is filled at the same book fraction, `⌊c_e·H_w/D_w⌋`, and an entry's accepted amount is the minimum of its rationed caps over the books it opposes. Rule (iii) is what makes the claim complete: without it, acceptance under a binding κ would fill in intra-block order and the priority race the vintage removes from vesting would reappear exactly where priority is most valuable. With it, settlement is invariant under permutation within a vintage in both vesting and acceptance (conformance cases V1-V3), so there is no intra-block *ordering* game and no priority auction over vintage. Order-free is not size-strategyproof: an entry's rationed share rises with its offered size while the refused remainder costs nothing, so under a binding κ the within-block contest becomes an oversizing game, the standard property of pro-rata matching everywhere it is used. A venue that needs to blunt it can fee or bond the offered amount rather than the accepted one, cap per-account offers, or run κ large per §4.3; the mechanism-level claim here is order-invariance only. A zero-accepted entry holds no claim and no capacity, and need not be stored (the reference keeps an inert record for accounting symmetry). Inter-block priority is a different matter and we treat it as a limitation, not a solved problem (§14).

Exit. Positions are transferable; a sale moves principal, vested claims and vintage intact without touching the pool. A position may be split or merged within its vintage class; integer division of vested claims uses the market's allocation rule, with sub-unit residue following the residue owner (§6).

Fees. The fee base is venue policy (§5.1); the reference convention, assumed throughout §13, charges the fee on the *accepted* stake at entry. Whatever the base, the refused remainder of a partial fill is returned gross, and the mechanism operates on net amounts.

4.5 The λ family

Vest a fraction λ of each stake by Rule 1 and settle the remaining 1−λ as a classic terminal pool. The composition with Rule 2 is part of the specification: **acceptance, capacity, and the seed clamp operate on the full stake exactly as at λ = 1; only the settlement of the accepted amount splits.** Because vesting is linear and every stake vests the same fraction, a winner's blended payout collapses to a closed form from one λ = 1 settlement, `Π_i(λ) = λ·(s_i + v_i) + (1−λ)·s_i·M` with `M` the classic multiple of the accepted pool, which is how the reference implements it and how the λ table in §13 is generated. λ = 0 is the classic parimutuel up to Rule 2's refusal (0.10% of gross in our study); λ = 1 the pure mechanism.

λ is a dial on a trilemma, and the honest statement is what each point buys and voids:

	no-lock corollary (P4)	creation floor (P6)	dilution protection	late hedging
λ = 1	exact: buzzer entry pays 1×	unconditional	maximal	none
interior λ	void: a buzzer snipe pays `λ + (1−λ)·M`	**void** (measured at λ = 0.5: seed min −16.8%)	`≈ (1−λ)` of classic dilution on average	partial
λ = 0	void	void	none	full (classic)

No point on the dial delivers all four; a venue must choose per market class, and §14's hedging paragraph names the consequence: a market run at interior λ for hedgers has given up the creation floor, so the R7/§10 seeding story does not apply to it, and its lock-window incentive is back in proportion to 1−λ. Degradation in between is **monotone but not uniform**: the late entrant's reward is exactly linear in λ (P7), while the mean dilution measured in §13 sits strictly below the linear reading at every interior λ, so the dial protects early capital better than a linear reading suggests, with the per-position caveat P7 states.

---

5. Properties

Stated for λ = 1 unless noted. Every hypothesis below is load-bearing: drop one and the property is false, in ways Appendix A makes concrete.

P1. Conservation. For every log, outcome and λ, payouts sum to the accepted pool exactly, in integer units. *Proof.* Fix ω and follow the contingent-ω ledger. Every accepted unit is placed, at arrival, into exactly one of: the principal of a position on ω, or the vested claims of a position on ω. No rule ever decrements either. This is a bijection between accepted units and payout units. What guarantees a unit can always be so placed is that every opposing book is non-empty: vintage 0 seeds every outcome and a seed that fails to voids the market (§4.4), and at finite κ Rule 2 independently refuses any stake an opposing book cannot cover. At unbounded κ, the regime §4.3 recommends for n-way markets, the creation rule alone carries the invariant, which is why it is specification and not implementation. ∎ *Pinned:* the conformance suite asserts the equality `Σ payouts = accepted pool`, per branch, in integer units, on every vector.

P2. Monotone win-branch floor. `F_i(t) = s_i + v_i(t)` is non-decreasing and equals the payout if `o_i` wins. *Proof.* `v_i` is a sum of non-negative increments; nothing subtracts. ∎ *This is a floor on the win branch only.* If the position's outcome does not realize it pays zero, and if the market voids the position refunds at principal (§12). Every use of the word "floor" in this paper is conditional in that sense.

P3. Accrued claims are invariant. Appending any stake to the log leaves every previously accrued `v_i` unchanged. *Proof.* Rule 1 writes only additively, into the books as they stood at the arriving stake's vintage, and never revisits an earlier allocation. ∎ *What this does and does not say.* A later same-side entrant shares only *future* opposing flow. It cannot touch what you have accrued, but it can reduce what you go on to accrue. §13 measures both under the full mechanism: under a 50%-of-pool late snipe, early winners' accrued claims are untouched by construction, and their total payout falls 0.19% on average (4.9% worst across 20 seeds), against 13.6% average and 64.4% worst under classic rules, measured in the same experiment on the same triggers.

P4. Late-entry neutrality. A stake accepted after the last opposing inflow is paid exactly its principal; for a trader with belief `p < 1`, its expected value is `−(1−p)s < 0`. *Proof.* No further increments arrive under Rule 1; apply P2. Under Rule 2 there is no unmatched residue to recover, so the negativity is unconditional. ∎ **Corollary: no lock window.** The buzzer snipe that forces venues to close entry early pays exactly 1× here, so the incentive has no target and markets can accept entries until resolution. §13 completes the statement with the unconditional number: the snipe's expected value is −1.05% pre-fee under this mechanism and +25.3% under classic rules, on the same triggers.

P5. Linearity, scale-invariance, and bounded leverage. Vesting is linear in principal, so splitting a stake across wallets at one vintage changes nothing, and return on capital is *scale-invariant*: every position on a book earns the same accumulator increment per unit of principal (§6), so a dust position earns exactly what honest capital entering at the same moment earns. Because acceptance is capped at the book level, that common multiple is bounded: while a book's principal is constant, accepted opposing inflow cannot exceed its remaining capacity `κ·P − V`, and each new unit of same-side principal adds at most `κ` more, priced at the enlarged book. Integrating, a winner's multiple is at most `1 + κ·(1 + ln(P_T/P_τ))`. *Verified:* every winner obeyed the bound over 14,520 positions in adversarial random markets; an ε-position probing a seeded market returns 9.0-9.7× at κ = 9 across stakes of 1 to 10,000 units (the small spread is acceptance, not distribution: larger probes grow the book they join), against 200,001× for the 1-unit probe under Rule 1 alone. *This is not a general sybil-resistance claim, and the wash-trading cost has a shape.* A trader staking `x` on each side into books owned by strangers vests to those strangers on both sides and receives nothing from their own opposing leg, since same-vintage entries do not vest to each other: pinned (conformance case P5w): a \$100/\$100 wash into an existing \$100/\$100 market returns −50% in both branches when the legs share a vintage, and −25%/−50% by branch when sequential, before fees. But the donation is `(1−f)`-scaled, `f` the washer's share of the opposing book: across vintages a wash leg vests pro-rata into a book that may include the washer's own earlier positions, so for a participant who dominates both books the mechanism-level cost of wash volume falls toward zero and the binding cost is the fee. Rule 2 does not impede sustained alternating self-play, because each leg grants κ times the headroom the next consumes. In the whale-seeder regime §9 flags as the mechanism's most likely drift, printed volume is therefore nearly free at the mechanism level, which is why §12 requires trust metrics to exclude self-vested flow rather than lean on volume.

P6. Creation floor. A creator who seeds every outcome in the reserved vintage 0 recovers at least the total seeded, in every branch and under every continuation. *Proof.* Immediately after the seed the creator owns every position in the market; by P1, every branch pays the whole accepted pool to its owner. By P3 later flow only adds. ∎ *Verified* over 8,910 branches at n = 2 to 4 with random later flow: zero violations, worst case exactly break-even; asymmetric seeds inherit the floor through the clamp (P1'/P6'). *The hypothesis is essential.* The floor does not hold across all histories without the reserved vintage: an implementation that admits any third party into vintage 0 breaks the floor in half of all branches, worst case

−50%. That failing case is pinned in the conformance suite, alongside the conformant behaviour (a stake arriving in the creation block is assigned vintage 1 and the floor holds).

P7. Blend degradation. At blend λ a late entrant's multiple is exactly `λ·1 + (1−λ)·M`, which is linear in λ. For dilution the correct statement is signed, not one-sided: under a late entry, position `i`'s dilution satisfies `D_i(λ) ⋚ (1−λ)·D_i(classic)` according as the position's pure multiple `1 + v_i/s_i` is above or below the classic multiple `M`, with equality at `M`; equivalently, the crossover is whether the position's entry accumulator lies below the principal-weighted mean entry accumulator of the winning book. No one-sided averaged bound survives either: in aggregate, the payout-weighted dilution of the pre-existing winning book under an entry appended after the last opposing inflow is *exactly* `(1−λ)` times its classic value (case P7c), while the principal-weighted mean of per-position dilutions exceeds that line whenever pure multiples differ, by Jensen's inequality on the convex per-position law (case P7c'). What the §13 λ table shows is the empirical fact that this crowd's mean dilution sits below the linear reading at every interior λ. An individual early position on the wrong side of the signed crossover can exceed the naive per-position bound, and the suite pins a log where a position earlier than the stake-weighted mean vintage does exactly that (cases P7a-P7b).

5.1 What the mechanism does not guarantee

The list is exhaustive over settlement-accounting guarantees; behavioural and microstructure costs are §14's subject, and the two lists cross-reference rather than repeat each other.

