The Vested Parimutuel

Paying prediction markets for information when it is worth the most

Hunch Research · Public whitepaper · v3.2 · August 2026 · playhunch.xyz

> **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 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 return set the instant you enter that can only rise; 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).

Everything here is checkable. A zero-dependency simulator, a conformance suite, and an O(1) reference settlement algorithm ship with this paper. CC BY 4.0 and MIT. **There is no token, and there will not be one.**

And we got things wrong. Version 2 of this paper claimed six properties. Adversarial review found that two were false as stated and that the cold-start rule admitted an unbounded-leverage attack that would have made the paper's headline strategy a way to extract value while supplying nothing. Appendix A is the record: what broke, the failing case, and what replaced it. We would rather publish that than have you find it.

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 a bound on the residue introduced by fixed-point arithmetic. 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 — against the classic parimutuel, which prices every entrant at the closing ratio. 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, 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.

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. **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. **Scoring-rule market makers** (LMSR and its relatives) solve pricing elegantly and their worst-case loss is *bounded* at `b·ln n` — but that bound is a subsidy someone must fund, per market, forever, which is exactly what a venue hosting millions of markets cannot do. **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 — continuous trading, first-dollar cold start, no subsidy, entry-time price integrity, 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 one-line consequences and their hypotheses stated exactly (§5).
2. A constant-time formulation (§6). The rule reads as an $O(n)$ loop per stake and therefore $O(n^2)$ per market; it collapses to one running accumulator per outcome, updated in $O(1)$. This 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)`, verified numerically to 0.002%, with the corollary that when a pool grows e-fold after your entry, your break-even belief is exactly the pool ratio you entered at.
4. The λ family, measured across its range rather than at two points (§13).
5. A published account of the mechanism's failure modes, including the two false claims and one unbounded-leverage attack we found in our own previous version (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 version 2 of this paper claimed novelty it did not have. Nor do we claim an equilibrium result: §9 and Appendix B are incentive arguments and conjectures, and we mark which is which throughout.

---

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-derivations 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. 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. We claim only that a bounded per-market subsidy is still a per-market subsidy, which is the binding constraint in the long tail.

2.3 Parimutuel microstructure. Thaler & Ziemba (1988) is the source of the favourite–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, Bharadwaj & Robinson (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 and belongs in any honest comparison.

2.5 Resolution. §12 does not propose a new oracle. Augur (Peterson & Krug), UMA's Optimistic Oracle, and Kleros 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. Their cold-start answer is an optional presale phase clearing at a constant price (p. 5); ours needs no extra phase. Their own limitations section notes that without a presale, *"the earliest entrants receive the entire share of counterparty-liquidity rewards"* (p. 10) — a concentration effect that converges 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 $F = (e_1 \dots e_m)$ 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.

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, $s_i + v_i(t)$
$B_w(t)$, $P_w(t)$	book on outcome w at t, and its total principal
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 and total vested-in
λ	vesting fraction; $\lambda = 0$ classic, $\lambda = 1$ pure
$\Pi(t)$	accepted pool at t; q a book's share of it
Π_i	payout to i at resolution

Two 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. 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.

---

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

Rule 1 alone is what version 2 of this paper specified, and it is unsound: with no cap on what a single position 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 on our own simulator, a seeding straddle returned +155% at $50 a leg and +82,037% at one cent a leg — meaning the profit-maximizing strategy was 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 + κ(1 + ln g)`, `g` the book's subsequent principal growth (§5, P5); it is *not* a flat `1 + κ` 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.1). Rule 2 is what a betting exchange already does — a matching constraint, with the partial fill as its familiar consequence.

4.2 A worked example

Binary market, `κ = 9`, fees zero, `λ = 1`. The creator seeds $25 on each outcome in a reserved vintage 0 whose legs match each other (§4.4). Five ordinary stakes follow.

