



Markets That Do Not Move

A news-ingestion engine matched headlines to prediction markets correctly, and then over half of the resulting positions did not move at all.

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: news-driven trading of non-crypto prediction markets

Programme: internal quantitative research — paper trading only, no capital at any stage

Status: version 1 closed on execution-realistic return evidence

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

PART ONE

In plain English

No mathematics in this section. We built a system to trade news, it identified the news correctly, and the market did not care.

1. The idea

On **Polymarket** you can buy a contract that pays out if something happens — "will China invade Taiwan", "will this candidate win". Prices sit between 0 and 100 cents and read as a probability.

If news arrives making an event more likely, the price should rise. So: read the news faster than the market does, and buy before it moves.

We built exactly that. Ingest news continuously, work out which contract each headline actually bears on, verify the connection with a second independent model, and record what would have been traded. It never touched real money — everything below is paper trading marked against real prices.

2. What happened

The pipeline worked. It read the news, matched headlines to contracts, and opened positions. Across seven recording runs it produced **31 trades** that ran a full 24 hours.

Then we scored them honestly, and the result was stranger than being wrong about direction.

16 of 31

where the price 24 hours later was *identical* — not a single cent of movement

0.31c

average movement across all 31, in either direction

1.00c

what it costs to buy and sell one contract

Over half the positions did not move at all. Not "moved slightly against us" — the price a day later was the same number.

THE FINDING THAT CLOSED IT

The signal is about **one third the size of the cost of trading it**, and the typical trade does not move. You cannot fix that with better direction, because there is no movement to be directionally right about.

Scored at prices you could actually get, 3 of the 31 made money. A **10% win rate**.



3. Where the flatness comes from

The obvious explanation is that we were trading obscure junk — thin contracts nobody watches, where nothing happens because nobody is there. We checked. That is not it.

Almost every trade landed in a contract resolving **more than 120 days away**, and long-dated contracts turn out to be structurally unresponsive to news. Something that shifts the odds of an event eight months out barely moves the price, because the event was always uncertain and still is.

The clearest case in the set is a large geopolitical contract with **\$684,000 of liquidity**, a spread of a tenth of a cent, and \$48,900 of daily volume. On every measure of quality it is an excellent market. It is also completely flat. **Flatness is a property of long-dated contracts, not thin ones.**

We also checked whether trading volume could separate the movers from the dead, since that is the obvious next filter. It cannot: the deadest positions included some of the highest-volume contracts in the set, and some movers were among the quietest. The same contract appears on both lists at different times.

4. Two ways we were scoring ourselves wrongly

Both flattered the result. Both were corrected before we closed.

Pricing at the middle of the spread

It is conventional to value a position halfway between the buy and sell price. On a market with a one-cent spread that is a fiction — you cannot trade at the middle. Scoring at prices actually available took the win rate from 23% down to **10%**, and the total from mildly negative to firmly negative.

The lesson is blunt: cross the spread twice, once in and once out. Anything else scores you against prices nobody offered.

Reporting gains as percentages

On a market where prices live between 0 and 100 cents, percentages mislead. A move from four cents to two is a two-cent move — but it prints as **minus forty-five per cent**, and dominates any average it lands in. One position had its significance overstated **twenty-four fold** this way. We now report cents, not percentages.

5. What we would keep

This was not an engineering failure. Reconciliation between what the system thought it held and what it actually held was exact. The daemon ran for days unattended. The news ingestion did what it was built to do.

What failed was the premise, and it failed for a reason worth more than the system: **long-dated prediction markets do not reprice on news**. That is a durable, checkable statement about a market structure, and it rules out a whole family of ideas that look attractive from outside.

6. Where it left us

The engine was stopped and archived with everything preserved. Live trading was requested and refused, and the refusal stands — the evidence that closed the paper record is the evidence that would have to be overturned first.