- The floor is **conditional on your outcome realizing.** It is not downside protection.
- Late pool entry is **supposed** to die. A VPM market's late pool volume should be read as noise, not information, and the pool's late-phase forecast quality measurably degrades (§13.2, §14).
- There is **no equilibrium theorem.** §9 and Appendix B are arguments, conjectures, and simulation evidence.
- **Large early capital compresses everyone else's return**, and nothing in the mechanism prevents it (§9).
- **$\lambda < 1$ voids P6** and re-opens the lock-window incentive in proportion to `1−λ` (§4.5). The classic component pays the seeder less than its stake whenever its realized-side share falls below its pool share; the measured minimum at $\lambda = 0.5$ is −16.8%, and even $\lambda = 0.75$ dips to −3.0% (§13).
- `κ` and `λ` are **venue policy**, published like a fee schedule, not identities; so are the fee base (§4.4), the minimum stake unit, and the residue owner (§6).
- The mechanism says nothing about **resolution**. A perfect settlement identity resolved by a liar pays the liar's friends exactly, and §12 names the two resolution attacks this mechanism makes *sharper* than a classic pool's, with the mechanical fixes.

5.2 Conformance cases beyond P1-P7

The paper's property namespace runs P1-P7; the conformance suite pins more than the properties, and the extra cases are cited throughout by these names:

- **P8**: the O(1) accumulator (§6) settles identically to the naive reference, exactly in rational arithmetic, within integer rounding in cent arithmetic.
- **P9**: the n-way thin-book coupling: a \$450 griefing stake blocks a \$2,000 informed entry at $\kappa = 9$ and the same entry clears at κ unbounded (§4.3).

- **P9c**: the n-way freeze as an *absorbing* state: post-seed headroom is exactly `S(κ − n + 1)` (pinned for n = 2..5), one stake of that size permanently blocks every outcome for n ≥ 3, the same construction is reachable mid-market, and n = 2 self-heals (§4.3).
- **P10 / P10b**: the binary squat, both branches: the squatter bears the blocked side's full risk when it loses, and pays nothing at the margin for exclusivity when it wins (§4.3).
- **P11**: seed validity: a seed leaving any outcome unbacked, including by integer flooring, voids at creation.
- **V1-V4**: block-vintage semantics (no intra-vintage vesting, permutation invariance, pro-rata rationing under a binding κ, §4.4) and the resolution-void refund identity (§12).
- **P7a-P7c'**: the λ family: exact linearity of the late multiple, the signed per-position dilution law with its counterexample, the aggregate equality, and the reverse-Jensen direction of the per-position mean.
- **P9b, P10b, P5w, A4, P1'/P6'**: the skewed n-way refusal construction, the winning branch of the squat, the wash-cost pins, the per-position-cap breakage that killed the A.4 design, and asymmetric-seed conservation and floor under the clamp.

An implementation claims conformance by passing the property suite and reproducing the 106 published fixture vectors exactly on acceptance and to one unit per position on payouts; the Artifacts note gives versions and hashes, and the served README and harness header document the settler interface.

---

6. Settlement in Constant Time

Rule 1 reads as a loop over every opposing position on every stake: $O(m^2)$ per market in the log length. On-chain that is fatal: at roughly 5,000 gas per position touched, a book of a few thousand positions already puts a single entry in the tens of millions of gas, beyond any per-transaction budget a user or an RPC will carry, and growing linearly with the book forever. It is tempting to describe the settlement identity as one loop over a public trade log, cheap to compute on-chain. That is exactly backwards.

It does not need the loop. The rule is linear in principal, so it collapses to a reward-per-share accumulator. Maintain per outcome the principal `P_w` and one scalar `A_w`. On a stake of `c` on `o`:

```
for each w ≠ o:   A_w += c / P_w         (O(1))
record position:  (o, s, A_o at entry, vintage)
P_o += c
```

and at resolution, position `i` on the realized outcome is paid

```
Π_i = s_i · (1 + A_ω(T) − A_ω(τ_i))
```

which telescopes to exactly the sum in §4.1. Block vintages batch cleanly: per vintage, snapshot each `A_w` and `P_w` at the vintage start, accept and record every entry against the snapshot (with the §4.4 rationing when headroom binds, one scale factor per book per vintage), then apply the summed increments and principals, still O(1) per outcome per vintage. Rule 2 costs nothing here: acceptance is `min over w ≠ o of (C_w − V_w)` with `C_w` and `V_w` two more running scalars per outcome, so the *entire* mechanism (matching and batching included) is O(1) per outcome per entry and O(1) per claim. State is a few words per outcome and four per position, and the four numbers are the complete transferable state of a position, which is what makes §10 and §11 possible.

Verified for the shipped mechanism, Rules 1 and 2 together: the accumulator and the naive reference agree **exactly, in rational arithmetic**, across 26,458 positions in 800 random markets at $\kappa = 3$ and $\kappa = 9$, including settlements where the capacity cap binds, partial fills occur, and multi-entry block vintages ration under the §4.4 rule; the same equivalence in integer-cent arithmetic holds to within rounding, bounded by one cent per opposing event (conformance case P8). One claim, two arithmetics: the rational form is the proof, the cent form is what an implementation actually runs.

This is also the sharper statement of R6. Conservation is a weak property: it holds for the classic parimutuel too. The distinctive one is that **each position's payout is a function of two scalars, its outcome's accumulator at entry and at resolution, so no participant's payout depends on any other participant's record.**

On-chain arithmetic. With fixed-point floor division the identity becomes an inequality that can never break in the dangerous direction: payouts plus refused stake plus a residue equal the pool, with residue ≥ `0`. The safe half is unconditional, because every floor under-pays: across 27,150 settlements at scales 10^{12}, 10^{18} and 10^{27} and 800 more with stakes to 4×10^{15} units, the accepted pool was never overpaid. The residue bound is *not* one unit per winner at all scales. The analytic bound is

```
residue  ≤  W  +  Σ_i s_i·m_i / S
```

with `W` the number of winning positions, `m_i` the opposing vesting events after position `i`'s entry, and `S` the fixed-point scale: one truncated unit per winner at claim time, plus each accumulator truncation, up to one `S`-scaled unit per opposing event, multiplied back by the position's stake. The "one unit per winner" reading holds exactly when `max s_i·m_i « S`, which our small-stake suite satisfies by construction and real deployments must engineer: **choose `S ≥ s_max·m_max`**, or store the accumulator as a rational. Concretely, `S = 10¹⁸` holds the one-unit reading for 6-decimal USDC up to roughly 10^9-unit positions and hundreds of events, and comfortably beyond; at our adversarial sizes, stakes to 4×10^{15} units, `S = 10¹²` strands tens of thousands of units per winner (measured: 68,000) while `S = 10¹⁸` still holds the one-unit reading, never insolvent in either case, only stranded. Dust-position spam cannot farm the residue: the truncation is global rather than per-position-per-event, so 300 one-unit positions against a whale extract at most 3 units. The residue is a single computable claim, `accepted pool − Σ payouts`, and it **must have a named owner fixed before the market opens** (the venue's fee sink, the resolver, or the last claimant, as policy): an unassigned residue is funds no one can withdraw; the suite's void case (V4) pins the refund identity, and every fixed-point check verifies the residue is the exact non-negative difference between the accepted pool and the paid claims.

Implementation requirements. Four settlement-path rules are specification for any on-chain deployment, stated here once because the arithmetic suite cannot execute them: **(a)** the settlement asset must be transfer-exact (no fee-on-transfer or rebasing), or all accounting must use measured balance deltas; **(b)** payouts and the residue claim must be **pull-based**: one blocked or blacklisted recipient must not be able to delay any other claim, and §6's per-claim O(1) form exists precisely so claims are independent; **(c)** all state updates (`P`, `A`, `C`, `V`, position records) must complete before any external transfer in an entry, and the same-transaction refund of a partial fill must revert the entry atomically if it cannot be delivered; **(d)**

batched or relayed entry paths must isolate per-stake failure. These are MUSTs in the sense of the integer-allocation rule in §3.

---

7. The Vesting Yield

The mechanism's pricing content is one formula, and its exactness has a scope. In full generality a stake `s` on side `o` accrues the path integral `V = s·∫ dΠ_opp / P_o`, the opposing inflow weighted by the side's principal as it stands. When the pool's composition `q` (your side's share) is constant over the interval, this closes to

```
V = s · ∫ (1−q)dΠ / (qΠ)  =  s · ((1−q)/q) · ln(Π_T / Π_t)
```

Write `L = ln(Π_T/Π_t)` for the pool's remaining log growth and `y = ((1−q)/q)·L` for the **vesting yield**. Break-even is `p(1+y) = 1`, giving:

> ***An entry is profitable iff*** **`p/(1−p) > [q/(1−q)] / L`.**

Verified: the published script (`vpm-yield.mjs`) holds composition exactly constant and measures pure discretization error: 0.005%-0.02% relative at 10,000 steps across `q ∈ [0.3, 0.9]` and `L ∈ [0.5, 2]`, converging as O(1/steps).

Three consequences, then the corrections.

Entry-time price integrity. Under the classic parimutuel, break-even is `p > q_T`: you are priced at the **closing** ratio, which is unknowable when you act. Here you are priced at the ratio **prevailing when you enter**, discounted by remaining growth. R4, precisely stated, is integrity with respect to *pool composition at entry*: what you are priced off is knowable when you act. It is not integrity with respect to the conditional law of future flow, which is §7.1's subject.

`L = 1` is the natural scale. If the pool will grow exactly e-fold after you, the condition collapses to `p > q`: enter if and only if you are more bullish than the pool is right now. That is textbook aggregation behaviour, and it falls out of the rule rather than being designed in.

`L > 1` opens a two-sided band. Both sides are simultaneously positive-EV whenever `p/(1−p)` lies in `((q/(1−q))/L, L·(q/(1−q)))`, non-empty exactly when `L > 1`. Early in a fast-growing market, the two implied prices do not sum to one. **The primary pool ratio is therefore not a probability**, and the paper should not be read as claiming it is. It is a pool composition. §8 is where a probability comes from, and §14 is honest about how well that works.

7.1 What the formula does not price

Two corrections separate the closed form from a usable entry rule, and both are measured in the published suite rather than waved at.

Composition drift. When `q` moves after entry the yield is the path integral, not the closed form. Under our study's crowd the closed form is *roughly unbiased in aggregate* (mean signed error −0.03 in yield units) but poor pointwise: the median absolute gap between realized and predicted yield is 55% of the prediction, with a mass point at 100% from positions whose side simply receives no further opposing flow.