#	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 2.470×):

position	staked	vested	payout	multiple	classic
creator (NO leg)	$25	$170.09	$195.09	**7.804**×	2.470×
B (t=30)	$200	$360.79	$560.79	**2.804**×	2.471×
D (t=90)	$200	$94.12	$294.12	**1.471**×	2.471×

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 breakeven 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 2.125× (immediately before E) to 1.680× — E takes 21% 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

Version 2 handled the empty-book case with an "unallocated bucket" metered out to later entrants at up to ρ× their stake, decaying with market age. That construction was the paper's only tuned parameter and it had three defects: it was strictly dominated by the ε-position attack above; its residue at resolution had to be refunded, which meant a losing stake was not fully at risk; and its decay was indexed to clock time while the risk it priced is informational, so a market that settles informationally early reinstated most of the attack. Rule 2 deletes all 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.

The cost is real and should be led with rather than buried, and it is larger than the binary case suggests. **In a binary market** the favoured side receives partial fills once the underdog book is saturated — an honest liquidity signal, and one that cannot be weaponized: to exhaust the headroom a buyer needs, you must stake their own side yourself, at full risk (conformance case P10 — blocking *is* exposure). **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 `(κ+1)·P_w > Π` in each opposing book `w`, so the *thinnest* book gates entry on all outcomes opposing it. That is both a liveness failure — a three-outcome market with one unloved branch refused 95.9% of gross volume at κ = 9 in our tests — and a griefing lever: a $450 stake on a side outcome blocked a $2,000 informed entry entirely (P9). 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 consequence for venues is stated as a rule rather than left to be discovered: **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 — 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 — 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 — each leg is accepted only up to κ times the capacity the other legs grant, so an asymmetric seed is partially refused rather than allowed to violate the leverage bound. This also gives §12 the two-sided bond its resolution story wants.

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: they vest to the books as of the previous block and not to each other, so there is no intra-block ordering game and no priority auction over vintage. 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.

Fees are charged on stake at entry; 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. λ = 0 is the classic parimutuel; λ = 1 the pure mechanism. §13 measures the interior rather than asserting it — and the measurement corrects a claim from version 2: degradation is **monotone**, not linear. The late entrant's reward is exactly linear in λ, while dilution protection is strictly *convex*, so the dial protects early capital better than a linear reading suggests.

---

5. Properties

Stated for λ = 1 unless noted. Each hypothesis is load-bearing; version 2 omitted several of them and two of its claims were false as a result (Appendix A).

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. Under Rule 2 nothing enters the pool that cannot be so placed. ∎

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. 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 a 50%-of-pool late snipe, early winners' accrued claims are untouched by construction, and their total payout falls 0.2% on average (2.6% worst) — against 13.0% average and 59.6% worst under classic rules.

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.

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 $\kappa \cdot P - V$, 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 $1 + \kappa \cdot (1 + \ln(P_T/P_\tau))$. *Verified:* every winner obeyed the bound over 14,520 positions in adversarial random markets; an ε-position probing a seeded market returns ≈ 9–10× at $\kappa = 9$ identically for stakes of 1 to 10,000 units — against 200,001× for the 1-unit probe under Rule 1 alone. *This is not a general sybil-resistance claim, and version 2's was wrong.* Wash-betting is strictly **costly**, not neutral: a trader staking x on each side into non-empty books vests to strangers on both sides and receives nothing from their own opposing leg, since same-vintage entries do not vest to each other. Measured: a \$100/\$100 wash into an existing \$100/\$100 market returns −25% or −50% before fees. Self-trading is a donation to the standing book.

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 11,982 branches at $n = 2\ldots4$ with random later flow: zero violations, worst case exactly breakeven. *The hypothesis is essential.* Version 2 claimed this "across all outcomes and all histories." It is false without the reserved vintage: if any third party enters before the creator completes its legs, the floor fails in 50% of branches, worst case −50%. That failing case is pinned in the conformance suite.

