THANAAYA TECHNOLOGIES
RESEARCH NOTE

Impossible Before We Started

The cost floor, the accuracy it implies and the latency budget — three constraints that can be settled with arithmetic before any research begins.

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: venue cost floors, required hit rates and measured latency budget
Programme: internal quantitative research — no client capital at any stage
Status: measured; constants carry forward to all later work
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 a question that can be answered with arithmetic before any research starts, and how much work gets done because nobody asks it.

1. The question nobody asks first

Before asking whether you can predict where **SOL perpetual futures** are going, it is worth asking whether the prediction would be worth anything if you could.

Those sound like the same question. They are not, and the second is far easier. It needs no model, no signal and no data science — just two numbers: how much the price typically moves over the time you intend to hold it, and how much it costs to get in and out.

If the second is bigger than the first, the strategy cannot work. Not "is unlikely to" — *cannot*, in the sense that it would need you to be right more than one hundred per cent of the time.

0.074%	**0.106%**	**121.5%**
average amount SOL moves in five minutes	cost of a fast round trip on Binance	hit rate you would need to break even

That comparison eliminates an entire family of strategies — anything holding for a few minutes and paying to trade immediately — no matter how good the signal is. It was knowable before we recorded a single trade.

2. What "being right enough" actually requires

If your signal is right some fraction of the time and wrong the rest, and the moves are about the same size either way, your average gain per trade is the size of the move multiplied by how much better than a coin flip you are. To cover a fixed cost you need a hit rate that depends only on the ratio between the cost and the move.

Which gives a table anyone can compute in advance, before writing any code:

How often you would have to be right, on SOL, to break even. Red entries are impossible.

IF YOU HOLD FOR	SOL MOVES, ON AVERAGE	PATIENT ORDERS, CHEAPEST FEES	TRADING IMMEDIATELY
5 minutes	0.074%	74.4%	121.5%
26 minutes	0.156%	61.5%	83.8%
1 hour 44	0.329%	55.5%	66.0%
6 hours 56	0.607%	53.0%	58.7%

The best hit rate we ever actually observed at five minutes, across 854 candidate rules, was **55.9%** — and it did not survive a test against scrambled data.

The gap between 55.9% and the 74.4% you would need is the real result. It is not a near miss to be closed with a better model. It is the difference between a coin that lands heads slightly too often to be interesting and a genuinely predictive instrument.

3. The other thing that costs money: being late

There is a second constraint, less obvious. Information in a market has a shelf life, and yours starts expiring the moment it leaves the exchange.

We measured how long Binance's data takes to reach our machine, using the exchange's own timestamps against a clock synchronised to within 31 *millionths* of a second. The answer: **137.8 milliseconds**. By the time an order could travel back, roughly **276 milliseconds** have passed.

Then we measured how fast the predictive content decays. It falls by a factor of **twenty-one** across those same 300 milliseconds.

THE CONSEQUENCE

The signal is real — at the instant it appears it is 170 times its own control. The problem is that **it is not ours**. About 95% of it has expired before an order from this machine could reach the exchange. What is left is worth **0.0097%** against a cost of 0.106% — roughly a tenth of what it costs to collect.

4. Why this is good news, sort of

A conclusion like this is far stronger than "our model was not good enough", because it names the conditions under which it would change. Three things could move, and all three are measurable rather than arguable:

- **Be faster.** To make the arithmetic work while trading immediately, the round trip would have to drop from ~276 ms to under ~67 ms. That is a question of where the machine sits, not how clever it is.
- **Pay less.** Fees are a published schedule tied to volume. At high enough volume they fall to nearly nothing.
- **Hold longer.** The move grows with time while the cost stays fixed — at the price of no longer being a fast strategy at all.

What none of them changes is the accuracy requirement, which is the one that actually failed.

5. The habit this produced

We now compute the cost floor and the required hit rate *before* starting a study, and write them into the specification next to the hypothesis. It takes about ten minutes and it has already stopped one line of work from beginning.

It is an uncomfortable discipline, because it occasionally kills an idea you were looking forward to. That is the entire value of it.

PART TWO

The technical record

Fee floors, required hit rates by horizon, the measured latency budget and the signal-decay constant, as recorded on SOL perpetual futures.

1. Cost floors

The venue's published VIP0 taker rate is inconsistently reported in secondary sources (0.04% legacy against 0.05% current), so every result is scored against all five candidate floors rather than one. A round trip is two fills; a taker also pays the spread it crosses, a maker does not. Median spread is **0.548 bps**, one tick at the prevailing price.

FLOOR	COMPOSITION	BPS / ROUND TRIP
Maker 0.018% (with exchange-token discount)	2 × 1.80	3.60
Maker 0.020%	2 × 2.00	4.00
Taker 0.040% + spread	2 × 4.00 + 0.548	8.55
Taker 0.045% + spread	2 × 4.50 + 0.548	9.55
Taker 0.050% + spread	2 × 5.00 + 0.548	10.55

Verified rate at the time of writing: **maker 0.0200% / taker 0.0500%**. No conclusion here depends on which figure is correct, because the cheapest conceivable floor — 3.60 bps — is the one used for the comparisons below.

