THANAAYA TECHNOLOGIES

RESEARCH NOTE

The Overlay Nobody Tested

A hedge was added because it seemed obviously correct, never validated, and cost more than the book it was protecting made.

ABOUT THANAAYA TECHNOLOGIES

Thanaaya Technologies is a software company licensed by the Dubai Department of Economy and Tourism, building since 2021. We design, build and operate production systems across financial software, distributed ledger, gaming, commerce, media and quantitative research — and we run several of them as our own products rather than only delivering them for clients.

We take a small number of engagements at a time, and the engineer who designed a system is the one who gets called when it misbehaves. This paper comes out of our internal research programme, which exists to find out cheaply whether an idea is real. We publish the negative results as well as the positive ones: an undocumented failure gets rediscovered, and a threshold nobody wrote down cannot be checked.

2021	24	80+	9
BUILDING SINCE	PRODUCTION DOMAINS	SERVICES IN PRODUCTION	INDUSTRIES SERVED

Subject: cross-venue binary pricing and delta-hedge overlay design
Programme: internal quantitative research — paper trading only, no capital at any stage
Status: post-mortem on first realised settlements; n = 8, single resolution date
Published: September 2026 · **Licence:** Dubai Department of Economy and Tourism 964356 · **Unified licence:** BL1484

PART ONE

In plain English

No mathematics in this section. A safety mechanism was added because it seemed obviously correct, never tested, and cost more than the thing it was protecting made.

1. The trade

On **Polymarket** you can buy a contract that pays out if "Bitcoin is above $75,000" on a given date. On **Hyperliquid** you can trade Bitcoin futures directly. The two describe the same world, so they ought to agree on the odds.

Our system worked out what the Polymarket contract *should* cost, using the futures price, the distance to the threshold, how volatile Bitcoin was, and how long was left. When Polymarket disagreed with that number, it bought the difference.

In one case the contract for "Bitcoin above $75,000" was trading at **11 cents** — an implied 11% chance — when the model said it was worth considerably more. So the system bought it.

Buying that contract leaves you exposed to Bitcoin itself moving, so the system also placed an offsetting short position in the futures. That is the **hedge**: keep the bet on the mispricing, remove the bet on direction.

Everything below is paper trading. No capital was committed, and the live trading path was never even switched on.

2. The first real results

Positions are held until the question resolves, so nothing was realised for weeks. Then eight settled — seven of them in the same second — and there was finally something to look at.

+$1,234	-$1,273	-$39
made by the Polymarket contracts	lost by the futures hedge protecting them	net result on about $800 of paper stake

> THE TEMPTING CONCLUSION
>
> "The strategy works — it made 154% — and the hedge is destroying it. Turn the hedge off." That reading is available, it is arithmetically correct as far as it goes, and it is **wrong**.

3. Why the obvious reading does not survive the data

Three reasons. Any one of them would be enough.

Eight positions is not evidence of a strategy

Eight. Seven settling in the same second, on the same day, in the same market conditions. That is not eight independent tests — it is closer to **one** test with eight parts. Whatever happened that day happened to all of them at once.

The two legs are the same bet, twice

This is the part that matters most and is easiest to miss. While these positions were open, **Bitcoin rallied from around $63,000–68,000 to $78,636.**

The winning contracts were mostly bets that Bitcoin would finish *above* a threshold. They won because Bitcoin went up. The hedge was short. It lost because Bitcoin went up.

The gain and the loss are the same market move, counted twice with opposite signs. So "remove the hedge and you would have made 154%" really says: an unhedged long position in a rising market went up. That is true, and it is not a finding.

It tested one kind of market

A single settlement date tests a single regime. Had Bitcoin fallen instead, the same eight positions would have produced the mirror image and we would be writing that the hedge saved us. Neither conclusion is supportable from one day.

4. So what *is* the finding?

Not a return. A methodology failure, and a clear one:

> THE FINDING
>
> The hedge was **calculated once, when the position opened, and never adjusted again** — for twenty-six days. No test or backtest ever checked whether that was sensible. It was wired in because it looked obviously right, and switched on.

Why that is a defect rather than a simplification is specific to this kind of bet. The right hedge size for a yes-or-no contract is not stable. It is largest when the price sits near the threshold — when the outcome is genuinely in doubt — and shrinks toward nothing as the answer becomes obvious either way. It also changes as time runs out.