P7 — Blend degradation. At blend λ a late entrant's multiple is exactly `λ·1 + (1−λ)·M_classic` — linear — while an early position's dilution under a late entry is at most `(1−λ)·D_classic`, with strict inequality for every position earlier than the pool's stake-weighted mean vintage. *Measured (§13): the late multiple falls 1.260 / 1.196 / 1.132 / 1.069 / 1.005 across λ = 0 … 1, exactly linear; dilution falls 13.0 / 8.7 / 5.3 / 2.5 / 0.2%, strictly below the linear reading.*

5.1 What the mechanism does not guarantee

Stated here rather than left to the reader to discover.

- 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.
- There is **no equilibrium theorem.** §9 and Appendix B are arguments and conjectures.
- **Large early capital compresses everyone else's return**, and nothing in the mechanism prevents it (§9).
- **λ < 1 voids P6.** The classic component pays the seeder less than its stake whenever its realized-side share falls below its pool share; the measured minimum at λ = 0.5 is −14.0%.
- κ and λ are **venue policy**, published like a fee schedule — not identities.
- The mechanism says nothing about **resolution**. A perfect settlement identity resolved by a liar pays the liar's friends exactly (§12).

---

6. Settlement in Constant Time

Rule 1 reads as a loop over every opposing position on every stake — $O(n^2)$ per market. On-chain that is fatal: at roughly 5,000 gas per position touched, a book of a few thousand positions makes a single entry exceed Base's block gas limit and the market becomes permanently un-enterable. Version 2 described the settlement identity as "one loop over a public trade log — cheap to compute on-chain," which was 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. 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 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: the accumulator form and the reference settler agree to within integer-rounding across 7,691 positions in adversarial random markets — conformance property P8.)*

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 and partial fills occur (`vpm-accumulator.mjs`).

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`. We verified across 27,150 settlements of the full mechanism — partial fills included — at scales 10^{12}, 10^{18} and 10^{27}, at $\kappa = 3$ and $\kappa = 9$, that the accepted pool is never overpaid and the residue is bounded by **one smallest unit per winning position**, identically at all scales. Dust-position spam cannot farm it: the truncation is global rather than per-position-per-event, so 300 one-unit positions against a whale extract at most 3 units — which accrue to the pool, not the spammer. The residue must nonetheless have a named owner in any implementation; an unassigned residue is funds no one can withdraw.