The formula is a pricing model, not an oracle, and an agent's edge in using it is exactly as good as its flow forecast.

The winner's curse. The break-even treats the yield and the win event as independent. They are negatively correlated by construction: vesting comes from opposing inflow, opposing inflow is disproportionately informed, and informed opposing flow concentrates in the histories where your side loses. Payment is contingent on winning, so the decision-relevant quantity is `E[y | win]`, and in our crowd model it is roughly **half** the unconditional yield: `E[y|win]/E[y] = 0.48` overall (0.57 for first-tercile entries, degrading to 0.40 late). For a first-tercile entrant the naive break-even belief of 0.337 corrects to 0.473, an understatement of 13.5 percentage points. The correction shrinks as the opposing flow becomes noise (for pure-noise flow it vanishes), and §10's operators should treat it as a floor on required edge, not a refinement. The `L = 1` corollary above survives as the zero-drift, uninformed-flow benchmark, which is what "textbook" was doing in that sentence.

---

8. The Separation Principle

The primary layer settles. No payout depends on the final pool ratio, so end-of-life manipulation of that ratio has nothing to grab: you cannot expropriate settled claims by trading against them. Manipulating the displayed ratio late is *strictly more expensive* here than in a classic parimutuel or a book, because the manipulator's stake earns nothing even when it is correct.

The secondary layer prices. A position is four numbers, so it trades cleanly. An informed trader late in a market does not inject into the pool, since P4 makes that pointless, but buys positions from holders. That reroutes late information through prices paid *to* risk-bearers rather than taken from them.

We are obliged to say how strong this is, because the rest of the paper leans on it, and the answer now has two halves. The structural worry stands: there is no guaranteed source of uninformed selling late in a market, and a standing cash-out vault is a designated adverse-selection sink that will widen or decline exactly when information is most valuable; B.4(iii) is the pricing question. But the *capacity* of a secondary price to carry the late forecast is no longer a confession, it is a measured bound: §13.2 adds a stylized dealer to the behavioural study (every late-arriving informed agent prints a transaction at their belief shaded by half the spread), and the resulting composite forecast's final-phase Brier is 0.054-0.057 at every spread from 2% to 20% (0.075-0.077 in the sparse arm), against 0.174 for the mechanism's own pool ratio and 0.148 for an unlocked classic pool. The model is deliberately generous (no inventory risk, no quote withdrawal, every informed arrival prints), so read it as an upper bound on what the layer can deliver and a lower bound on nothing; what it establishes is that the late information exists and a thin transaction layer suffices to surface it, so the open question is the adverse-selection economics of who quotes, not whether a quote would help. For short-duration recurring markets no position market can form inside the round, and §15 resolves the deployment tension that creates by specifying the venue-side RFQ cash-out as the degenerate secondary layer there. §11's collateral and index constructions remain gated on this section maturing from a bound into a design.

---

9. Incentives

For a stake s on o with belief p, $EV = p\cdot y\cdot s - (1-p)\cdot s$, with y from §7 and the §7.1 corrections applied (the honest form uses $E[y|win]$). Three observations.

Flow exposure. That is what being the counterparty means, made explicit. Any mechanism paying counterparty-liquidity rewards must fund them from flow that has not arrived yet; that dependence is a property of the problem, not of a particular design. The difference is where it lives: here it is the second term of the payout, where an agent can model it, and §7.1 quantifies what the modelling must include.

Timing. For a trader committed to a side and a size, entering earlier captures every intervening allocation, all non-negative. This is close to tautological and we label it as such. It is *not* a claim that waiting is irrational: waiting buys information, and a trader with market impact facing a concave $V(s)$ will generally split and delay, so strategic withholding is generically profitable relative to immediate full entry, not a knife-edge case; B.4(v) states what withholding can and cannot touch. A metering whale also shapes the public q-path other entrants condition on, an aggregation channel this paper does not model. The unconditional timing problem is open (Appendix B).

Crowding out. Vesting is pro-rata by principal, so a large enough first vintage absorbs most future flow. Measured under the shipped mechanism (the seed-size sweep in `vpm-study.mjs`): growing the seed from \$50 to \$5,000 a leg, about three times the organic pool, compresses ordinary early winners from 1.70× to 1.09×, while the seeder sits at its floor. Under a binding κ the same concentration appears within a side as a race for headroom (§4.3). Three forces push back: vintage 0 is contestable in principle, crowded-out traders can buy the seeder's positions rather than disappearing, and $\lambda < 1$ keeps small entrants' returns alive. None is an identity, and a venue whose creators seed heavily will look like a market-maker venue with floors. We regard that as an acceptable degenerate case: the market maker is permissionless, floored, and earns no information rent, but it is not the time-priced ideal, and it is the mechanism's most likely real-world drift. The drift has a second cost, named in P5: a participant who dominates both books can print wash volume at near-zero mechanism cost, so in exactly the regime the drift produces, headline volume stops being evidence of anything; §12 draws the metric consequence.

The reserved vintage. Vintage 0 cannot be competed for *within* a market by construction: that is what makes P6 a theorem, so the creation floor is a rent the mechanism grants rather than one the market prices. Appendix B.3 states the tension exactly: a contested seeding race and an unconditional creation floor are mutually exclusive, and this design picks the floor. Three reasons, none decisive alone: the seed doubles as the resolution bond §12 wants; a floored creation position is what makes market-making the long tail a computable business (§10); and competition does not disappear: it moves *across* markets, where creators compete for flow with the quality and resolution record of what they create. The cost is equally plain: in our crowd model the reserved seed settles positive in 100% of 40,000 markets (a zero-event rate whose 95% upper bound is 0.0075%), and a rational creator population will treat that as a subsidy schedule. A venue that prefers contested seeding can open vintage 0 and accept a conditional floor; that is a legitimate point in the design space, and P6 tells it exactly what it gives up.

Time-priority ordering. The funding ledger deserves to be stated exactly, because a loose statement of it concedes a structure the mechanism does not have. A position's return above principal is funded *exclusively by subsequent opposing stakes*: event-contingent counterparties who, in the other branch, collect the position's own principal. Later **same-side** entrants, the population a pyramid harms, fund earlier same-side entrants with nothing: their accrued claims are untouchable (P3) and the last same-side cohort is paid exactly its principal (P4), disclosed as the design goal rather than discovered at the exit. Nobody's payout improves by recruiting anyone; self-referral is measured strictly costly outside the book-dominance regime P5 names, and the event contingency is exogenous. The honest analogy is a bookmaker's book that pays its counterparties by arrival time, not a scheme that pays its earlier adherents with its later ones. What remains true, and what a venue must own, is that expected seeding returns are a claim on future two-sided flow, so the venue's marketing incentive is to recruit counterparties for standing positions; a venue's headline volume is a poor trust signal when the best-rewarded strategy is being early with capital (and worthless in the self-dealing regime of P5), so §12's track records weight distinct, bonded, non-self-vested counterparties rather than handle.

---

10. Agents

R7 asked who arrives first at the ten-thousandth market. The mechanism's answer is that seeding is a *legible job*: a creation floor (P6), closed-form position value (§6), and economics that depend on forecastable quantities, flow volume and balance, rather than on out-quoting a professional. Machine-payable rails make the marginal cost of discovering and reaching a new market near zero; the cost of *seeding* it is the opportunity cost of floored capital plus resolution risk, which is what the position is compensated for.

Three limits on that claim, two old and one measured. The competition for vintage 0 is a conjecture about behaviour, not a mechanism property. Whether the race clears at agent speed is an empirical question this paper cannot answer from simulation. And the entry rule's profitability is exactly as good as its flow forecast, with teeth: in the sparse long-tail regime this section is about (§13.2's low-arrival arm, roughly 27 stakes a market), agents using the naive growth estimate $\hat{L} = \ln(T/t)$ realize **−11.5%** per entry post-fee while the same beliefs under classic rules earn +11.2%; the naive estimator wildly overstates growth precisely where flow is thin, and §7.1's winner's-curse correction is not optional there. Operators should form $\hat{L}$ from realized per-class growth curves (the fixed-point iteration of §13.2 is the template, and it also shows the correction converging rather than unravelling), never from promotional volume, which P5's wash analysis shows can be manufactured at near-zero cost by a book-dominant party. Under free entry the seeding rent should compete away toward the cost of floored capital, which is the design working as intended: the profit competes away, the seeded market remains.

---

11. Positions as Primitives

A position is a deterministic transferable claim with a monotone win-branch floor, which the classic parimutuel never had. It maps onto semi-fungible tokens (one class per market × outcome × vintage), and same-vintage positions genuinely are fungible under §6 because they share an accumulator snapshot. Its win-case payout splits into a **floor tranche** (principal plus accrued claims: a digital option with a known payout) and a **flow tranche** (whatever vests later), letting a holder keep conviction and sell activity, or the reverse.

The pricing *inputs* require no new trust: position state is recomputable from the public log. The *instruments* do: tokenization adds contract risk, collateralization adds dependence on whatever probability feed marks the position, and the flow tranche adds counterparty risk on an unsettled claim. This section describes what the primitive makes *representable*; it is not a proposal to offer instruments. And all of it is gated on §8 maturing from a measured bound into a working layer. We are not claiming yield instruments here.

---

12. Resolution

Liquidity is half of permissionless; the other half is who says what happened. We propose nothing new mechanically: staked resolution bonds with slashing, escalation games, and dispute juries are the deployed designs of Augur, UMA and Kleros, and have been in production for years. What this section asks of a venue is operational: every creator carries a public history (markets resolved, volume settled, disputes, time-to-resolution) surfaced via API so agents can price resolution risk before staking; unresolved markets auto-refund; a winning outcome with no backers voids. Because P5 shows a book-dominant party can print volume at near-zero cost, **track-record metrics must exclude self-vested flow**, flow whose vesting lands on positions of the same funding cluster, which is identifiable from the public log, and weight bonded, aged, distinct counterparties instead; distinct-counterparty counting alone is sybil-purchasable and volume alone is worse.

This mechanism changes the resolver's incentive landscape in three specific ways, two of them for the worse, and naming them is the point of this section.

