HCP

HERETIC CAPITAL WORKING PAPER · CRYPTO SERIES

Real-World Assets: Cutting Promotional Waste with Blockchain Burn

Instrument-Bound Settlement, Item-Level Precision, and the Thermodynamics of Enrollment

No. 002 · KRW-C-2026-002 · July 2026

S. Kaza · J. Rosenberg · T. Wichmann

HCP Institute for Quantum Finance · Williams College of Business, Xavier University

Case study: Aladdin's Castle Meme Coin ($ACM), Solana · first measured intervention

ABSTRACT

American promotional spending is aimed at items and settled against brands. The instruments that carry it — coupons, rebates, sampling offers, loyalty credits — are validated at the brand level and reconciled weeks later in aggregate, so a promotion funding one product routinely moves a different one. This is neither fraud nor breakage. It is unmeasured by construction: no victim, no report, no line item. We describe a real-world-asset settlement architecture in which each promise is bound one-to-one to a unique instrument, extinguished exactly once by an irreversible commitment at the point of settlement, and recorded as an event. Item-level precision follows, and so does something less obvious — a complete census of promotional outcomes as a byproduct of settlement rather than as a separate data product. We then show that the enrollment required to operate such a system is thermodynamically identical to the cooling process of KRW-C-2026-001: each enrolled consumer moves supply out of tradeable float into a durable credential. The counterintuitive result is that a consumer-products marketing budget, spent for entirely unrelated reasons, becomes the mechanism that lowers the temperature of the token that gates it. We report the first measured intervention — a founder lock of 50% of supply — against twenty days of on-chain history, and pre-register four falsifiable predictions.

01

The Blunt Instrument

A promotional coupon looks like a precise object. It names a product, states a value, carries a barcode, and expires on a date. Nearly everything about that appearance is misleading at the point where it matters.

Coupon barcodes have historically carried a manufacturer identifier and a short family code — a category-level designator, not a unique product number. The family code addresses a group of items, and one common value addresses all of a manufacturer's items in that group. A modern symbology exists that can carry the full item number and encode purchase requirements, but enforcement at the lane is uneven, and cashier override remains available at every register in the country. The instrument is printed with item-level intent and validated with brand-level precision.

The consequence is structural. A promotion funded to move one product will be redeemed against whichever product in that family the household was already buying. The shopper is not cheating; the coupon is valid, the register accepts it correctly, and the retailer bills correctly. The money simply lands somewhere other than where it was aimed.

One of the present authors managed Cascade at Procter & Gamble in the 1990s and later acquired the Oxydol brand. In that period a brand-level coupon dropped against a Cascade item was redeemable across the Cascade family, and the redemption data came back showing the family had moved. It had. The intended item frequently had not.

> **The measured outcome of a brand-level promotion is that the brand moved. That is also the only outcome the instrument is capable of measuring.**

This leakage does not appear in any published figure, and the reason is worth stating plainly. Direct coupon fraud in the United States is estimated in the range of several hundred million dollars annually, with broader promotion-abuse losses estimated in the tens of billions; a substantial share of merchants report knowingly tolerating some level of abuse because detection costs more than recovery. Breakage — issued value never redeemed — runs to roughly a quarter of loyalty points unspent and around a tenth expiring outright, against unredeemed balances estimated in the hundreds of billions of dollars. Both categories are counted because both have a claimant. Item-level leakage has none. It is booked as trade spend against volume that would have sold regardless, and it is invisible not because it is small but because nothing in the system was built to see it.

The cost is highest exactly where precision matters most. A newly listed item lives or dies on velocity in its first two quarters. It requires trial *on that item* to survive a line review. Promotional dollars that leak to the established item in the same family fund the incumbent, generate no incremental trial on the new one, and produce a delisting for underperformance. The spending worked against the objective it was funded to achieve, and the reporting could not distinguish the two outcomes.