---

7. The Vesting Yield

The mechanism's pricing content is one formula. Let `q` be your side's share of the pool, approximately constant over the interval, and `Π` the pool. A stake `s` entered at `t` accrues

```
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 numerically: relative error 0.002%–0.009% against the settlement rule in the continuous limit, converging as O(1/steps).

Three consequences.

This is what "entry-time price integrity" means. 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 is not asserted in this paper; it is this line.

$L = 1$ 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, and it is a real cost. Both sides are simultaneously positive-EV whenever $p/(1-p)$ lies in $((q/(1-q))/L,\ L\cdot(q/(1-q)))$, 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.

---

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 — 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 plainly how strong this is, because the rest of the paper leans on it. It is the weakest section here. There is no structural source of uninformed selling pressure 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. Worse, the deployment we propose first (§15) is short-duration recurring markets, where a secondary market in positions cannot form inside a fifteen-minute round. **We therefore treat §8 as a research direction rather than a result**, and §11's collateral and index constructions are gated on it.

---

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. Three honest observations.

Entering is a bet on the outcome and on future flow. 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.

Early entry weakly dominates, conditional on entering. 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. The unconditional timing problem is open (Appendix B).

Scale crowds out time, and we cannot fix it. Vesting is pro-rata by principal, so a large enough first vintage absorbs most future flow. Measured: growing the seed from $50 to $5,000 a leg — about three times the organic pool — compresses ordinary early winners from 1.694× to 1.089×, while the seeder sits at its floor. 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 reserved vintage is a granted privilege, and we choose it with eyes open. 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, 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.

A structural objection, stated rather than dodged. An entrant's return above principal comes from the deposits of subsequent entrants, and returns are monotone decreasing in arrival order with the last cohort receiving exactly zero. That is true and follows directly from Rule 1. It is not a fraud — the pool is zero-sum, payment is genuinely contingent on the event, conservation is exact, and nobody is promised a return — but the time-priority ordering deserves to be named by us rather than discovered by a critic. Two practical consequences: a venue's headline volume is a poor trust signal when the best-rewarded strategy is being early with capital, so §12's track records should weight distinct counterparties and bonded stake rather than handle; and any venue deploying this should take advice on how profit-from-later-inflows is characterized in its jurisdiction, particularly for the tokenized constructions of §11.

---

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.

Two corrections to how version 2 put this. 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 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. And all of it is gated on §8 actually working. We are not claiming yield instruments here; we are noting what the primitive makes possible if the secondary layer materializes.

---

12. Resolution

Liquidity is half of permissionless; the other half is who says what happened. We propose nothing new. 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 Hunch Bazaar contributes 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.

Two things this mechanism changes and §12 must therefore handle. A creator's seed is now a two-sided bond, which is a better starting point for the escalation path than reputation alone. And because accrued claims are monotone, 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 rather than behavioural: fix the resolution timestamp at creation, freeze the accumulator there, and let the resolution transaction land whenever it lands. Then latency has zero payoff impact.

---

13. Simulation

Two studies. The first 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. The main study runs the actual mechanism — Rules 1 *and* 2, κ = 9, creator-seeded vintage 0 — over 2,000 fifteen-minute binary markets per seed, **20 seeds, reported as mean ± 95% CI**: 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   # main study:
capacity rule, CIs, fees, κ sweep, n=3
node docs/whitepaper/sim/vpm-behavior.mjs --markets 2000           # §13.2: responsive
agents, Brier, calibration
node docs/whitepaper/sim/vpm-sim.mjs --markets 2000 --seed 42 [--lambda L]   # λ sweep
(uncapped; see §13.1)
node docs/whitepaper/sim/vpm-capacity.mjs      # capacity-rule property tests
node docs/whitepaper/sim/vpm-accumulator.mjs   # O(1) equivalence + dust bound
node docs/whitepaper/sim/vpm-edge-tests.mjs    # conformance suite, incl. every failing
case
```

Payout by entry time (median multiple on winning positions, by decile; mean of per-seed medians, 20 seeds — Figure 2). Every CI is at most ±0.017:

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	1.13	1.09	1.06	1.03	**1.004**

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 κ = 9 — the cap exists for the adversarial case (§5, P5) and barely touches ordinary flow. At κ = 3 it refuses 10.1% and the decile pattern is unchanged; above κ = 9 it never binds in these markets. The same pattern holds at n = 3 outcomes (500 markets, conservation exact in all, deciles 4.50 → 1.01).

Post-fee, because the fee is not a detail. At a 2% entry fee the last-decile multiple is 0.984 — 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 breakeven minus fees, i.e. −2%.

The λ dial (Figure 4). Every column is the same 2,000 markets:

λ	sniper's multiple	early-winner dilution (mean / max)	seed straddle mean / min / % positive
0 (classic)	1.260×	13.0% / 59.6%	−19.5% / −45.8% / 15.8%
0.25	1.196×	8.7% / 32.6%	+25.0% / −28.7% / 68.0%
0.5	1.132×	5.3% / 18.6%	+69.6% / −14.0% / 96.7%
0.75	1.069×	2.5% / 8.9%	+114.1% / +0.7% / 100%
1	**1.005×**	**0.2% / 2.6%**	+158.7% / **+15.3%** / **100%**