Delay farming. Accrued claims are monotone, so a resolver holding a winning position has a weakly dominant incentive to *delay*: every extra dollar of opposing flow vests to them at no risk. That is new relative to the classic parimutuel, where delay dilutes the winner. The fix is mechanical for the markets §15 puts first: **fix the resolution timestamp at creation, freeze the accumulator there, and let the resolution transaction land whenever it lands**; entries with a vintage after the freeze are refused in full, which is the market ending at its declared time, not a lock window (P4's corollary is about markets *before* their declared end). Then latency has zero payoff impact. The scope is real: this fix is exact only for known-schedule markets. For event-driven markets, whose resolution time is unknowable at creation, the freeze must trigger at the earliest of a fixed outside date and the first bonded resolution attestation, so the resolver can accelerate the freeze but never extend it; between attestation and finality the accumulator is frozen and post-freeze entries are refused. An implementation that instead accepts post-freeze entries has broken P1, because a frozen accumulator orphans their vesting. The attestation trigger is itself a lever and must be priced: any bonded party can halt entry by attesting, so a defeated attestation must un-freeze the market and slash the attestor's bond, sized at least to the blocking value of the halt (§4.3 gives the shape), and the residual honesty note is that event-driven markets do close entry for the length of a challenge window, a bounded reintroduction of the lock that P4's corollary should be read as excluding.

Stall-into-void. The auto-refund safety net is itself an attack surface, and the two must be designed together. A resolver on the losing side of a large vested payout has a strictly profitable strategy of never resolving: at the refund deadline they recover their full losing principal and every counterparty's accrued claim evaporates. The vested structure makes this sharper than in a classic pool, because accrued claims are public, so the staller knows exactly what refusing to resolve saves them, and can attempt to extort large vested winners inside the window. Three rules close it, and they compose without contradiction. First, **a timeout forfeits the resolver's seed**: the seed legs' principal, a branch-independent amount, funds the dispute layer, so the stall costs 2× seed and rises with whatever escalation has attached; the legs are extinguished as positions, and anything they would have collected beyond principal in a later resolution redistributes pro-rata to that branch's other winners, so the premium below never accrues to the seed or to the dispute layer as such (a fallback resolver who happens to hold positions in the declared branch still inherits a pro-rata slice like any other branch winner, the generic resolver-conflict the escalation game must price). Second, **fallback resolution is permissionless** (anyone may post a bond to resolve after the timeout, with the standard escalation game), so a void is the outcome of genuine ambiguity rather than a strategy; and because the timeout alone cannot distinguish a stall from honest ambiguity, the venue should provide a **bonded void declaration** through the same escalation path, letting an honest creator of a genuinely unresolvable market resolve to void without forfeiture, so the forfeiture prices only silence. Third, Rule 1's "irrevocable" is, as §4.1 now states, conditional on the market resolving: in the void branch every non-seed position refunds at its accepted principal, exactly, which conserves the pool and is pinned in the suite (case V4); the resolver-seed's refund is what the timeout redirects.

The direction premium. A creator-resolver holds the seed on every outcome, so whichever way they resolve, they hold the winning leg, and the legs do not pay alike: the winning seed leg farms the *losing* flow, so the seed always pays more in the branch with more opposing volume, typically the minority outcome. The gap, `seed·(A_w(T) − A_w'(T))` across the two declarations, is a closed-form, publicly computable premium the mechanism itself attaches to resolving toward one side, and it is steeper here than in a classic pool because vintage 0 captures an outsized share of losing flow. Venues should display each market's resolution premium live (it is two accumulator reads), score creators on resolution direction against the pool-implied favourite over their history, and make vintage 0 slashable on adverse escalation rulings, which converts the premium from a bribe into a bond at risk. A creator's seed is a better starting rung for the escalation path than reputation alone; it is only a rung. The bond is fixed at creation while what it secures grows with the market (the behavioural study's $100 seed against its $2,254 mean accepted pool is 4.4%, and the ratio only falls as a market succeeds), so **escalation bonds must scale with the disputed declared-branch payout**, a quantity §6 makes O(1)-computable per position, which is a concrete advantage over classic pools worth claiming: the dispute layer can price exactly what is at stake.

---

13. Simulation

Two studies. The first (§13.1) settles a **fixed flow sequence** under both rule sets, isolating settlement; the second (§13.2) lets the crowd respond to the rules it faces, measuring participation and forecast quality. Threats to validity are §13.3. Every number regenerates from the published scripts, and every table names its experiment; the main study runs the actual mechanism (Rules 1 *and* 2, $\kappa = 9$, creator-seeded vintage 0) over 2,000 fifteen-minute binary markets per seed, **20 seeds, reported as mean ± 95% CI** (Student t, 19 degrees of freedom, sample standard deviation): a 900-step walk, an informed-plus-noise crowd (75% noisy-informed, 25% noise; ~180 stakes; log-normal sizes, ~$10 median). Integer-unit accounting throughout; conservation asserted per market.

```
node docs/whitepaper/sim/vpm-study.mjs --markets 2000 --seeds 20    # §13.1: capacity rule, CIs, fees, κ sweep, n=3
node docs/whitepaper/sim/vpm-behavior.mjs                           # §13.2: responsive agents, Brier, dealer, decomposition
node docs/whitepaper/sim/vpm-behavior.mjs --sparse                  # §13.2: long-tail arm (arrival 0.03)
node docs/whitepaper/sim/vpm-behavior.mjs --grid                    # §13.2: sensitivity grid
node docs/whitepaper/sim/vpm-behavior.mjs --fixedpoint              # B.4(i): L-hat iterated to self-consistency
node docs/whitepaper/sim/vpm-yield.mjs                              # §7: discretization, drift error, winner's curse
node docs/whitepaper/sim/vpm-estimators.mjs                         # §13.3: primary-layer late-price candidates (all negative)
node docs/whitepaper/sim/vpm-figures.mjs --markets 2000 --seeds 20  # figures + the λ table
node docs/whitepaper/sim/vpm-capacity.mjs                           # properties P1–P7, cases P8–P11, V1–V4, P7a–P7c', P9b, P9c, P10b, P5w, A4
node docs/whitepaper/sim/vpm-accumulator.mjs                        # O(1) equivalence, residue bounds, seed voiding
node docs/whitepaper/sim/vpm-conformance.mjs                        # 106 fixture vectors, settler-injection harness
```

13.1 The main study

Payout by entry time (median multiple on winning positions, by decile; mean of per-seed medians, 20 seeds; Figure 1). Every CI in the mechanism column is at most ±0.019:

decile	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
classic	1.40	1.37	1.36	1.36	1.36	1.37	1.38	1.39	1.40	1.42
this mechanism, κ = 9	**2.54**	1.76	1.45	1.29	1.20	1.14	1.09	1.06	1.03	**1.004**

Figure 1. Median winning multiple by entry-time decile: the capacity mechanism against the classic parimutuel, 20 seeds by 2,000 markets.

The classic row is the late free-ride, drawn. Last-decile entrants capture 13.20% ± 0.07 of the losing pool under classic rules and 0.28% here. Rule 2's partial fills refuse 0.10% ± 0.02 of gross stake at $\kappa = 9$: the cap exists for the adversarial case (§5, P5) and barely touches ordinary flow. In the single-seed κ sweep (500 markets, seed 42, a smaller sample of the same generator: its point estimate at $\kappa = 9$ is 0.05%, inside sampling variation of the 20-seed figure, and the statistic is dominated by rare near-saturated markets), $\kappa = 3$ refuses 10.1% with the decile pattern unchanged, and measured refusal is zero to two decimals from $\kappa = 12$ up. The same pattern holds at $n = 3$ outcomes (500 markets, seed 42, refusal 1.03%, conservation exact in all, deciles 4.50 → 1.01).

The buzzer snipe, completed. The sniper stakes 50% of the gross pool on a >85% leader at $t = 0.95T$, inserted at its timestamp (Appendix A.5 explains why appending instead would fake the number). Restricted to the triggers where the leader goes on to win, the sniper's multiple is **1.005 ± 0.00006** under the mechanism (Rule 2 fills 97.7% ± 0.1 of intended size) against **1.273 ± 0.003** classic, and early winners' payouts fall 0.19% ± 0.002 on average (4.9% worst across seeds) against classic's 13.64% ± 0.13 average and 64.4% worst, same experiment, same triggers. The win-branch numbers alone are an upper bound on the attack's value, so the study also counts every trigger: the leader wins 98.42% ± 0.13 of the time it qualifies, which makes the snipe's **unconditional expected value −1.05% ± 0.13 pre-fee here (−3.03% post-fee), against +25.27% ± 0.33 pre-fee (+22.76% post-fee) under classic rules.** That pair of signed numbers, not the lock rule, is what deletes the buzzer strategy; the residual 0.19% is early winners sharing the last 5% of opposing flow with a newcomer, which is P3 working as stated, not leaking.

Post-fee. At a 2% entry fee the last-decile multiple is 0.984, so rational late pool entry is dead *post-fee by an even wider margin*, which is the design working. The creation floor survives fees in these markets: worst case across 40,000 settled markets, **+8.6% after fees**, because the floor already contains vested flow by the time fees matter. The unconditional statement stays honest: in a market that attracts *zero* flow, the seed returns break-even minus fees, i.e. −2%.