So a hedge set on day one and left alone does not become slightly wrong. It **stops being a hedge** and quietly becomes an ordinary directional bet, in whichever direction it happened to point. That is what these numbers show: a stale short position held through a 15–25% rally, with nothing in the code to notice.

5. Three other things the same post-mortem found

The emergency stop tripped on bookkeeping order, not on losses

A risk limit halted the system when losses appeared to breach its budget. But seven positions were settling in the same instant, and the limit was checked partway through processing them. At that moment the running total looked far worse than the batch's actual outcome, which finished comfortably inside budget. **The kill switch fired on the order settlements were processed in**, and left the system halted for about a day for nothing.

The system had been holding nothing for a month

A filter on which prices the system would accept had been tightened until almost nothing qualified: **45,985 status messages in a row reporting zero positions, across 33 days.** An earlier note in our own log had predicted exactly this, named the setting to change, and recorded that it was adjustable "for exactly this reason". Nobody applied it.

Nobody looked for nine days

The only place these results existed was the system's own log files. They sat unexamined for nine days after the first real outcome this project had ever produced. Not a modelling failure — an operational one, and probably the most embarrassing item here.

6. What we take from it

The rule we adopted is narrow: **a protective overlay is a strategy component and gets tested like one.** If a hedge has never been backtested, it is an untested position, however standard the textbook says it is.

And the smaller one, which cost nothing to fix: a system whose results live only in log files will produce results nobody reads.

PART TWO

The technical record

The settled book, the hedge-disabled re-score and why it is arithmetically admissible, and the six findings from the post-mortem.

1. Subject

Construction	Price a Polymarket binary from a futures-derived barrier model on Hyperliquid; take the mispricing; hedge the resulting delta with a perpetual futures position
Realised history	`8 settled positions — 7 at 2026-09-01 00:08:02, 1 at 2026-09-04 16:29`
Stake deployed	`$800.42, paper only`
Live execution	`never configured — signing key empty, wallet unfunded`

2. The settled book

The entire realised history of this strategy. Seven of eight settled in the same second.

MARKET	RES	ENTRY	STAKE	PM LEG	HEDGE LEG	NET
BTC YES above $75,000	1	0.110	100.00	+809.00	-712.49	+96.51
BTC YES above $72,500	1	0.200	100.00	+399.90	-609.55	-209.65
ETH YES above $2,500	1	0.267	99.99	+274.42	-158.73	+115.69
ETH NO below $2,300	1	0.901	100.10	+11.01	+0.31	+11.32
BTC YES above $76,000	1	0.711	100.11	+40.75	0.00	+40.75
BTC NO above $65,000	0	0.360	100.01	-100.11	+889.97	+789.86
ETH NO above $2,200	0	—	100.10	-100.10	+81.31	-18.79
XRP NO below $1	0	—	100.11	-100.11	-764.45	-864.56
Total			800.42	+1,234.76	-1,273.63	-38.87

Prediction-market leg win rate 5/8; mean win +$307.02; mean loss −$100.11, a total loser being the full stake. Entry rows for the two oldest positions had rotated out of the process logs; their stake is recovered exactly from their legs, both being total losses.

3. The hedge-disabled re-score

Removing the hedge turns −$38.87 into **+$1,234.76** at the *same* 5/8 win rate. This substitution is arithmetically admissible for one specific reason: the hedge routine has exactly two call sites, at entry and at settlement, and there is no rebalancing loop anywhere in the runner. It therefore **cannot influence which prediction-market positions were opened**, so removing it does not change the book being scored.

Replaying the realised-loss budget over both paths: as run, cumulative loss breaches the −750 limit at the final settlement (−883.35) and fires the risk halt. Hedge-disabled, the cumulative path never goes negative and the halt never fires.

> WHAT THIS FIGURE IS NOT
>
> It is not a return, an annualised anything, or evidence that the pricing thesis is profitable. **n = 8, seven of them settling simultaneously on a single date.** It is a counterfactual on one batch, and it is quoted here only to size the overlay's cost against the book it was applied to.

4. The two legs are not independent

This is the caveat that governs every number above, and it is not visible from the totals.