The sniper stakes 50% of the pool on an >85% leader at t = 0.95T, in the ~1,600 markets per seed where that leader goes on to win. Under the full mechanism (20 seeds): early-winner dilution **0.19% ± 0.00 mean, 4.9% worst across all seeds**, against classic's 13.0% / 59.6%; the sniper's own multiple is **1.005 ± 0.000**, and Rule 2 caps its fill at 97.7% ± 0.1 of intended size. Two honesty notes. The subsample is conditioned on the snipe being *correct*, so it measures the attack's payoff when it succeeds — an upper bound on its profitability. And the residual 0.19% is not a violation of P3: accrued claims are untouched by construction; it is early winners sharing the last 5% of opposing flow with the newcomer, which is the rule working as intended.

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%**	−14.0%	−45.8%
worst case, post-2%-fee	**+8.6%**	—	—
markets settling positive	**100.00% ± 0.00**	96.7%	15.8%
mean	+164.7% ± 1.2	+69.6%	−19.5%

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: under classic rules, enter when belief clears the pool ratio plus fee and edge; under this mechanism, enter when `b·(1+ŷ) > 1 + fee + edge` with ŷ from §7 and remaining growth estimated as `L̂ = ln(T/t)`. A third arm is the actual deployed baseline: a classic pool **with a lock at 0.95T**. Same belief stream, same arrival opportunities, same stake sizes; only the entry decision differs. 2,000 markets, 2% fee (`vpm-behavior.mjs`).

Volume moves where the theory says it moves. Share of volume by phase of market life:

arm	first third	middle third	67–95%	final 5%
classic (no lock)	32.3%	33.1%	29.3%	5.3%
classic + lock	34.1%	34.9%	30.9%	0.1%
this mechanism	**43.7%**	34.7%	19.9%	**1.7%**

The mechanism does with incentives what the lock does with a rule: late pool volume collapses (5.3% → 1.7%) without prohibiting anything — while pulling a third more volume into the market's most informative phase. Three robustness notes, because the result is only as good as the agents. First, the growth estimator `L̂ = ln(T/t)` is *optimistic* early (it overestimates realized log growth by roughly half a unit in the first third), so a fourth arm re-runs the mechanism with agents using `L̂/2`: the shift not only survives, it strengthens (48.2% of volume in the first third) — the early migration is the incentive, not the estimator. Second, total participation falls: \$2,730 → \$2,261 per market under the naive estimator, \$1,757 under the conservative one. The mechanism trades volume for time-placement, and we report the level rather than only the shares. Third, realized per-entry PnL, settled under each arm's own mechanism post-fee: classic +6.4%, lock +5.7%, mechanism with naive `L̂` **+0.9%**, with conservative `L̂` +5.1%. Agents trusting the optimistic estimator overtrade to near-breakeven — a caution for §10's agent operators that the entry rule's profitability is exactly as good as its flow forecast, which is §9's point made empirical.

And the cost is paid where the theory says it is paid. Brier score of the pool ratio against the realized outcome, by phase:

arm	first third	middle	67–95%	final 5%
classic	0.238	0.201	0.167	**0.149**
this mechanism	0.239	0.202	0.172	**0.176**

Early and mid-life the pools are equivalent forecasters. In the final phase the mechanism's pool ratio is a *worse* forecaster than an unlocked classic pool — 0.176 against 0.149 — because rational late flow, which would have corrected it, rationally stays out. This is the measured price of dilution protection, reported rather than hidden, and it is the quantitative case for §8: the late price must come from somewhere else. The favourite–longshot calibration table from the same run is discussed in §14.

13.1 Threats to validity

What v3.0 owed and what remains.

- ~~Finite κ~~, ~~single seed~~, ~~pre-fee~~, ~~binary only~~, ~~no behavioural response~~, ~~no forecast accuracy~~ — closed above.
- **The secondary layer is still not simulated.** §13.2 measures the hole (late Brier 0.176 vs 0.149); nothing here measures whether the position market fills it. Backtesting the composite estimator against recorded live tapes remains the most valuable missing experiment.