A separate finding from the final run, about why the system's own quality filter was rejecting nearly everything, is in the companion paper *The Matcher, Not the Verifier*. For anyone building something similar, that one is the more useful of the two.

PART TWO

The technical record

Return evidence at both marking methods, the flatness diagnosis, and the measurement defects corrected before close.

1. Return evidence — 24-hour horizon, all runs, n = 31

The same book under two marking conventions. The difference between the rows is not opinion; it is whether the prices used were available.

MARKING METHOD	MEAN	MEDIAN	WIN RATE	NET
Mid-to-mid (as originally run)	-0.089c	+0.000c	23%	-2.75c
Real bid/ask (execution-realistic)	-0.747c	-0.500c	10%	-23.15c

DISTRIBUTION	VALUE
Median full round-trip spread	1.00c
Mean absolute 24-hour move	0.308c
Profitable at real fills	3 of 31
Moved exactly 0.00c	16 of 31 (52%)
Moved $\geq 1.00c$	3 of 31 (10%)

Mean absolute move 0.308c against a 1.00c round-trip cost. The signal is approximately one third the size of the friction, and the modal outcome is no movement whatsoever. No improvement in directional accuracy addresses this, because the quantity being predicted is absent rather than mis-signed.

2. The flatness is structural, not a liquidity artefact

2.1 Maturity dominates

24 of 26 realised trades sat in markets resolving more than **120 days** out. Within that group the mean absolute 24-hour move was **0.175c**, and **0 of 24 cleared a single cent**. Long-dated contracts do not reprice on news at a magnitude that survives the spread.

2.2 Liquidity does not explain it

A large geopolitical market in the set clears every quality threshold — \$684k liquidity, 0.001 spread, \$48.9k 24-hour volume — and is completely flat over the horizon. Flatness is a property of the contract's maturity, not of book depth.

2.3 Volume does not discriminate

Dead positions ran to 24-hour volumes of 263,679; movers ran as low as 259. The same token appears in both categories on different occasions. Volume is not a usable filter for separating responsive markets from unresponsive ones, and gating on it would have removed movers while retaining dead tape.

3. Measurement defects, found and corrected

DEFECT	EFFECT	CORRECTION
Mid-to-mid marking	Flattered the book from a 10% to a 23% win rate, and the net from -23.15c to -2.75c	Cross the spread twice; mark at prices actually available
Percent-of-entry returns	Overstated one tail case 23.7×; printed -44.83% on a 2c move	Report price points, never percentages, on bounded contracts
Model self-reported confidence	0.90+ on 77 of 100 trades, of which a second model rejected 75	Never size on a model's own confidence score

The first two are the reason this paper exists in its current form. A book that reads as mildly disappointing under mid-to-mid marking reads as decisively closed under execution-realistic marking, and the difference is entirely in the measurement rather than in the trading.

4. What was sound

Position reconciliation between the engine's internal state and the recorded book was exact across the full set. The daemon ran multi-day recording runs without intervention. Ingestion and routing performed as designed. **The failure is located in the premise, not in the implementation** — which is worth stating precisely, because "the system did not work" and "the system worked and the idea was wrong" call for entirely different responses.

5. Disposition

Version 1 was stopped, frozen and archived in full. All 31 trades reached final marked state before the grader was stopped, so the dataset cannot drift. Live trading was requested and **refused**; the refusal stands and nothing was ever armed. No capital was committed at any stage.

6. Settled findings carried forward

- **Long-dated prediction markets do not reprice on news** at a magnitude that clears the spread.
- **Liquidity is not the explanation for flatness** — deep, tight, high-volume long-dated markets are flat too.
- **Volume does not discriminate movers from dead tape** and is not a usable gate.
- **Mid-to-mid marking lies** on any market with a meaningful spread.
- **Percent-of-entry lies** on bounded contracts and should not be reported.
- **A model's self-reported confidence carries no signal** and must never drive position size.

Any successor that proposes trading news into long-dated contracts has to overturn the first of these before anything else in its design matters.

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