02

What Settlement Actually Requires

Under prevailing practice a redemption occurs at a point of sale, evidence travels to a clearinghouse, and reconciliation among the funding manufacturer, the honoring retailer and intermediaries completes days or weeks later. During that interval every redeemed instrument sits in an indeterminate state: it may prove valid, duplicated, counterfeit, misreported or disputed. Precision at the item level is impossible in such a system for a simple reason — the system does not track items. It tracks entries in mutable, balance-based records reconciled after the fact.

Three families of alternative have been proposed and none resolves it. Distributed-ledger loyalty systems are predominantly balance-based: points are recorded as fungible balances, transferred between accounts and reconciled in aggregate. They relocate the account book without changing its structure, so no single promise is individually verifiable or individually final. Centralized single-use codes do enforce once-only redemption, but only inside one issuer's private database, which every other settling party must trust. Payment token systems fail differently and more instructively: payment presumes a payer who has surrendered value and may require recourse, so reversal, refund and chargeback machinery is layered on by design.

> **A promotional instrument has no payer. The consumer surrendered nothing to receive it. Reversal infrastructure therefore protects no one — while creating the exact pathway that duplicate redemption exploits.**

Removing that pathway is not a security feature added to the design. It is the design. What remains is an architecture in which each promise is bound one-to-one to a unique settlement instrument; the instrument is non-transferable and time-limited; validation occurs at the point of settlement using a single-use code derived on the consumer's own device without network connectivity; commitment is atomic, so of any number of simultaneous redemption attempts exactly one succeeds; and the committed transition is verifiable by parties other than the issuer. No operation anywhere in the system returns an extinguished instrument to a redeemable state.

Because the offer definition attaches to the individual instrument rather than to a family code, the instrument can specify a single item. The scalpel is available only once settlement stops being an aggregate.

A provisional application covering this architecture was filed with the United States Patent and Trademark Office on 22 July 2026, application 64/116,829, establishing priority as of that date. The system is patent pending; no patent has been granted, and claim scope will be determined in prosecution.

03

The Node

Settlement requires a place. A promotional credit redeemed at a supermarket lane inherits that lane's constraints — no dwell time, no explanation, no staff able to answer a question about an unfamiliar product, and a retailer who owns the resulting data.

The alternative is a venue that already assembles the right audience and holds it still. A family entertainment center — an arcade — holds a parent in place for the length of a child's play session. Warehouse-club demonstration programs, the most effective in-person sampling format in American retail, work with twenty to forty seconds against a moving shopper. The same conversation in a venue where the family has already committed to staying is not a longer version of the same thing; it is a different instrument. It is long enough to state a benefit, give the reason it is credible, and answer an objection.

Three functions then occupy one location. **Fulfillment**: the credit is redeemed and the product is handed over. **Demonstration**: a trained person explains it while the family is stationary. **Information**: the redemption commits an event, and the event is the record.

These are not three businesses. They are one loop. A brand funds the sample; the sample draws a family to the venue; the venue receives traffic it did not previously have and sells its own admissions and play; the redemption produces the record; the record is what the brand was unable to buy anywhere else; the brand funds the next offer. Each leg pays for the following one, which is why no single participant can build the system alone. A brand assembling all three would be constructing a media network, which is not its business and whose economics only close when amortized across many participants. A venue operator cannot originate brand demand. The system is licensed into, not built.

> **It also sits on neutral ground. A manufacturer cannot assemble a cross-retailer view of its own consumers inside any single retailer's walls. Outside them, it is simply a venue.**

No such location is operating today. Instar Brands holds the United States trademark application for Aladdin's Castle, serial 98123992, and has formed a joint venture with Heretic Capital Partners to evaluate a relaunch. What exists at the date of this paper is the trademark, the filed patent application, the token, and software under construction. The venues, the operator network and the brand programs are forward-looking. This paper describes an architecture and the foundation being laid for it; it does not describe an operating network.

04

The Layer That Does Not Exist

Fulfillment can be bought today. Demonstration can be bought today. The third function cannot be bought at any price, and it is the reason to build the first two.