- **The behavioural agents are myopic.** Entry rules are one-shot threshold rules with a stated constant-arrival growth estimate, not equilibrium strategies; the unravelling question (Appendix B) is untouched by construction.
- **The simulator processes entries sequentially**, approximating rather than implementing block-vintage batching.
- The λ-sweep table retains the uncapped rule; at κ = 9 the capacity cap refuses ~0.05% of gross stake in these markets, so the approximation is measured to be immaterial there, and the main table above is fully capacity-ruled.

---

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). 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.

Late-stage pool odds are not a price, and the cost is now measured. Beyond §7's two-sided band, unrewarded late entry means the ratio stops tracking probability near resolution. §13.2 puts a number on it: with agents responding to the rules they face, the pool ratio's final-phase Brier score is 0.176 under the mechanism against 0.149 under an unlocked classic pool — the mechanism *pays* forecast quality in the pool's last act to buy dilution protection for everyone before it. The live late price is §8's job, and §8 is the weakest part of this design.

Inter-block MEV is real and this mechanism creates it. 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 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. A venue on a centralized sequencer should treat flow confidentiality as a mechanism parameter and publish a non-extraction commitment covering ordering and flow visibility. We claim the narrow thing — no intra-block ordering game — not the broad one.

Hedging near resolution is unpurchasable at λ = 1. 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 that is the honest reason λ exists.

Carry is unmodelled and binds on long horizons. Capital is locked from entry to resolution. At 4% and six months the hurdle exceeds the entire observed late-tercile yield. λ = 1 is a short-horizon mechanism.

The favourite–longshot bias gets measurably worse, as predicted. Since $y \propto (1-q)/q$, the mechanism pays a subsidy for backing the minority side that increases as that side shrinks. If the bias is driven by probability misperception rather than risk preference (Snowberg & Wolfers 2010), the misperception is untouched and a mechanical inducement is added pointing the same way. §13.2 measures it with responsive agents: in the 0.2–0.3 pool-implied bin, outcomes realize at 0.01 under classic rules and 0.05 under the mechanism; in the 0.7–0.8 bin, at 0.98 versus 0.95. The direction is as theory says — the pool tilts toward longshots — and the magnitude in this crowd model is small. Realized longshot returns rise while the price anomaly worsens slightly: a cost the venue accepts knowingly, because the pool ratio was never the published probability in the first place (§7, §8).

Cascades. 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. And the pool's first stake is a creator's seed carrying no information at all.

No equilibrium theorem, and the rational-flow horizon may unravel: 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. Whether that terminates at an interior fixed point pinned by the noise rate, or collapses to vintage 0 plus noise, is open and is the sharpest question we can hand a theorist.

Verifiability is not total. §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 and fee schedule, which should be committed on-chain per market. A paper that claimed otherwise would be doing what it criticizes.

---

15. Deployment

Hunch operates real-money parimutuel markets on Base today. The mechanism here is not live; we would rather publish it before shipping it than the reverse. The order its own risk structure dictates:

1. **Paper first.** Flow-vesting settlement in the paper-money twins, behind the existing payout authority.
2. **The secondary layer**, which is additive under any settlement mode and is the source of the "cash out now" affordance no pure parimutuel offers. Ship it before, not after — §8 is a precondition, not a follow-up.
3. **Real money where P4 bites:** short-duration recurring markets, where $\lambda = 1$ deletes the lock window. Note the tension with step 2 and resolve it before shipping, not after.
4. **Creator-seeded permissionless creation**, with the reserved vintage 0 as the resolution bond.