2. Required directional accuracy

With edge ≈ $(2p - 1) \times E|\text{move}|$, the hit rate required to clear a floor is $p = \frac{1}{2}(1 + \text{floor} \div E|\text{move}|)$. Measured unconditional move sizes on the registered window:

Bold entries require a hit rate above 100% and are therefore arithmetically impossible.

HORIZON	E|MOVE| (BPS)	MAKER 3.60	MAKER 4.00	TAKER 8.55	TAKER 10.55
5 min	7.37	74.4%	77.1%	108.0%	121.5%
10 min	10.29	67.5%	69.4%	91.5%	101.2%
26 min	15.61	61.5%	62.8%	77.4%	83.8%
52 min	22.44	58.0%	58.9%	69.1%	73.5%
104 min	32.90	55.5%	56.1%	63.0%	66.0%
208 min	46.50	53.9%	54.3%	59.2%	61.3%
416 min	60.68	53.0%	53.3%	57.0%	58.7%

Taker trading at the five-minute scale cannot work at any signal quality. Conditioning does not rescue it: restricting to the top 2% of the strongest magnitude feature raises E|move| only to 9.46 bps, still requiring a 105.7% hit rate. This is arithmetic, not an empirical result, and no model changes it.

Against the cheapest maker floor at the same horizon the requirement is **69.0%**. The best hit rate observed anywhere at that horizon across an 854-cell search was **55.9%**, and it sat inside the noise distribution of that search.

3. The latency budget

MEASUREMENT	VALUE
Feed latency, exchange stamp → host	median 137.8 ms · p90 150.1 ms · p99 167.8 ms
Sample	1,719,883 depth events across 51 hours
Clock trust	NTP-locked, 31 µs from reference — true one-way latency, not skew
Stability	hourly median 130.3–147.4 ms, sd 5.1 ms
Round-trip reaction	~276 ms; specifications evaluate at a 3-step, 300 ms entry lag

The clock figure matters more than it looks. Without it, a latency measurement is indistinguishable from clock skew between the venue and the host, and the entire budget becomes unfalsifiable.

4. Signal decay

The incremental predictive contribution of the strongest feature block falls from **0.005437 to 0.000254 across 300 ms** — a factor of 21.4, implying an exponential time constant of **τ ≈ 97.9 ms**. At the primary 300 ms lag, the strongest decile spread is **0.97 bps** against a 10.55 bps taker round trip.

Extrapolating the measured decile spread back along that decay gives the latency each floor would require:

Extrapolation from two measured points under an exponential-decay assumption. An order of magnitude, not a threshold to engineer against.

FLOOR	REQUIRED UPLIFT	MAX ROUND TRIP	≈ ONE-WAY
Maker 3.60 bps	3.71×	171.6 ms	85.8 ms
Maker 4.00 bps	4.12×	161.3 ms	80.6 ms
Taker 8.55 bps	8.81×	86.9 ms	43.4 ms
Taker 10.55 bps	10.88×	66.3 ms	33.1 ms

This host sits at 137.8 ms one-way. Clearing the maker floor needs a 1.6× improvement; clearing the taker floor needs about 4×. Both are purchasable by colocation in the venue's own region; neither is achievable from where this ran.

5. What the decay figure actually says

It is worth being precise, because this is the most misreadable result in the study. The signal is **real**: the incremental contribution at zero lag is 170 times its own control. The problem is not that it is spurious. The problem is that **it is not ours** — 95% of it has expired before an order from this host could reach the exchange, and what remains is worth roughly a tenth of the cost of acting on it.

Two separate feature blocks showed the same 300 ms cliff independently, and in one of them the edge *inverts* once the round trip completes: +0.299 bps at the decision instant becomes −0.277 bps 300 ms later on the test split. A strategy built on it would not merely fail to pay; it would pay to be wrong.

6. Reusable constants

CONSTANT	VALUE
Feed latency, this host	137.8 ms one-way median; ~276 ms round trip; NTP-verified to 31 μs
Signal decay constant	τ ≈ 97.9 ms (0.005437 → 0.000254 over 300 ms)
Median spread	0.548 bps
Cheapest conceivable round trip	3.60 bps (maker, with token discount)
Surviving edge at primary lag	0.97 bps — about one tenth of the taker cost floor

The first two are properties of the infrastructure rather than of any strategy, so they carry to every subsequent study run from the same place. Any proposal whose edge decays faster than τ, or which requires acting inside 276 ms, can be declined without being tested.

7. What this does not close

Every figure here assumes crossing the spread. The passive side — resting an order and being traded against — has different arithmetic, and the quantity that decides it is the adverse-selection cost:

conditional on a resting order being filled, what does the price do next? That has not been measured, and nothing in this paper licenses a claim about it in either direction.