Consumer-goods measurement currently derives from household panels of tens of thousands projected to a national population, or from retailer loyalty data that the retailer owns and that does not follow the shopper across banners. Neither answers the question a brand manager actually has: did *this item* get trial, from whom, and did they return. Panels cannot resolve individual items reliably at low volumes — precisely the condition of every new launch. Retailer data cannot cross the banner boundary.

In an instrument-bound system every lifecycle transition is committed as an event: issuance, redemption, expiration, reimbursement, funding, verification attestation. The record is complete rather than sampled, so redemption velocity, breakage rate and per-offer consumption are computed over every instrument rather than estimated from a panel.

Reimbursement obligations between a funding manufacturer and an honoring venue are computed from the committed events themselves, without a clearing intermediary and without either party trusting the other's private records. Impression-type events are committed only as batched pseudonymous digests and identity verification only as a Boolean attestation, so the record is a record of settlement facts about instruments, not behavioral profiles about persons.

> **The measurement is not a product built alongside the system. It is what the system leaves behind. The exhaust is the thing the data industry cannot manufacture.**

05

The Two-Asset Separation

A settlement system needs to know who is eligible. It does not need that credential to be spendable, and there are strong reasons for it not to be.

The architecture therefore separates two assets. Eligibility is evidenced by holding at least one unit of a membership asset recorded on a distributed ledger, verified through a read-only query, never spent, transferred, staked or consumed by any operation of the system, and in disclosed embodiments held in a vault that prevents transfer for a period of years. Spendable value exists only as the credits: non-transferable, time-limited, redeemable exactly once, extinguished on redemption. Eligibility is durable and non-monetary. Consumption is unique and terminal.

In the commercial embodiment described here the membership asset is $ACM and the credit is a Meme Token Credit. The separation is what allows continuing benefit to accrue to holders without the membership asset ever functioning as a medium of exchange inside the system.

06

The Cooling Machine

KRW-C-2026-001 established that price is governed by two state variables — pool depth, the system's heat capacity, and the distribution of holder reservation prices, its temperature — and that effective float, not total supply, is what prices against the pool. It identified three durable levers: deepen the pool, lock float, and convert paper holders into conviction holders through real utility. The third was described as the only lever that is unbounded and free.

Enrollment operates the second and third simultaneously, and it does so mechanically.

A consumer who wishes to receive credits acquires eligibility, and a unit of the membership asset moves from treasury into a time-locked vault associated with that account, where it cannot be transferred, sold or withdrawn for a period of years. That unit has left the tradeable float — forced cold, in the language of the first paper. The same act converts the holder: a person who enrolled to receive product has a reservation price far above spot, because the reason they hold has nothing to do with spot. Voluntary cold. One operation, both levers.

> **The consumer-products company funding the offer is not trying to affect the token. It is trying to sell detergent. The cooling is a byproduct of a marketing budget spent for entirely unrelated reasons.**

Recall the reflexive supply function of the first paper, in which released selling scales with the unlocked float value F:

$$S(P) = \alpha \cdot (P / P_0 - 1) \cdot F$$

Reducing F reduces the reflexive response at every price level. Figure 3 traces effective float against enrollment at constant founder lock and constant hot fraction. The curve is a straight line because the mechanism is a straight line: one enrollment, one vaulting operation, one unit of float removed. Nothing about it depends on sentiment, narrative or attention.

Figure 3. Effective float against enrolment, at constant founder lock and constant hot fraction.

