



Beaten by the Relay

The opportunity lasted a second. The feed that defined it arrived 1.6 seconds late. A structural result, not a sample-size one.

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

BUILDING SINCE

24

PRODUCTION DOMAINS

80+

SERVICES IN PRODUCTION

9

INDUSTRIES SERVED

Subject: short-dated crypto prediction markets and their settlement feeds — BTC, ETH, SOL

Programme: internal quantitative research — no client capital at any stage

Status: closed on a pre-registered decision gate

Published: September 2026 · **Licence:** Dubai Department of Economy and Tourism 964356 · **Unified licence:** BL1484

PART ONE

In plain English

No mathematics in this section. It is about an opportunity that was real, visible and repeatable — and already over by the time we could see it existed.

1. The trade

On **Polymarket** you can bet on whether Bitcoin will be higher or lower five minutes from now. Meanwhile, on **Binance**, Bitcoin is trading right there in front of you.

So: the bet says Bitcoin finishes up. Binance says Bitcoin is already down. One of them is wrong, and there is a five-minute window in which to be right about which.

That is not a subtle idea. It is the most obvious trade in the market, which is exactly why it needs checking rather than assuming — if it were free money it would not be sitting there.

A gap on a screen can be manufactured by several dull mechanisms: one feed being slower than the other, a price being averaged over a window, somebody pulling their order the instant you reach for it. Every one of those produces something that looks precisely like an edge and cannot be captured.

2. What we built

We recorded three of these markets — Bitcoin, Ethereum and Solana — continuously and in parallel, for **twenty-five days**. Every price tick, the full order book, and every feed involved in settling the bet, each stamped the microsecond it arrived on our machine rather than when it claimed to have been sent.

That distinction is the whole study. A timestamp written by the sender tells you when they say they sent it. A timestamp written by the receiver tells you when you could actually have acted.

Before any of it ran we wrote down six conditions the result would have to meet, with the numbers fixed, and a rule that failing any one ends the project. That document was dated and never edited.

3. What we found

The recording worked. Coverage was 99.5%, and all **21,269 completed bets** could be rebuilt from our own tape and matched against the official outcome with **zero mismatches**. The data was never the problem.

The problem was condition two, which required the opportunity to last at least three seconds — long enough for an order to actually arrive.



1.0 s

how long the typical opportunity lasted

1.6 s

how long the settlement feed took to tell us it existed

0.15 s

how long the Binance feed took, on the same clock

Read those in order. **The typical opportunity is over before the feed that defines it reports that it began.**

Not slightly too slow. Not a badly placed machine. Not too small a sample. The information needed to identify the trade arrives, on average, six hundred milliseconds after the trade has stopped existing.

4. The question this settled

We had written an open question into the project plan before we could answer it: *is the settlement feed's delay caused by transport — how long the data takes to reach us — or by cadence, meaning the feed simply updates infrequently?*

The difference is enormous. Cadence is survivable: if a feed publishes every two seconds, everyone waits the same two seconds and you plan around it. Transport is not: if the data takes 1.6 seconds to reach *you specifically*, anyone closer to the source has already traded.

It was transport — on all three feeds and all three markets. That is a property of the wire, not of how much data we collected, so more recording cannot change it. Which is why this is a closure rather than a pause.

5. The condition that passed, and why we ignore it

One condition did pass, by an enormous margin: it required at least 100 qualifying opportunities and found **30,420** on Bitcoin alone.

We record it as **meaningless**, because it could not have failed. The rule that opened an opportunity used a wider threshold than the fee it was tested against, so every opportunity that opened cleared the profitability test automatically. The condition was testing its own definition.

We flagged this two weeks before the close and repeat it here so nobody can later cite thirty thousand passing events as encouragement. **A test that cannot fail has not been passed.**

6. What it cost, and one process failure

No capital, at any stage. No trading code was ever written. The cost was recording time and disk space.

One failure is worth publishing. The project ran **eleven days past its own deadline**. The review date fell, the recorders were still running when someone noticed during an unrelated disk audit, and they had quietly written another seven gigabytes.

The cause was mundane: the project was not listed in the summary we read at the start of each session, so when its notes aged out, nothing prompted anyone to stop it. Recording past a pre-registered endpoint is a genuine hazard — it lets you pick your stopping point after seeing the data. It did not

corrupt this result, because the condition fails at every window length. It could easily have corrupted a different one.

PART TWO

The technical record

Capture quality, the pre-registered gate as scored, and the transport measurement that decided it.

1. Subject and instrumentation

Instrument	Polymarket 5-minute Up/Down binary markets — BTC (primary), ETH, SOL
Recording window	2026-08-15 19:19Z → 2026-09-09 19:19Z = 25 days (requirement ≥ 14 days, met 1.8×)
Feeds recorded	market book and tape; binance reference; settlement feeds chainlink, twap60, twap30
Stamping	on receipt, microsecond resolution, against a clock disciplined to the same reference as the latency measurement
Gate	6 criteria, thresholds frozen 2026-08-14, never re-tuned

2. Capture quality

Recording integrity across the full window. Capture was never the constraint.