The λ dial (Figure 2; capacity mechanism at every λ via the §4.5 composition, 20 seeds, sniper inserted by timestamp; λ = 0 differs from the true classic only by Rule 2's 0.10% refusal; the sniper's own multiple runs 1.273× to 1.005× linearly across the same rows, by the composition identity):

λ	last-decile multiple	early-winner dilution (mean)	seed straddle mean / min across seeds / % positive
0	1.416×	13.5%	−18.0% / −44.4% / 16.6%
0.25	1.315×	9.0%	+27.7% / −30.6% / 70.8%
0.5	1.213×	5.5%	+73.4% / −16.8% / 97.0%
0.75	1.112×	2.6%	+119.0% / −3.0% / >99.9%
1	**1.004×**	**0.2%**	+164.7% / **+10.9%** / **100%**

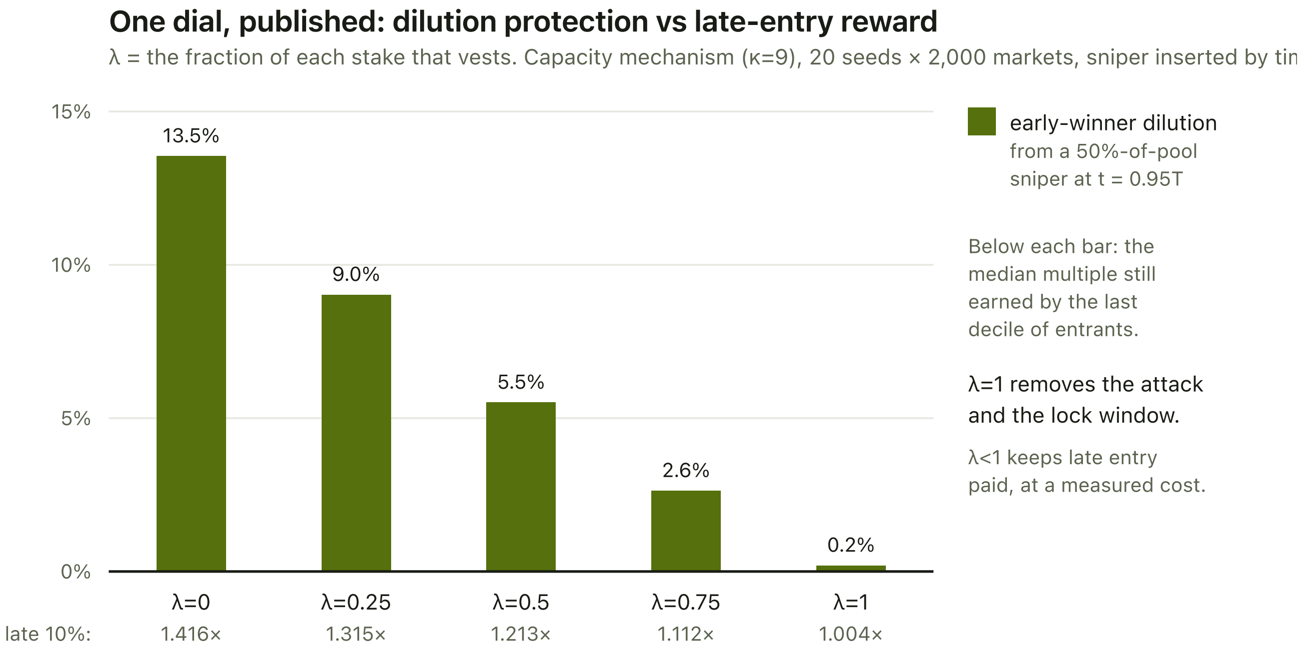

Figure 2. The λ dial: early-winner dilution from a 50%-of-pool sniper, and the multiple still paid to the last decile, per λ.

Seeding (creator seed, matched vintage 0, 20 seeds). Read the floor rows, not the mean:

	this mechanism	λ = 0.5 blend	classic
worst case, 40,000 markets, pre-fee	**+10.9%**	−16.8%	−46.8%
worst case, post-2%-fee	**+8.6%**	n/a	n/a
markets settling negative	**0 of 40,000** (95% upper bound 0.0075%)	3.0%	83.4%
mean	+164.7% ± 1.3	+73.4%	−18.0%

The mean is a property of our crowd model and would differ under yours. The floor is a property of the mechanism, and P6's unconditional form is exact: a market with zero subsequent flow returns the seed to the cent, minus fees. At λ = 0.5 the floor provably lapses, exactly as theory says it must, which is the more useful thing the middle column shows. Returns fall as the seed grows relative to the flow behind it while the floor does not move. **This is payment for cold-start risk-bearing, concentrated in whoever bears it first. It is not a yield, and quoting it as an APY misquotes us.**

13.2 What happens when the crowd responds

Settling a fixed flow sequence under two rule sets isolates settlement but cannot measure what the rules do to *participation* or to *forecast quality*: the two things a prediction-market paper must measure. So a second study lets each trader decide whether to enter under the mechanism they actually face, with the entry rules in the same units across arms (return-EV thresholds throughout): under classic rules, enter the better side when its return EV clears fee plus edge; under this mechanism, enter when `p̂·(1+ŷ) > 1 + fee + edge`, `p̂` the trader's private belief, with `ŷ` from §7 and remaining growth estimated as `L̂ = ln(T/t)`. A third arm is the deployed baseline, a classic pool **with a lock at 0.95T**, and a fourth (`vpm-cons`) repeats the mechanism with the conservative estimate `L̂/2`. Same belief stream, same arrival opportunities, same stake sizes; only the entry decision differs. 20 seeds × 2,000 markets, 2% fee, all four arms printed for every metric (`vpm-behavior.mjs`).

Volume. Share of volume by phase of market life, with the level, since the mechanism trades volume for time-placement:

arm	first third	middle third	67-95%	final 5%	total $/market
classic (no lock)	32.9%	33.3%	28.7%	5.1%	2,845
classic + lock	34.6%	35.1%	30.2%	0.1%	2,702
this mechanism	**43.9%**	34.7%	19.7%	**1.7%**	2,254
conservative L̂/2	48.5%	30.6%	18.8%	2.1%	1,748

The mechanism does with incentives what the lock does with a rule: late pool volume collapses (5.1% → 1.7%) without prohibiting anything, while pulling a third more volume into the market's opening phase, and the conservative-estimator arm strengthens rather than reverses the shift, so the migration is the incentive, not the estimator.

Forecast quality. Brier score of the pool ratio against the realized outcome, by phase, all four arms (lower is better):

arm	first third	middle	67-95%	final 5%
classic	0.239	0.202	0.168	**0.148**
classic + lock	0.239	0.202	0.168	0.150
this mechanism	0.239	0.201	0.171	**0.174**
conservative L̂/2	0.240	0.197	0.183	0.180

And the mean absolute gap to the generator's true probability, `|q − p|`, same arms, which is the harsher metric because it does not reward being lucky:

arm	first third	middle	67-95%	final 5%
classic	0.101	0.164	0.244	0.321
classic + lock	0.101	0.164	0.244	0.300
this mechanism	0.102	0.173	0.277	0.353
conservative $\hat{L}/2$	0.102	0.190	0.285	0.363

The sampling scheme matters and is checked: those pool-quality tables weight each arm's own transaction instants, so a fixed-clock re-measurement (one end-of-phase snapshot per market, identical instants across arms; printed by the same script) is the control, and it agrees: final-phase Brier 0.167 for the mechanism against 0.145 unlocked classic and 0.152 locked, so the late-phase comparison is not an artifact of arm-specific sampling. Three honest readings. First, the volume migration buys **no measurable early forecast gain**: the first-third Brier is identical across arms, and `|q − p|` is weakly worse under the mechanism in every phase; within this crowd model, paying for time-priority does not improve the price, it protects the payout. The instrument case for the mechanism is dilution protection, the deleted lock, and the seeding economics, not forecast quality, and this paper stops implying otherwise. Second, the late-phase cost against the *deployed* baseline is the fair headline: 0.174 against the lock arm's 0.150 (the unlocked ideal's 0.148 is the bound, not the practice). Third, the conservative arm degrades earlier (0.183 from midlife), which is the other half of its robustness and is printed, not footnoted.

Who gets paid. Realized per-entry PnL, settled under each arm's own mechanism post-fee, split by trader type, with the first-third winning multiple by type as the direct test of whether payment tracks information *within* a moment:

arm	informed PnL	noise PnL	spread	first-third win multiple, informed / noise
classic	+5.9%	−20.7%	26.6pp	1.58× / 1.63×
classic + lock	+5.4%	−19.2%	24.6pp	1.61× / 1.66×
this mechanism	+1.4%	−26.1%	27.5pp	2.10× / 2.22×
conservative $\hat{L}/2$	+5.7%	−28.4%	34.1pp	2.03× / 2.07×

Two things this table settles. Payment is information-blind *conditional on the moment*, exactly as P5 requires: an early noise winner's multiple is slightly **above** an early informed winner's (noise contrarians land on longshot sides more often). The informed-noise PnL differential is at least as wide as classic's (the spread column), and it arises through selection, informed traders choosing moments and sides better, which is the only channel any parimutuel has; but note the direction of the extra spread: under the naive estimator the noise side's deeper losses accrue to the seed and to fees, not to informed entrants, whose own return falls to +1.4%. And agents trusting the naive optimistic estimator overtrade to +1.4%: a caution §10 turns into the sparse-market warning, where the same rule goes to **−11.5%** while classic earns +11.2% (arrival 0.03, the long-tail regime; `--sparse`). Sensitivity (`--grid`; single seed, 800 markets, point estimates): the late Brier cost ranges +0.011 to +0.042 across the grid, increasing in informed share and decreasing in belief noise, and the mechanism's informed PnL falls as informed trading crowds (to roughly −5% at 90% informed), both directions a deployer should expect.

The stylized dealer (§8's bound). Adding a transaction layer in which each late informed arrival prints at its belief shaded by half a spread, the composite forecast's final-phase Brier is **0.054-0.057 for every spread from 2% to 20%** (0.075-0.077 in the sparse arm), against 0.174 (this mechanism's pool), 0.150

(lock), 0.148 (classic). The late information exists; the pool is simply no longer the place it shows up. §8 states what the model omits.

The favorite-longshot calibration, snapshotted at 0.95T: in the 0.2-0.3 pool-implied bin, outcomes realize at 0.02 under classic rules and 0.03 under the mechanism (binomial SE ±0.002-0.003); in the 0.7-0.8 bin, 0.99 against 0.97. The mechanism's displayed ratio is mildly *compressed toward the seed's 50/50 prior* relative to classic, a stale-price effect of rational late flow staying out, visible as thinner tail bins; it is not a longshot-subsidy tilt in realized frequencies, and the sparse arm shows the same compression more strongly. §14 restates the favorite-longshot paragraph on this corrected footing.