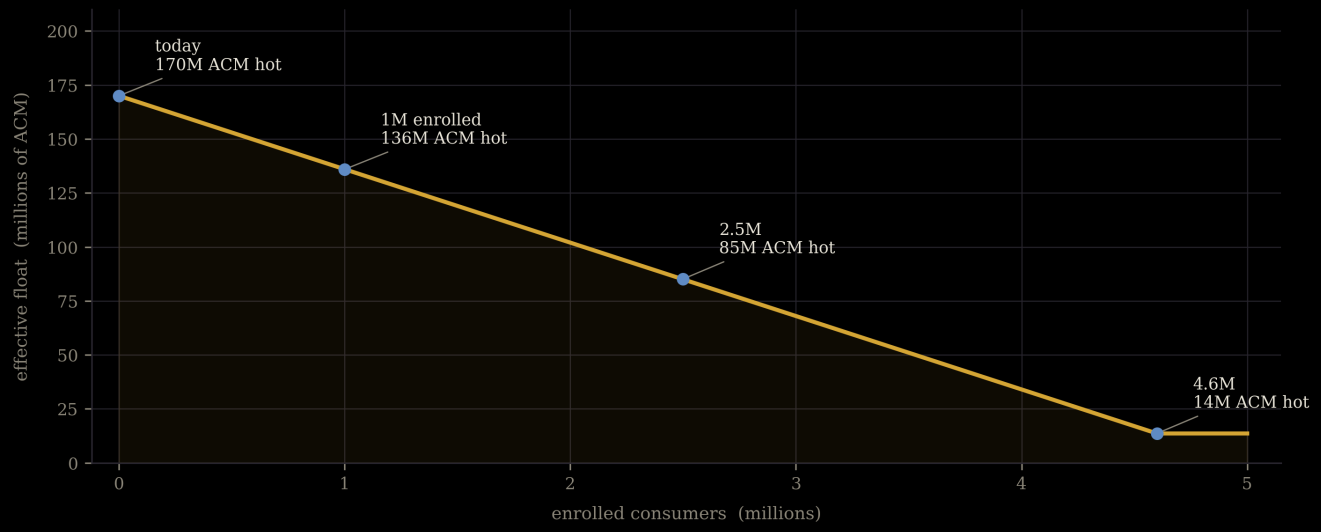

Figure 3. Effective float against enrollment. Founder lock held at 50%; hot fraction held at its measured value of 0.34. The vaulting rate per enrollment is illustrative and not a committed program parameter.

07

The First Intervention

The founder position — 50% of total supply — is locked on-chain. This is the first deliberate change to a state variable in a system whose other variables are publicly observable, which makes it an experiment rather than an observation. What follows is the pre-intervention record against which subsequent papers will measure.

Figure 1. $ACM market capitalization, 3–23 July 2026. Logarithmic axis.