Why Base. The mechanism is chain-agnostic — the settlement identity is one accumulator update per outcome — and we would rather it ran everywhere. But 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. At two seconds they are small; at twelve they are not. The rest of §10's assumptions are, on Base, already true rather than forthcoming: settlement in native USDC with per-transaction builder attribution, and x402 — the machine-payable standard this paper leans on — live against a working agent rail alongside MCP servers and typed SDKs. We did not design the mechanism for this chain; we designed it against a set of assumptions and then noticed we were already standing inside them. The reference implementation ships where we operate. 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 — 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 are not claiming that insight is ours. Pennock had the essential half of it in 2004, and the honest description of this work is that it makes entry-time pricing an accounting identity, gives it a constant-time form that can actually run on-chain, derives what an entrant is priced at, and supplies a dial to the classic mechanism — while being specific about 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. Late information should enter through the secondary layer, where it pays risk-bearers instead of taking from them. Whether that layer actually forms is the open question in §8.

Doesn't that destroy the price signal? It moves it, at a real cost. 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.

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.

What stops a whale from seeding everything? Nothing. §9 measures it: a seed three times the organic pool takes ordinary early winners from 1.69× to 1.09×. We flag it as the most likely real-world drift.

Is there MEV? Not within a block. Between blocks, yes, and §14 says so plainly — this mechanism makes advance sight of order flow directly monetizable, which is a new risk class relative to a classic parimutuel.

Is there a token? No, and there will not be one. There is nowhere in the mechanism to put one.

Can I implement it? Yes. CC BY 4.0 on the text, MIT on the code, any chain, no permission. If you call it a Vested Parimutuel, pass the conformance suite.

How do I break it? Please do. We found three real breaks in our own design this way; they are in Appendix A with their failing cases preserved.

Appendix A — What We Got Wrong

Version 2 of this paper was reviewed adversarially before publication. Four defects in the artifacts and three in the claims survived that version. All are fixed above; the failing cases are pinned in the conformance suite so no implementation can regress into them.

A.1 — The ε-position attack (fatal, redesigned). Vesting was pro-rata over the opposing book with no per-position cap, so the first position on an empty book received 100% of opposing inflow until a second joined. Return on capital diverged as the position shrank: the seed straddle returned +155% at $50 a leg, +1,104% at $1, and +82,037% at one cent. The dominant strategy was to post dust on every new market and toll organic flow while supplying nothing — which would have destroyed the paper's entire argument that agents supply liquidity. **Fix:** Rule 2, applied at the book level so distribution stays pro-rata: return on capital is scale-invariant (dust earns what any same-moment capital earns) and bounded by `1 + κ(1 + ln growth)` (P5). An earlier draft of the fix capped positions individually; that closed the attack too, but broke the O(1) accumulator — allocation stopped being pro-rata whenever a cap bound, which it did in 82% of test settlements — so it was replaced before publication, and the seed was additionally required to obey Rule 2 among its own legs (an asymmetric seed had violated the bound at 13.5×).

A.2 — The bootstrap bounty (deleted). The ρ-metered bucket was strictly dominated by A.1, left losing stakes partially recoverable through its residual refund (so late entry was not unconditionally negative-EV), and decayed in clock time while the risk it priced is informational — a market settling informationally early recovered most of the attack the decay was introduced to close. Rule 2 removes the bucket, ρ, the decay schedule and the refund together, taking the paper's only tuned parameter with them.

A.3 — P6 was false as stated. Version 2 claimed the seeding floor held "across all outcomes and all histories." A third party entering before the seeder completes its legs breaks it in 50% of branches, worst case −50% (verified over 8,000 branches). The true hypothesis — the seeder is first on every outcome, in a reserved vintage — holds with zero violations over 11,982 branches at n = 2…4.

A.4 — P5's wash claim was false. "Wash-betting both sides returns your own money minus fees" is true only against empty books. Into an existing $100/$100 market a $100/$100 wash returns −25% or −50% before fees. Wash-betting is costly, not neutral — which is a better property than the one we claimed.