The growth estimate iterated to a fixed point (`--fixedpoint`; 5 seeds × 1,000 markets, mechanism arm only, so its iteration-0 baseline differs slightly from the four-arm table; simulation evidence on B.4(i), not a theorem). Replacing $\hat{L}$ = `ln(T/t)` with the realized mean growth curve of the previous iteration and repeating: the phase-volume profile converges in four iterations under a 0.3pp criterion (the growth curve's nodes are still moving about 7% per iteration at stop, which the script prints), not toward collapse but toward an interior point. Volume contracts about 20% ($2,250 → $1,812 a market), the first-third share *rises* to 50.4%, informed PnL improves from +1.3% to +4.6%, the final-phase Brier drifts from 0.176 to 0.184, and the fixed-point growth curve sits well below the naive one ($\hat{L}$ at t = 0.1T: 1.61 against 2.30). Within this crowd model, rational flow does not unravel to vintage 0: it thins, arrives earlier, and stabilizes, because the noise flow keeps `L` positive and the informed edge keeps entry worthwhile. The theorem remains open; the simulation now points at an interior fixed point rather than at nothing.

13.3 Threats to validity

What the studies above do not establish.

- **The secondary layer is bounded, not designed.** The dealer arm shows a thin transaction layer *would* carry the late forecast (0.054-0.077 against 0.174); it does not show who quotes it or at what adverse-selection cost (B.4(iii)). Backtesting the composite estimator against recorded live tapes remains the most valuable missing experiment.
- **The cheaper repair does not exist, which is why §8 is load-bearing.** Before accepting that the late price must come from a position market, we tested whether it can be recovered from primary-layer state alone, which is free: every position's vintage and stake are stored anyway (§6). Three candidate estimators, against the classic arm as a control under identical treatment (`vpm-estimators.mjs`, 20 seeds x 2000 markets). Recency-weighting the pool by vintage does improve the final-phase Brier (0.1736 to 0.1395 at weight $(t/T)^4$), but it improves the *classic* arm more (0.1478 to 0.0942), so the vpm-minus-classic gap widens from 0.0258 ± 0.0007 to 0.0453 ± 0.0014: a generic recency result owned by neither mechanism. Weighting toward early vintages halves the gap to 0.0126 ± 0.0006 but degrades both arms absolutely (vpm 0.1948), which is convergence, not repair. Inverting the entry rule into a revealed-belief bound is destroyed by noise entries, which reveal no bound and are indistinguishable from those that do: final-phase Brier 0.4700, far worse than doing nothing. Against the dealer arm's 0.054 on the same instants, no primary-layer construction we found is competitive. The negative result is the point: §8 is not a convenience.
- **The behavioural agents are near-myopic.** Entry rules are threshold rules; the fixed-point mode iterates the growth estimate to self-consistency, which is an equilibrium in the estimator only, not in strategies. The unravelling question is narrowed by simulation evidence, not closed.

- **One crowd model.** The grid sweeps informed share and belief noise, and the sparse arm covers the thin-market regime, but arrival process, stake sizing, and the belief-noise family are single choices; the seeding *mean* and every PnL level inherit them. The floor rows and the identity-based claims do not.
- **Two stated conventions.** The behavioural study charges the fee on top of the gross stake (a loser realizes −(1+fee)) while §13.1's post-fee restatement nets it from the multiple; aggregates agree and splits differ by under a point, and each output states its convention. Pool-quality metrics are transaction-instant-weighted, with the fixed-clock control printed alongside (§13.2).
- **Fifteen-minute binary markets.** Duration enters only through the growth path; carry (§14) is not simulated.

---

14. Limitations

Finite κ does not survive n-way markets (§4.3): acceptance couples across branches through the thinnest book, which chokes liveness in skewed markets and hands a small stake a blocking lever (P9). For `n ≥ 3` the failure is worse than throttling: one stake of `S(κ − n + 1)` renders the market permanently unenterable on every outcome, an absorbing state with no in-mechanism recovery, reachable at any time and costing $350 at the reference parameters of §4.3 (P9c). Venues must treat κ as a binary-market instrument and run n-way markets at large or unbounded κ, where the mechanism's soundness rests on the seed and scale-invariance instead. We note but do not develop the obvious repair, a matched all-outcome top-up obeying the vintage-0 rule, which would restore `c(κ − n + 1)` of headroom per leg: it is a mechanism extension rather than a policy choice, the shipped settler admits no second matched vintage, and its interaction with P6 and P3 is unanalyzed. Appendix B.5 states it as an open problem.

Late-stage pool odds. Beyond §7's two-sided band, unrewarded late entry means the ratio stops tracking probability near resolution. §13.2 puts the fair pair of numbers on it: with agents responding to the rules they face, the final-phase Brier is 0.174 under the mechanism against 0.150 under the *deployed* lock baseline and 0.148 under the unlocked ideal: the mechanism pays forecast quality in the pool's last act to buy dilution protection for everyone before it. The measured dealer bound says the information is recoverable off-pool (§8, §13.2); building the venue that recovers it is the open design problem, and until then the live late price is simply worse here.

Inter-block MEV. Vintage determines payout, so a party who sees a stake before inclusion (a sequencer, an RPC or API operator, a peered searcher) can take the opposing vintage one block earlier and capture that stake's vesting with little principal and little uncertainty. The capture is computable from Rule 1 directly: front-running a victim stake `c` into an opposing book of principal `P` with a stake of `f` captures `c·f/(P+f)` of contingent claim, so the marginal value of one block of foresight on one stake is ≈ `c·f/(P+f)`, increasing in the victim's size and the thinness of the book, and the front-runner's optimum trades that against `(1−p)·f` of event risk. The victim is indifferent; the loss falls on the honest early holders who would have shared that flow. Equally: delaying someone by one block is economically inert in a classic parimutuel and a direct transfer here, and under a binding κ the same foresight steals *acceptance* headroom, a second surface whose victim is not indifferent, they are refused. Mechanical mitigations exist and compose, each with a cost: commit-reveal entry (removes flow visibility, costs a round trip and a griefing bond design), threshold-encrypted mempools where available, and widening the vintage from one block to a k-block epoch, which is a dial trading time resolution against ordering value, the natural extension of the batching §4.4 already specifies. A venue on a centralized sequencer should treat flow confidentiality as a mechanism parameter, publish a non-extraction commitment covering ordering and flow visibility, and understand that a commitment is weaker than any of the mechanical options. We claim the narrow thing, and after §4.4's rationing rule it is now actually true in both vesting and acceptance: no intra-block ordering game. Between blocks the game is real and priced above.

Hedging. A party with real exposure wanting protection late faces +0% on a win and −100% on a loss. In a book or an LMSR they buy at 0.93 and are covered. This matters twice: it is a genuine loss of function, and hedging demand is a principal source of the uninformed flow §8 needs in order to exist. Hedging-relevant markets should run λ well below 1, and §4.5's trilemma row states the price of that prescription plainly: interior λ voids the creation floor and partially re-opens the lock window, so a hedging market is an operator-seeded or floor-waived market. That is the honest reason λ exists, and no λ delivers everything.

Carry. Capital is locked from entry to resolution. At 4% and six months the 2% hurdle is the same order as the entire observed late-tercile win-conditional yield (median 2.7%, `vpm-yield.mjs`) and exceeds the median final-decile winner's post-fee return, and it applies to the seed too, which turns the nominal floor into a real loss on slow markets. λ = 1 is a short-horizon mechanism.

Favorite-longshot bias. Since `y ∝ (1−q)/q`, the mechanism pays more per winning dollar for backing the minority side, and that payment is information-blind: in any crowd where contrarians skew noise, the minority premium pays noise more per winning dollar than information (§13.2 measures exactly this: early noise winners at 2.22× against informed 2.10×). If the bias is driven by probability misperception (Snowberg & Wolfers 2010), the misperception is untouched and a mechanical inducement is added pointing the same way. What the corrected calibration shows, however, is not a realized-frequency tilt: at 0.95T the mechanism's bins sit within one to two points of classic's (0.03 vs 0.02 in the 0.2-0.3 bin), and the visible effect is *compression*, thinner tail bins and a displayed ratio dragged toward the seed's uninformative 50/50 prior as rational late flow stays out. The displayed-price distortion `|q − p|` does worsen in every phase (§13.2), and that, not a longshot tilt, is the measured price anomaly; the venue accepts it knowingly because the pool ratio was never the published probability (§7, §8).

Cascades and manufactured growth. The mechanism penalizes joining the crowded side, which is a brake. But §7's two-sided band means that early in a growing market a trader with a signal weaker than the pool's still finds entry profitable: the cascade-enabling condition. Worse, entry is a strategic complement in expected growth: entry EV rises in `L̂` and realized `L` rises in entry, so the same market can rationally clear as a frenzy or a ghost town on volume expectations alone, a coordination surface the classic parimutuel does not have (an early classic entrant gains nothing from volume per se). The pool's first stake is a creator's seed carrying no information at all; a venue that features a market is transferring expected value toward its earliest vintages and should disclose that; and agents must anchor `L̂` on realized base rates, not on observed early growth, which P5's wash analysis shows is manufacturable at near-zero cost by a book-dominant party.

No equilibrium theorem. The unravelling worry is that if no rational flow arrives after some `t*`, then `L` just before `t*` is generated by noise alone, shrinking `y` and pushing the horizon earlier. Two things temper the interior-fixed-point-or-collapse dichotomy. The horizon is side-dependent: `y` is `((1−q)/q)·L` on one side and `(q/(1−q))·L` on the other, an 81-fold gap at q = 0.9, so favourite-side rational flow dies far earlier than longshot-side flow and the object of study is a pair of cutoffs `t*(side, q)`, not a scalar. And abstention is self-limiting: while rational flow stays out, arriving information accumulates as a widening `|p − q|` edge that eventually clears any entry threshold. §13.2's fixed-point iteration is consistent with both: volume thins and shifts early but stabilizes interior, in that crowd model. The theorem, side-dependent cutoffs in a game whose state is (staleness edge, remaining growth), stays open and is the sharpest question we can hand a theorist (B.4).

Limits of verifiability. §6 removes *pricing* trust. It does not remove trust in the completeness and ordering of the trade log: a property of the chain and, on a centralized-sequencer L2, of the sequencer, nor in resolution, nor that the venue applied the published κ, λ, fee schedule, minimum stake, and residue owner, which should be committed on-chain per market, nor that an implementation honors the settlement-path MUSTs of §6 (transfer-exact asset, pull claims, state-before-transfer), which a JavaScript suite cannot execute for it. A paper that claimed otherwise would be doing what it criticizes.