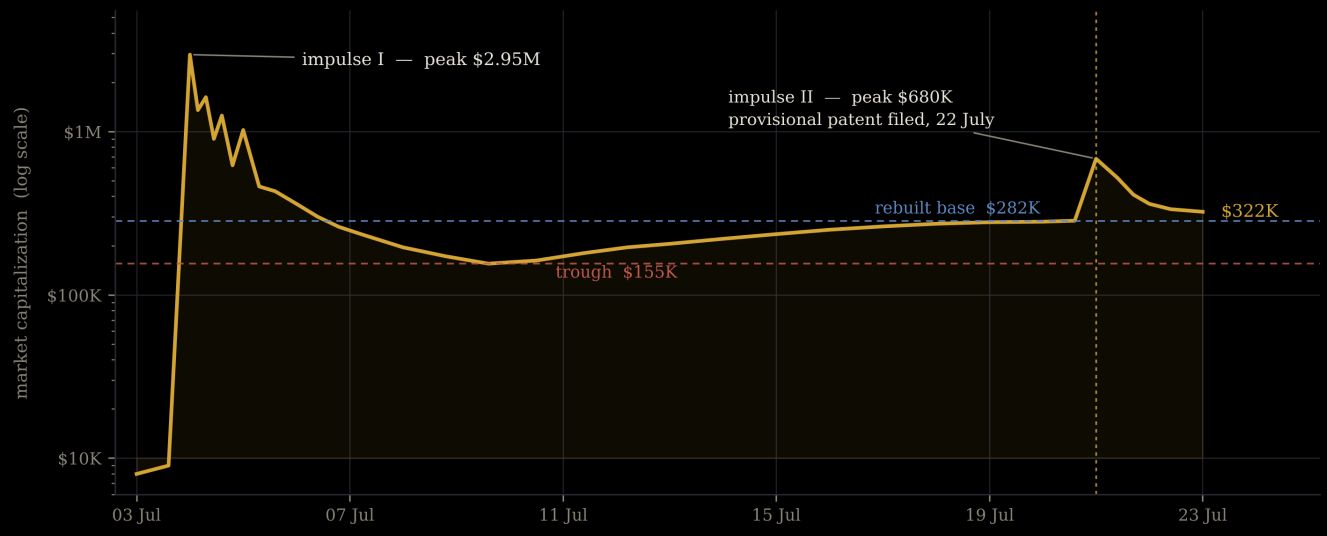

Figure 1. $ACM market capitalization, 3–23 July 2026, logarithmic axis. Reconstructed from public price history; intraday resolution is approximate.

Two demand impulses appear in the window, and they are not the same kind of event. The first, on 4 July, was a launch: attention arriving at a newly created pool. It carried the system to a peak near $2.95M and then produced the oscillating

decay characteristic of a thin pool — successive lower highs at roughly $1.6M, $1.25M and $1.0M — before settling into a trough near $155K by 12 July.

The second, on 22 July, was the filing of the provisional patent application described in section 02. It carried the system from a base near $282K to approximately $680K and settled at $322K. This distinction matters more than the dates. The first impulse was generated by attention to an asset; the second by the arrival of a fact about the underlying business.

Comparing them requires care. The first impulse has no meaningful pre-impulse base: the token had just launched, so any ratio to a starting price is an artifact of when one starts measuring. Retention of peak avoids the problem entirely and requires no assumption.

MEASURE	IMPULSE I · 4 JUL	IMPULSE II · 22 JUL
Peak market capitalization	$2.95M	$680K
Level retained	$155K	$322K
Retention of peak	5.3%	47.4%
Settles relative to pre-impulse base	far below	1.14×
Pool depth at impulse	thin	≈$93K

Figure 2. The same system, eighteen days apart. Overshoot fell 8×; the fade stopped below base.

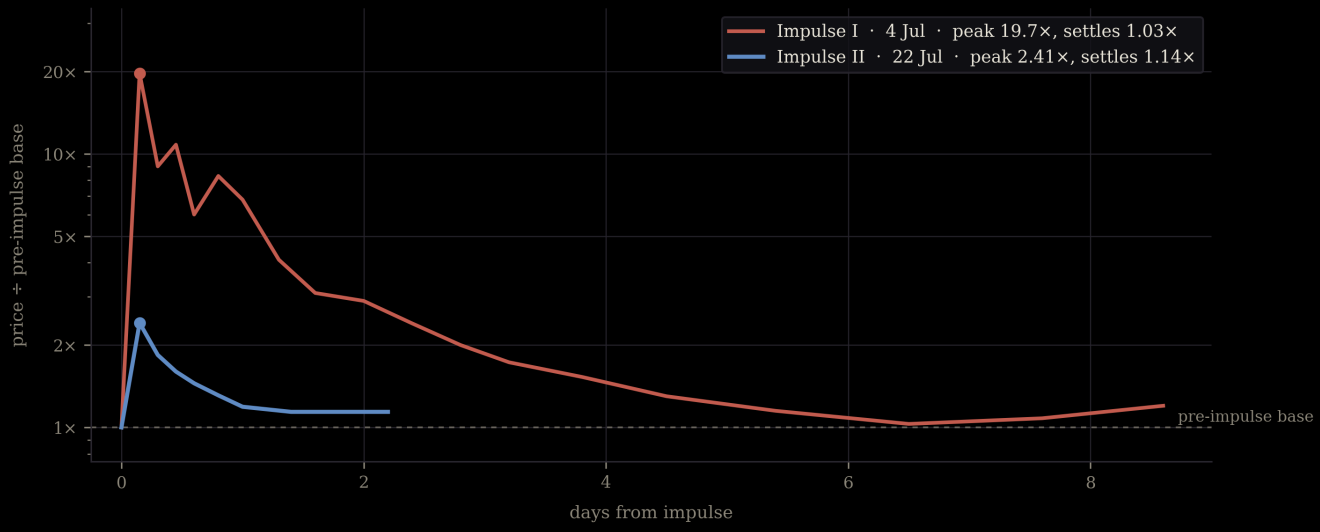

Figure 2. Retention of peak for both impulses. Each curve is scaled to its own maximum, so no base price is assumed.