METRIC	BTC	ETH	SOL
Ticks	2,148,515	2,148,592	2,148,672
Coverage @ 1 Hz	99.47%	99.47%	99.48%
Priced (settlement-stamped)	99.2%	99.2%	99.2%
Complementary pair sum off by > 2c	0.031%	0.052%	0.050%
Settlement rule reconstruction	100.0%	100.0%	100.0%
Book present	78.9%	79.9%	80.2%

6,445,779 ticks; 21,269 resolved windows; zero rule-reconstruction mismatches on any asset. Book presence near 79% independently reconfirms a book-withdrawal effect found in an earlier audit — the ladder is frequently not there to be hit, which is a second, separate obstacle to the same trade.



3. The gate as scored

#	CRITERION	BTC (PRIMARY)	ETH	SOL	RESULT
1	≥ 100 episodes, net edge ≥ 0.05	30,420	11,408	7,142	PASS — <i>vacuous, see §5</i>
2	Median episode duration ≥ 3 s	1 s	1 s	1 s	FAIL
3	$\geq \$250$ fillable after walking the ladder	med \$250	med \$250	med \$250	PASS
4	Direction, 98.33% lower bound	—	—	—	moot
5	Null / blind-Up / blind-Down controls	—	—	—	moot
6	Settlement rule resolved	7088/7088	7093/7093	7088/7088	PASS

The gate states that if any one criterion fails, the next phase is not authorised. Criterion 2 fails on all three assets, and the pre-registration additionally names BTC the primary asset, so a pass elsewhere would authorise nothing regardless.

Criteria 4 and 5 were **deliberately not computed**. One failure ends the gate, and computing the remaining criteria after the fact could only invite re-reading a closed result. The 98.33% bound in criterion 4 is itself the pre-declared correction for testing three assets simultaneously, fixed in advance rather than chosen afterwards.

4. The transport measurement

Feed delay measured against each feed's own publication timestamp, on a common clock.

FEED	MEDIAN DELAY	HEALTH VERDICT
chainlink (settlement)	1,627 ms	TRANSPORT
twap60 (settlement)	1,702 ms	TRANSPORT
twap30 (settlement)	1,704 ms	TRANSPORT
binance (execution reference)	152 ms	healthy

Episode durations: median **1,000 ms** on all three assets; p90 6 s (BTC), 3 s (ETH), 4 s (SOL). Set against a settlement feed 1.6 s behind, **the median episode is over before the settlement feed reports it exists**.

The verdict TRANSPORT is the discriminating output. It resolves an open question carried in the project plan since before the data existed — whether the observed lag was transport delay or oracle round cadence — empirically, on 6.4 M ticks, in the unfavourable direction. Transport delay is a property of the relay, not of the sample, so additional recording cannot change it. That is why this is a closure rather than a pause.

5. The vacuous pass, recorded as such

Criterion 1 required at least 100 episodes carrying a net edge of at least 0.05, and returned tens of thousands. It is recorded as vacuous because it could not have failed: **episodes open at a gap threshold**

of 0.12 while the fee peaks at 0.0175, so every episode that opens clears the net-edge test by construction. The criterion tested its own opening condition.

This was flagged in an audit before the close and is repeated in the closed record so it can never be cited as support. A test that cannot fail has not been passed, and a large number produced by one is not evidence of anything.

6. Disposition

Recorders stopped by graceful signal with no tape lost, deleted from the process manager and removed from the saved process list so they cannot resurrect on reboot. 16 GB of raw tape deleted; 12 KB of results preserved because this record cites them. No trading code was ever written and no capital was involved at any stage.

Do not propose a re-run with a longer window, a different asset set, or a looser persistence threshold. The transport lag is the binding constraint and none of those touch it.

7. Reusable constants

CONSTANT	VALUE
Settlement-feed transport delay	1,627–1,704 ms median across three oracle feeds
Execution-feed delay, same clock	152 ms
Median episode duration	1,000 ms (p90 3–6 s)
Book presence	~79% – the ladder is often absent when needed

The first two are properties of the venues and the relay path, not of any strategy. Any proposal requiring a reaction to a settlement-feed update inside two seconds can be declined on these figures without being tested.

8. What this does not close

It closes taking a short-dated position that depends on reacting to the settlement feed. It says nothing about strategies that do not require that reaction — a resting quote, for instance, is not subject to the transport constraint in the same way, because it does not need to arrive in time. That has not been measured here.

Scope. Thanaaya Technologies is a Dubai Department of Economy and Tourism licensed company (licence 964356). This paper describes an internal research programme. No client capital was involved at any stage, no strategy or signal design is published, and nothing here constitutes investment advice or an offer of any service relating to trading. Enquiries: info@thanaaya.com · thanaaya.com