---

15. Deployment

Nothing in this specification requires a particular venue, chain, or frontend. What the mechanism's own risk structure does dictate is an order of adoption, least-recoverable risk last:

1. **Paper first.** Flow-vesting settlement in a paper-money twin, behind the venue's existing payout authority.
2. **The secondary layer where it can exist.** For markets long enough for a position market to form, ship transferable positions and the cash-out affordance before real money, because §8 is a precondition, not a follow-up.
3. **Real money where P4 bites, with the degenerate secondary layer built in.** Short-duration recurring markets get $\lambda = 1$ (deleting the lock window) *plus a venue-side RFQ cash-out at a published spread* as the in-round secondary layer, since no position market can form inside a fifteen-minute round (§8). The vault's adverse-selection pricing is B.4(iii), it is the acknowledged cost of this step, and its quotes are the composite's late price. One requirement is structural rather than economic: the vault's quote must never be a mechanical function of the current pool ratio, or the manipulation §8 disarms comes back through the vault at P5's near-zero printing cost for a book-dominant party; it must price from the venue's own model over wash-excluded flow (§12's filter) and off-market inputs, which makes manipulation-robustness the second half of B.4(iii)'s pricing problem. A venue unwilling to run the vault should run these markets at $\lambda < 1$ and accept §4.5's trade instead.
4. **Creator-seeded permissionless creation**, with the reserved vintage 0 as the first rung of the resolution bond and §12's timeout-forfeiture and scaling-bond rules in place from day one.

Chain requirements. The mechanism is chain-agnostic: the settlement identity is one accumulator update per outcome, and nothing in it prefers one execution environment to another. It has exactly one chain sensitivity worth naming: vintages batch per block, so block time *is* the mechanism's time resolution, and the ordering games removed at the intra-block level reappear between blocks (§14 prices them; k-block epochs widen the batch where a venue wants to). At two seconds they are small; at twelve they are not. The rest of §10's assumptions are assumptions about rails rather than about consensus: settlement in a native stablecoin, per-transaction attribution, and a machine-payable request standard (x402 is the deployed one) alongside agent-readable market discovery. Any chain that offers short blocks and those rails will carry this. The specification belongs to whoever wants it.

---

16. Conclusion

Information has a value curve: worth the most when few people have it, worth nothing when everyone does. A market that pays a flat multiple to everyone who was right, whenever they arrived, is mismeasuring the thing it exists to price. The rule in this paper, that losing flow vests to the opposing book at the moment it arrives and is accepted only as far as that book can cover it, produces a payout schedule with the shape of that curve, and produces it as arithmetic rather than as a promise.

We close on the two honesty notes the measurements force. What the arithmetic guarantees is payment for *early risk-bearing*; it pays information only through the hypothesis that informed capital arrives before the crowd, and in our own simulations the payment shape moves volume earlier without making the early price measurably better. The measured case for the mechanism is the other column: exact dilution protection, a deleted lock window with the buzzer strategy at negative expected value, a floored and computable seeding job for the long tail, and a late price that must, and demonstrably can, come from a layer above the pool. And none of the novelty is the gradient itself: Pennock had the essential half in 2004. The honest description of this work is that it makes entry-time pricing an accounting identity, gives it a constant-time form that runs on-chain, derives what an entrant is priced at and what that price omits, supplies a dial to the classic mechanism with its trilemma stated, and publishes the cases where it fails.

Prediction markets number in the low thousands because each one must be worth a market maker's attention or an operator's subsidy. A market settled by arithmetic and seeded by its own creator carries no such per-market cost. We do not know what the ceiling is. We know it stops being the number of market makers.

FAQ

Why would anyone bet late? Into the pool, they shouldn't, and that is the design: the buzzer strategy's unconditional EV is measured negative here and strongly positive under classic rules (§13.1). Late information should enter through the secondary layer, where it pays risk-bearers instead of taking from them; §13.2's dealer bound shows a thin layer would carry it, and §15 says what stands in for it in short rounds.

Doesn't that destroy the price signal? It moves it, at a real, measured cost: final-phase Brier 0.174 against the deployed lock baseline's 0.150 (§13.2). Pool ratios are informative early and stale late; §7 shows they are not even a probability in fast-growing markets. If you need live late pool odds more than dilution protection, that is what $\lambda < 1$ is for, at the price §4.5 states.

Is this a Ponzi? No, and the distinction is structural, not rhetorical (§9). Returns above principal are funded exclusively by *opposing*, event-contingent stakes, counterparties who collect your principal in the other branch, never by later members of your own side, whose entry cannot touch your claim (P3) and whose own worst case is exact principal back (P4). Payment is contingent on an exogenous verifiable event; conservation is exact with no operator skim; no return is promised anywhere and the marketing surface is a conditional floor; and no participant's payout improves by recruiting anyone, with self-referral measured strictly costly outside the disclosed book-dominance regime (P5). The honest analogy is a bookmaker's book that prices counterparties by arrival time. What is true, and disclosed, is that early seeding returns depend on future two-sided flow arriving (§9).

Isn't this just a parimutuel with extra steps? It is a parimutuel with one moment moved: the losing pool is assigned at entry rather than at resolution, plus a matching constraint that keeps that assignment sound.

Who puts up the first dollar? A venue will not seed thousands of markets, and a creator posting a question will not stake it. Nobody's first arrival is a bet anywhere — an order book holds it as an unmatched quote, an AMM meets it with a curve someone funded, a classic pool escrows it — so the real question, asked of every structure, is who funds the first counterparty and at what expected cost. §1.3 prices the three answers, and this mechanism's is the only floored one: worst case nominal recovery in every branch (P6), which turns programmatic venue seeding into a revolving float rather than a subsidy budget. The roles also separate: a market can be posted as an unseeded listing that opens when anyone — including a first bettor tilting the seed to the κ bound — posts vintage 0. A market nobody will seed at any tilt has no believer willing to stand on any side of it, and the mechanism declines to pretend otherwise (P11).

What stops a whale from seeding everything? Nothing. §9 measures it: a seed three times the organic pool takes ordinary early winners from 1.70× to 1.09×. We flag it as the most likely real-world drift, and §12 requires trust metrics that its wash-printing corollary cannot game.

Is there MEV? Not within a block, in either vesting or acceptance (§4.4). Between blocks, yes: §14 gives the capture formula and the mechanical mitigations, and says plainly that this mechanism makes advance sight of order flow directly monetizable, a new risk class relative to a classic parimutuel.

Can I implement it? Yes. CC BY 4.0 on the text, MIT on the code, any chain, no permission. We ask that implementations describing themselves as a Vested Parimutuel pass the published conformance suite and state the suite version and vectors hash they passed; the Artifacts note points at the files, and the README served beside the suite documents the interface.

How do I break it? Please do. Appendix A is the set of attacks that already killed earlier candidate designs, every one pinned in the shipped code; structured adversarial review has already broken two of this paper's own formal claims, both restated above with their counterexamples pinned in the suite. A case the suite does not cover is the interesting one; the Artifacts note says where to send it.

Appendix A. Design Alternatives, and How They Fail

The mechanism in §4 is the survivor of several designs that look reasonable and do not work. Each failure below is invisible under ordinary flow and appears only under adversarial input, which is why each one is pinned in the shipped code rather than left as prose: A.1's scale-invariance fix, A.3, A.4, A.5 and A.6's seed-validity pitfall in the conformance suite; A.2's grabs and A.6's float-arithmetic trap in the historical record (`vpm-edge-tests.mjs`).

A.1. Uncapped vesting. Rule 1 without Rule 2: distribute each stake pro-rata over the opposing book, with no ceiling on what the book may absorb. The first position on an empty book then receives 100% of every opposing inflow until a second position joins, so return on capital diverges as that position shrinks. In the design-phase measurements a seeding straddle returned +155% at $50 a leg, +1,104% at $1, and +82,037% at one cent; the shipped suite pins the divergence as P5's 200,001× single-probe comparison. The profit-maximizing strategy is to post dust on every new market and toll the organic flow while supplying no liquidity, which is the opposite of what a cold-start mechanism is for. Rule 2 fixes it at the book level, so distribution stays pro-rata: return on capital becomes scale-invariant (dust earns what any same-moment capital earns, never more) and is bounded by $1 + \kappa(1 + \ln g)$.

A.2. A metered bootstrap bounty. Hold unmatchable flow in a bucket and release it to later entrants at up to $\rho\times$ their stake, decaying as $\rho(t) = \rho(1 - t/T)$. Three defects. It is strictly dominated by A.1's dust position, so it never becomes the best strategy it was designed to be. Its residue must be refunded at resolution, which means a stake placed while an opposing book was empty is not fully at risk, and late entry stops being unconditionally negative-EV. And its decay is indexed to clock time while the risk it prices is informational: in a market that resolves informationally early, a grab at $t = T/30$ still returns 3.90× against the 4.00× the decay was introduced to remove (pinned in `vpm-edge-tests.mjs`). A capacity rule needs no schedule and therefore has no schedule to mis-index.

A.3. An unreserved seeding vintage. If the creator's legs are not alone in vintage 0, the creation floor (P6) does not exist. An implementation that admits a third party into vintage 0 takes the floor to −50% in half of all branches; the conformant entrypoint (a same-block stake is assigned vintage 1) preserves it. Both behaviours are pinned against the current mechanism. The floor is a theorem about atomicity, not about capacity: it needs the creator to be first on every outcome, which is what the reserved vintage buys.

A.4. Per-position capacity caps. Capping each position at $\kappa \cdot s_i$ also closes A.1, and it is the more obvious fix. It is the wrong one: the cap binds in 87% of test settlements at $\kappa = 3$ (pinned, with the breakage, as suite case A4), and whenever it binds the allocation is no longer pro-rata, which breaks the O(1) accumulator of §6 and with it the property that a payout is a function of two scalars. Constraining acceptance instead of distribution closes the same attack and keeps the constant-time form exact.