Retention rose from 5.3% to 47.4% of peak, and the second impulse settled above its starting level rather than below it. The floor across the window roughly doubled, from about $155K to about $282K, while peak amplitude fell by 77%. Under the framework of the first paper this is the signature of standing supply that does not wake: less of the float is hot, so a comparable impulse produces less overshoot and less released supply on the way down.

> **The decline between 4 and 12 July was not the failure of the system. It was the system expelling heat. The fade is what the cooling looks like from the outside.**

One further observation strengthens the method of the first paper. KRW-C-2026-001 conceded in its limits that the reservation-price distribution is not directly observable and that temperature must be inferred. That concession was too strong. Public pair data report unique buyers over a window and holder count at its end; the difference between them is the

fraction of participants who acquired and did not keep. On the day of the second impulse, roughly 3,800 unique buyers transacted against a holder count near 1,600, implying that more than half of the window's buyers had already exited. Buyers-minus-holders is a coarse but public, per-token, time-stamped proxy for the hot fraction, available on every asset on the chain and requiring no privileged data. We adopt it here and use 0.34 as the working hot-fraction estimate for circulating supply.

08

Does the Model Fit?

A series that reports only confirmations is not measuring anything. The question is whether the price history is consistent with the reflexive model of KRW-C-2026-001, and the answer is that it fits in shape, fits partially in magnitude, and fails in a way that identifies what the first paper is missing.

Shape fits. Figure 1 of the first paper contrasts an identical demand impulse through a thin pool — overshoot, then fade below the starting level — against the same impulse through a deep pool, which climbs smoothly to a modest multiple and holds. Impulse I is the first archetype and impulse II is the second, and no parameter fitting is required to see it.

Magnitude fits partially, and yields a first estimate of α. The constant product invariant gives $P'/P = (1 + C/R)^2$. At impulse II the quote-side reserve was approximately $46K, so a peak of 2.41× base implies a net inflow of about $25.4K. Holding the settled level of 1.14× requires a residual of roughly $3.2K, so approximately $22.3K was absorbed by released supply. Substituting into the reflexive supply function of the first paper and evaluating at the peak deviation gives:

$$\alpha \approx 0.11 \ (F = \text{unlocked float}) \quad \alpha \approx 0.33 \ (F = \text{hot float})$$

The first paper left α unfitted. This is the first empirical estimate of it on any asset, and it is an order-of-magnitude estimate rather than a measurement.

> **A model that only ever agrees with the data is not a model. Two things here do not agree, and both are informative.**

First failure: the static form is path-dependent and does not say so. S(P) is written as a function of instantaneous price. But released supply during a fade depends on the whole path the price traverses, not on its maximum. Solving the model for its self-consistent equilibrium using the peak-evaluated α returns a retained level near 1.71×, against 1.14× observed. The model therefore under-predicts the fade unless α is understood as a path integral rather than a point evaluation. The first paper does not distinguish the two, and this data shows the distinction is material.

Second failure: there is no relaxation time. Impulse I required roughly eight days to settle. Impulse II settled in under one. The model predicts terminal levels and says nothing about how quickly a system reaches them, so an eightfold difference in relaxation rate is entirely outside it. A future paper should introduce a relaxation constant, and the obvious candidate is the ratio of pool depth to hot float — the same quantity that governs heat capacity elsewhere in the framework.