Entry references on the three BTC legs were 63,057 / 68,375 / 68,722; settlement was at **78,636.50**. Over the holding period the underlying rallied 15–25%.

The prediction-market gains are concentrated in three *above-threshold* positions, which resolved in the money *because of that rally*. The hedge was short, and lost *because of that rally*. **The +1,234.76 and the −1,273.63 are the same market move counted twice with opposite signs** — which is exactly what a hedge is supposed to do, and is why the near-cancellation should not be read as the overlay "erasing" an independent edge.

The correct statement is narrower: *on this one date, in this one direction, the overlay was mis-sized badly enough to overshoot the exposure it was hedging.* A move of the same magnitude in the opposite direction would have produced the mirror image and an equally unsupportable opposite conclusion. Nothing here establishes that the pricing thesis has an edge.

5. Finding 1 — a static hedge on a non-stationary greek

The hedge target is computed from a binary-option delta at fill time and never recomputed. A binary option's delta is violently non-stationary: it peaks at the strike, decays toward zero as the option moves deep in or out of the money, and shifts with time to expiry. A hedge set once therefore stops hedging within days and degrades into a stale directional perpetual position.

The settle prices show the cost directly — short into a 15–25% rally, unadjusted for 26 days. Integration documentation exists for the plumbing (per-market attribution, ledger), but **no test, backtest or study ever validated that set-and-forget delta hedging was the right overlay.** It was wired in and switched on.

This is the transferable result: the overlay carried an unexamined modelling assumption of the same weight as the strategy's own, and was exempted from the validation the strategy would have had to pass.

6. Findings 2–6

#	FINDING	STATUS
2	**No exit framework.** Removed 2026-07-10 and never rebuilt; named as the binding constraint in every entry since, and the reason a losing position has no mechanism to be cut.	open
3	**Risk halt fires on an intra-batch transient.** Halt logged at 00:08:02.818 on cumulative −883.35 while seven simultaneous settlements were still processing; the same batch finished at −79.63, well inside the −750 budget. Tripped on settlement *ordering*, not outcome; left the runner halted ~24 h.	open
4	**Entry band starves the book.** The consolidation to [0.10, 0.90] is arithmetically the same filter that produced a measured drought: 80 tokens → 25 pass spread → 19 unheld → 5 pass band → **0 fills**. The fix was named in our own log and recorded as env-tunable "for exactly this reason". Never applied. **45,985 heartbeats at zero exposure across 33 days.**	open
5	Second strategy (attributed-fill following) never built — blocked on a fill-attribution source the venue's socket does not carry.	not started
6	Live execution never tested; signing key empty, wallet unfunded. Everything above is paper.	by design

7. Why the project went quiet

Removing the exit rule made the strategy hold-to-resolution. The venue lists only two expiry clusters — verified across 9,645 sampling rows, with nothing between 30 and 120 days — so the first realised result *could not exist* before the settlement date it arrived on. The gate that was supposed to review it had been set in calendar weeks, against a strategy whose feedback interval was structurally months.

The results then sat unread in process logs for nine days. Both halves of that are process failures rather than research failures, and both are cheap to prevent: match the review cadence to the strategy's actual settlement cadence, and never let a result exist only in a log.

8. What this establishes, and what it does not

ESTABLISHED	NOT ESTABLISHED
A hedge overlay was deployed with no validation of its own, and its sizing assumption is demonstrably wrong for this instrument class.	That the pricing thesis has an edge. n = 8 on one resolution date, dominated by a single directional move.
The overlay's cost on this batch exceeded the book's gross gain.	That removing the hedge is the correct action. Unhedged, the book is a directional crypto position, which is a different strategy.
Three independent operational defects — halt ordering, entry-band starvation, unread results.	Any statement about performance in a falling market, which this sample does not contain.

9. Reusable rules

- **An overlay is a strategy component.** A hedge that has never been backtested is an untested position, however standard its textbook form.
- **Check whether a hedge's greek is stationary before setting it once.** Barrier and binary payoffs are the case where it is not.
- **Evaluate a risk limit against a settled batch, not mid-batch.** Simultaneous settlements make intra-batch cumulative totals meaningless.
- **Match review cadence to settlement cadence.** A monthly gate on a strategy that realises twice a year reviews nothing.
- **A result that exists only in a process log will not be read.**