A.5 — The sniper result was a measurement artifact. The simulator appended the sniper to the end of the event array rather than inserting it at its timestamp, so it measured a buzzer entry and reported exactly 1.000× and exactly 0.0% dilution. Inserted honestly at $t = 0.95T$ the numbers are 1.005× and 0.2% mean / 2.6% max. The comparison against classic (13.0% / 59.6%) is barely changed; the claim is now defensible.

A.6 — Three implementation defects. The bounty capacity was computed in floating point and disagreed with the published integer rule on 147 of 2,700 grid points; the λ-blend did not forward the market horizon, silently using a 900-step default for any market; and the settler hard-coded two outcomes, so every `n > 2` claim in a paper specifying `n ≥ 2` was untested and in fact threw.

A.6b — Round two. The re-review of v3.1 found: the O(1) equivalence and dust bounds were proven against the *pre-Rule-2* settler — evidence for a mechanism the paper no longer specified — now re-proved for the shipped mechanism, exactly, with partial fills binding (§6); finite κ's min-over-branches acceptance chokes n-way markets with a dead branch (95.9% of volume refused) and admits a small-stake blocking lever, now analyzed in §4.3 and pinned as P9 rather than discovered by a reader; an incomplete or zero-leg seed silently bricked a market instead of voiding it (now voids at creation, P11); and a §4.1 sentence still claimed the flat `1 + κ` bound that P5 had already replaced with `1 + κ(1 + ln g)`. The behavioral study also gained a conservative-estimator arm and per-arm realized PnL after the reviewer showed the naive growth estimator was optimistic early — the volume-shift conclusion survives both.

A.7 — Claims withdrawn. That the mechanism is "the first market structure whose payout schedule is that value curve" (§2.1 — Pennock). That guarantees "degrade linearly" (they degrade monotonically; dilution is convex, which is better). That there is "nothing undisclosed" (κ, λ and fees are venue policy). That solvency has "nothing to verify" (§14). And "riskless-floored," which is not a phrase that should appear in a mechanism paper.

Appendix B — Toward an Equilibrium Analysis

What can be said in a line, what we conjecture, 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) Does the rational-flow horizon unravel to vintage 0, or terminate at an interior fixed point pinned by the noise rate? (ii) Bound the divergence between pool ratio and consensus probability as a function of flow predictability and λ — §7 gives the leading term, `logit(q) = logit(p) + ln L`. (iii)

Adverse-selection pricing for the cash-out vault, and whether a competitive RFQ set converges. (iv) Optimal λ per market class. (v) Whether strategic flow withholding is profitable outside knife-edge cases — delay forfeits your own vesting and cannot touch accrued claims, so the only available harm is abstention.

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).
- Melee Markets (2026). *The Parimutuel Market Maker: A New Market Structure for Permissionless Prediction Markets*. Public litepaper.
- 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.
- 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).
- White, D., Bharadwaj, D., & Robinson, D. (2021). *TWAMM*. Paradigm.

---

*Reproducibility: every number in §13 regenerates from the published scripts with the seeds and flags shown (`vpm-study.mjs` for the main tables, `vpm-behavior.mjs` for §13.2). **The conformance suite for the mechanism this paper specifies is `vpm-capacity.mjs` (properties P1–P11) together with**

`vpm-accumulator.mjs` (exact O(1) equivalence, dust bounds, seed-validity voiding) — an implementation claiming the name must pass both. `vpm-edge-tests.mjs` preserves the failing cases of the *deleted* v2 mechanism (the bounty's late grab, the residual-refund leak, the unreserved-vintage seed break) as a historical record, labeled as such. The mechanism is specified in §4, including its integer allocation rule. Three quantities are venue policy rather than mechanism — κ, λ, and fees — and, like a fee table, belong to whoever runs the market and must be published before it opens.*

Cite as: Hunch Research (2026). "The Vested Parimutuel: Paying prediction markets for information when it is worth the most." v3.2, August 2026.