And a confound that cannot be resolved at n = 2. The higher retention of impulse II is consistent with a cooled float. It is equally consistent with section 05 of the first paper, which holds that marketing the mission recruits holders while marketing the price recruits sellers. Impulse II was a patent filing — a fact about the business, not a price event — so the particles it recruited should have had higher reservation prices for reasons having nothing to do with float. Two mechanisms predict the same observation and this data cannot separate them. Prediction P4 below is designed to.

09

Pre-Registered Predictions

These are stated before the next observation window and are structural rather than directional. None is a price forecast, and none can be satisfied by the price rising.

- **P1 — Overshoot compression.** For a demand impulse of comparable size relative to pool depth, occurring after the lock, peak-to-base overshoot will be lower than the 2.41× of impulse II. If overshoot rises at constant relative impulse size, the mechanism is wrong.
- **P2 — Retention improvement.** Retention of peak will exceed 47.4%. A fall below that level at constant or greater pool depth falsifies the float-reduction argument.
- **P3 — Hot fraction decline.** The buyers-minus-holders proxy will fall below its present value as enrollment begins. If enrollment proceeds and the proxy does not fall, the claim that utility converts holders is wrong.
- **P4 — The null case.** If enrollment does not occur, none of the above should be observed, and the lock alone should produce a one-time step in effective float with no continuing trend. Observing continued cooling without enrollment would indicate the mechanism proposed here is not the operative one.

Each is measurable from public data at no cost, on a stated schedule, by any reader. Subsequent papers in this series will report measured against predicted, including where the predictions fail.

10

Honest Limits

- No venue is operating. There is a trademark application, a filed provisional patent application, a token, and software under development. Everything concerning physical locations, operator agreements and brand programs is forward-looking.
- The patent application is provisional and pending. No claim has been examined, allowed or granted, and scope will change in prosecution.
- Enrollment is the load-bearing assumption of the entire cooling argument. If consumers do not enroll at scale, no float moves, and the mechanism of section 06 produces nothing.
- The authors are simultaneously the experimenters, the intervention and holders of the asset under study. Publishing predictions about a system one is inside changes that system. This is a genuine methodological limit and readers should weight the results accordingly.
- Two impulses are two observations. The retention comparison is consistent with the model and is not, at n = 2, evidence for it. Pool depth, market conditions and attention all differed between the two events and are not controlled.
- The price history in Figure 1 is reconstructed from public charting at daily-to-hourly resolution; intraday extremes are approximate.
- The vaulting rate per enrollment in Figure 3 is illustrative. The supply price-sensitivity α remains unfitted, as in the first paper.
- Buyers-minus-holders is a proxy, not a measurement of reservation prices. It counts participants, not conviction, and will misstate the hot fraction in the presence of wallet fragmentation or automated trading.
- Impulse II coincided with the patent filing. Its higher retention may reflect the kind of news rather than the state of the float, and the two explanations cannot be separated at n = 2.
- The α estimate of section 08 rests on an inferred net inflow, a reconstructed price path and a single event. It is an order of magnitude, not a measurement, and it changes by a factor of three depending on whether F is taken as unlocked float or hot float.
- Item-level promotional leakage is asserted from operating experience and from the structure of coupon symbology. It is, by the argument of section 01, not directly measurable under current infrastructure — which is the claim, and also the reason the claim cannot yet be quantified.

11

Coda

Every promotional dollar spent in America is aimed at an item and settled against a household, and nobody can say which item it moved, because nothing in the chain was built to know. Build the instrument that knows and two things follow that have never before been connected. The money finds the shelf it was aimed at. And the float goes cold — not because anyone marketed the price, but because a stranger wanted free detergent and the system required a credential to give it to them. The measurement was never the cost of running the machine. The measurement is the machine.

RELATED WORKING PAPERS · KRW SERIES

KRW-C-2026-001 — Crypto Thermo: The Thermodynamics of Liquidity, Volatility, and Why Patience Is the Only Free Alpha
KRW-2026-002 — The Gibbs Free Energy of Capitalism
KRW-2026-003 — The Four Eigenstate Scenarios

HCP Institute for Quantum Finance · Williams College of Business, Xavier University · www.hereticcapital.com