A.5. A measurement pitfall. A settlement study that appends a late entrant to the end of its event array rather than inserting it at its timestamp measures a buzzer entry, not a late one, and will report exactly 1.000× and exactly 0.0% dilution. Inserted honestly at $t = 0.95T$ the same experiment gives 1.005× and 0.19%. The difference is small and the comparison against classic rules survives it, but the first pair of numbers is too clean and should be disbelieved on sight. (The figure generator asserts the honest insertion and refuses to run an appended sniper, so this warning is an executable check and not only prose.)

A.6. Implementation pitfalls. Four that a conformant implementation must avoid, each caught by the suite or stated as a MUST in §6. Compute capacity or the accumulator in floating point and the settlement disagrees with the published integer rule on grid points (the reference computes $\kappa \cdot s$, headroom, and rationing in integers; the float-vs-integer disagreement is pinned in `vpm-edge-tests.mjs`, where the float form disagrees with the published integer rule on 438 of 2,700 grid points). Hard-code two outcomes and every $n \geq 2$ claim in the specification goes untested. Accept a seed that leaves an outcome unbacked, including a dust leg floored to zero by integer allocation, and the market is silently un-enterable forever, because every book's capacity is zero: it must void at creation instead (P11). And pay winners in a resolution-time loop instead of pull-claims and one blacklisted address stalls every claim behind it (§6's MUSTs).

Appendix B. Toward an Equilibrium Analysis

What can be said in a line, what we conjecture, what simulation now indicates, and what is open.

B.1 Early entry weakly dominates conditional on entering with a fixed side and size, ignoring the option value of information acquired by waiting. Nearly tautological; stated for completeness.

B.2 Late entry is strictly dominated: after the last opposing inflow, $EV = -(1-p)s < 0$, unconditionally under Rule 2. Rational pool participation therefore has an endogenous horizon.

B.3 Seeding rents dissipate under free entry toward the cost of floored capital plus flow uncertainty. Note the tension with P6: a contested vintage 0 is precisely the configuration in which the reserved-vintage hypothesis fails, so a venue must choose between an unconditional creation floor and a contested seeding race. That trade-off is real and we have not resolved it.

B.4. Open. (i) The unravelling question, now with structure: the rational-flow horizon is side-dependent (the two sides' yields differ by $((1-q)/q)^2$, 81-fold at q = 0.9) and abstention is self-limiting (staleness accumulates as entry-relaxing edge), so the object is a pair of cutoff paths $t^*(\text{side}, q)$ in a game whose state is (staleness edge, remaining growth). §13.2's fixed-point iteration, the estimator equilibrium, contracts volume about 20% and stabilizes interior rather than collapsing to vintage 0 in our crowd model; whether strategic agents do the same is the open theorem, and lumpy entry (positions firing when accumulated edge crosses the threshold) is a plausible third shape the interior-or-collapse dichotomy misses. (ii) Bound the divergence between pool ratio and consensus probability as a function of flow predictability and λ. §7's band gives the static envelope, $|\text{logit}(q) - \text{logit}(p)| \le \ln L$ inside the two-sided region; which edge binds depends on the flow model, and the drift measurements of §7.1 are the empirical handle. (iii) Adverse-selection pricing for the cash-out vault, and whether a competitive RFQ set converges; the dealer bound of §13.2 prices the value of solving it. (iv) Optimal λ per market class given §4.5's trilemma. (v) Strategic flow withholding, stated exactly: withholding cannot revise any accrued claim (P3) and forfeits the withholder's own vesting, so its direct cost is real; but §9 concedes that a trader with impact will generally split and delay, so withholding is *generically* individually rational, and its unmodelled harm is informational, a metering whale controls the public q-path that every other entrant's rule conditions on. The open problem is the dynamic-execution equilibrium with that signalling channel included.

B.5 The n-way freeze of §4.3 has an obvious candidate repair and no analysis: a **matched all-outcome top-up**, obeying the vintage-0 rule (legs counterparties to one another, each accepted only against the others), restores $c(\kappa - n + 1)$ of headroom per book while consuming $(n - 1)c$, so it strictly unfreezes. Three questions block it. Whether a second matched vintage preserves P6, since the creation floor's proof relies on the creator owning the entire market at the *unique* reserved vintage. Whether it preserves P3 for holders who entered between the two matched vintages. And who may call it: a creator-only top-up is a dilution lever over standing positions, while a permissionless one is a new griefing surface, since the same construction that freezes a market is cheap for whoever wants it frozen. We state the arithmetic and decline the design.

References

- Abernethy, J., Chen, Y., & Wortman Vaughan, J. (2013). *Efficient Market Making via Convex Optimization*. ACM TEAC 1(2).
- Agrawal, S., Delage, E., Peters, M., Wang, Z., & Ye, Y. (2011). *A Unified Framework for Dynamic Prediction Market Design*. Operations Research 59(3).
- Ali, M. M. (1977). *Probability and Utility Estimates for Racetrack Bettors*. JPE 85(4).
- Angeris, G., & Chitra, T. (2020). *Improved Price Oracles: Constant Function Market Makers*. ACM AFT.

- Chen, Y., & Pennock, D. M. (2007). *A Utility Framework for Bounded-Loss Market Makers*. UAI.
- Hanson, R. (2003). *Combinatorial Information Market Design*. Information Systems Frontiers 5(1).
- Hanson, R. (2007). *Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation*. J. Prediction Markets 1(1).
- Lange, J., & Economides, N. (2005). *A Parimutuel Market Microstructure for Contingent Claims*. European Financial Management 11(1).
- Lesaege, C., Ast, F., & George, W. (2019). *Kleros: Short Paper v1.0.7*. kleros.io.
- Melee Markets (2026). *The Parimutuel Market Maker: A New Market Structure for Permissionless Prediction Markets*. Public litepaper; no stable public URL at the time of writing, so quotes and page numbers were verified against the August 2026 version, copy on file.
- Ottaviani, M., & Sørensen, P. N. (2008). *The Favorite-Longshot Bias: An Overview of the Main Explanations*. In Hausch & Ziemba, *Handbook of Sports and Lottery Markets*.
- Ottaviani, M., & Sørensen, P. N. (2010). *Noise, Information, and the Favorite-Longshot Bias in Parimutuel Betting*. AEJ: Microeconomics 2(1).
- Pennock, D. M. (2004). *A Dynamic Pari-Mutuel Market for Hedging, Wagering, and Information Aggregation*. ACM EC.
- Pennock, D. M., & Sami, R. (2007). *Computational Aspects of Prediction Markets*. In Nisan et al., *Algorithmic Game Theory*, Ch. 26.
- Peterson, J., & Krug, J. (2015). *Augur: A Decentralized, Open-Source Platform for Prediction Markets*. augur.net whitepaper.
- Plott, C. R., Wit, J., & Yang, W. C. (2003). *Parimutuel Betting Markets as Information Aggregation Devices*. Economic Theory 22(2).
- Snowberg, E., & Wolfers, J. (2010). *Explaining the Favorite-Longshot Bias: Is it Risk-Love or Misperceptions?* JPE 118(4).
- Thaler, R. H., & Ziemba, W. T. (1988). *Anomalies: Parimutuel Betting Markets: Racetracks and Lotteries*. JEP 2(2).
- UMA Project (2021). *UMA's Optimistic Oracle*. docs.uma.xyz.
- White, D., & Diamandis, T. (2024). *pm-AMM: A Uniform AMM for Prediction Markets*. Paradigm.
- White, D., Robinson, D., & Adams, H. (2021). *TWAMM*. Paradigm.

Artifacts and reproducibility. The canonical home of this paper is playhunch.xyz/vpm-whitepaper (this page, plus /vpm-whitepaper.md for machines and /vpm-whitepaper.pdf), and that page is where to report a break in the mechanism, the suite, or the paper's claims. A companion library of supporting articles — plain-language explanations, structure-by-structure comparisons, operator and implementer guides — is served at playhunch.xyz/vpm-whitepaper/articles (and as one document at /vpm-whitepaper/articles.md); it is explanatory only, quotes this paper's own figures, and specifies nothing: where the two disagree, this paper governs.

*Every number in §13 regenerates from the published scripts with the flags shown, and every script in `docs/whitepaper/sim/` is also served verbatim at `playhunch.xyz/vpm-whitepaper/sim/<file>` (including `LICENSE`, `README.md`, and the fixture vectors `vpm-vectors.json`). **The conformance suite for the mechanism this paper specifies is suite version 1.1.1: `vpm-capacity.mjs` (properties P1-P7 with their**

asymmetric-seed variants, cases P8-P11, V1-V4, P7a-P7c', P9b, P10b, P5w and A4) together with `vpm-accumulator.mjs` (exact O(1) equivalence including block vintages, residue bounds, seed-validity voiding) and the 106 fixture vectors in `vpm-vectors.json` (sha256 prefix `52b0fcdea345`) run through `vpm-conformance.mjs`, which accepts any implementation exporting the `settle(vector)` interface documented in the served `README.md` and the harness header. An implementation describing itself as a Vested Parimutuel should pass all of it and state the suite version and vectors hash it passed. `vpm-edge-tests.mjs` and `vpm-sim.mjs` preserve the *deleted* v2 mechanism (the ρ-decay bounty) and its failing cases (the late grab, the informationally-early grab, the residual-refund leak, the float-arithmetic trap) as a labelled historical record; the paper's λ table comes from `vpm-figures.mjs` and the capacity settler, never from the historical simulator. The mechanism is specified in §4, including its integer-allocation, block-vintage, and λ-composition rules; four quantities are venue policy rather than mechanism (κ, λ, fees and their base, and the residue owner) and, like a fee table, they belong to whoever runs the market and must be published before it opens.*

Scope of conformance. Passing the suite verifies one thing: that an implementation's entry-and-settlement arithmetic matches §4; the exit semantics of §4.4 (transfer, split, merge) and §12's freeze rule are specified but not exercised by the executable suite. It is not an endorsement of any venue and says nothing about a venue's resolution integrity, custody, or solvency (§14). The specification and code are provided as is, without warranty of any kind, and every simulated return in §13 is a property of the stated crowd models, not a projection.

Cite as: Hunch Research (2026). "The Vested Parimutuel: Paying prediction markets for information when it is worth the most." August 2026.