

# Touchline

## *NAV Methodology & Tier Variations*

An inventory-managed market-making vault on fifteen-minute crypto binaries — NAV construction, three risk-tier variations, and capacity limits.

**Produced by Itô Markets**

Touchline · Quantitative Research · Version 1.0 · Q2 2026

# Contents

<b>1</b>	<b>Executive Summary</b>	2
<b>2</b>	<b>The Strategy</b>	2
2.1	Instrument & Market	2
2.2	Quoting & Inventory	2
2.3	Fill Model	3
2.4	Edge & Significance	3
<b>3</b>	<b>NAV Methodology</b>	3
3.1	The Vault & Share Price	3
3.2	NAV Definition	3
3.3	Initialization & Accrual	3
<b>4</b>	<b>Tier Variations</b>	3
4.1	Same Engine, Different Cushion	3
4.2	The Three Tiers	4
4.3	Note on TL-3	4
<b>5</b>	<b>NAV by Tier — Results</b>	4
<b>6</b>	<b>Capacity Limits &amp; Maximum Vault Size</b>	5
6.1	The Working-Capital Cap	5
6.2	Equal Dollars at Full Capacity	5
6.3	Measured Capacity	5
<b>7</b>	<b>Operating Notes</b>	6
<b>8</b>	<b>Risk Factors &amp; Limitations</b>	6
8.1	Significance is borderline	6
8.2	Drawdown is understated by the short sample	6
8.3	Capacity caps the working unit, not the cushion	6
8.4	Do not full-reinvest (volatility drag)	6
8.5	Live execution will differ	6
<b>9</b>	<b>Governance</b>	6
	<b>Changelog</b>	7
	<b>Approvals</b>	7

# 1 Executive Summary

**Touchline** is a tokenized, inventory-managed market-making vault built on Polymarket fifteen-minute BTC and ETH "Up or Down" binary markets. One share is a pro-rata claim on everything the vault holds, and its price is the net asset value (NAV) per share — depositors buy and redeem at NAV, not at a price they pick. This document specifies how the strategy works, how the share price is constructed, how the three risk-tier variations differ, and where the capacity limits sit.

The engine quotes both sides of each fifteen-minute window, manages inventory with a soft band, and flattens at the mid before the discontinuous 0/1 settlement. The edge is small and substantially rebate-driven: mean return is approximately **+\$1.17 per window**, the window win rate is approximately **54.8%**, and the window-clustered *t*-statistic is approximately **1.94** — approaching, but not yet past, the conventional 2.0 bar.

The three variations are not different position sizes. They are the *same* working unit inside three vault sizes — i.e. how many basis points of the vault are committed to risk, with the remainder held as cushion. The same settled profit is packaged at three different risk dials: more cushion means a smaller percentage return and a far smaller drawdown; less cushion means a larger percentage return and a larger drawdown.

**Snapshot as of 2026-06-15.** 431 settled windows over 2.25 days; total **+\$504.25** (trading +\$344.58 + maker rebate +\$159.67). Every figure in this document is derived deterministically from the real settled per-window profit-and-loss of the live run.

## 2 The Strategy

### 2.1 Instrument & Market

Touchline trades Polymarket fifteen-minute BTC and ETH up/down binary contracts. Each is a cash-or-nothing digital option that pays \$1 if the underlying is up at the close of the window and \$0 otherwise; its quoted price lies in the unit interval and equals the implied probability of the in-the-money event. Settlement is mechanical and oracle-free, and the thin, penny-tick books are the regime in which a patient passive maker can harvest spread and the venue rebate.

### 2.2 Quoting & Inventory

Per fifteen-minute window the engine quotes both YES and NO at the touch during a defined window before expiry, caps risk at roughly **\$30 per side** — about **\$60 of working capital per window**, the "working unit" — and flattens all open inventory at the mid shortly before close rather than holding to the 0/1 settlement. Inventory is managed with an Avellaneda–Stoikov style reservation-price skew expressed in cents; the maker joins the touch but cuts size on the inventory-adding and toxic side. The specific timing and coefficient parameters are proprietary.

## 2.3 Fill Model

The fill model is flow-and-queue gated: a fill occurs only when a real trade prints through the engine's queue position, and is capped by the real trade size. Fills the tape did not produce are never assumed. Maker entries are slippage-free by construction; the flatten exit is marked at the mid (on Polymarket the associated slippage is small — cents per window — and immaterial to the result).

## 2.4 Edge & Significance

The income is the bid/ask spread plus the maker rebate (a function of fill price, maximised at-the-money). On the current sample the edge is small and mostly rebate-carried: mean  $\approx$  **+\$1.17/window**, win rate  $\approx$  **54.8%**, window-clustered  $t \approx$  **1.94**. Significance is borderline and is addressed in Section 8.

# 3 NAV Methodology

## 3.1 The Vault & Share Price

Touchline is a tokenized vault (ERC-4626-style). One share is a pro-rata claim on everything the vault holds; the share price is the NAV per share, and deposits and redemptions transact at NAV. Open inventory is marked to the current mid — not only to settled cash — so that depositors are fair to each other and no one captures or absorbs in-flight profit-and-loss they did not fund.

## 3.2 NAV Definition

NAV = idle cash +  $\Sigma$ (open inventory shares  $\times$  current mid) + accrued rebate – fees  
 NAV/share = NAV  $\div$  shares outstanding

## 3.3 Initialization & Accrual

NAV/share is initialized at **\$100.00**, so shares outstanding = AUM / 100. The strategy is capacity-capped — it deploys the same  $\sim$ \$60 working unit per window regardless of vault cash — so profit accrues as cash and the share price moves additively:

$$NAV/share(t) = \$100 \times ( AUM + cumulative P\&L(t) ) / AUM$$

The full bankroll is deliberately *not* reinvested into one larger window (see the volatility-drag note in Section 8); the working size is held constant and the cushion absorbs the variance. The cumulative P&L series is the real settled (trading + rebate) series, re-derivable from the run's event log.

# 4 Tier Variations

## 4.1 Same Engine, Different Cushion

The variations are not different position sizes. They are the same  $\sim$ \$60 working unit inside three vault sizes — i.e. how many basis points of the vault are committed to risk. The remainder is cushion: idle cash that damps the percentage swings. Committed basis points =  $\$60 / AUM \times 10,000$ .

## 4.2 The Three Tiers

Tier	Vault (AUM)	Working unit	Committed	Cushion (idle)
TL-1	\$10,000	~\$60/window	60 bps	~99.4%
TL-2	\$1,000	~\$60/window	600 bps	~94%
TL-3	\$100	~\$60/window	6,000 bps	~40%

## 4.3 Note on TL-3


A "60,000 bps" tier would imply a \$10 vault holding \$60 of risk — 6× leverage, which is not fundable. TL-3 is therefore **6,000 bps** (a \$100 vault, 60% committed): the top of the unlevered ladder. A levered tier is possible but is a different, far riskier product and is out of scope here.

## 5 NAV by Tier — Results

The same **+\$504.25** of settled profit appears in all three tiers — the tier only changes how it is packaged.

Tier	AUM	NAV/share start	NAV/share end	Return	Max drawdown
TL-1 — 60 bps	\$10,000	\$100.00	\$105.04	+5.0%	-1.4%
TL-2 — 600 bps	\$1,000	\$100.00	\$150.42	+50.4%	-10.0%
TL-3 — 6,000 bps	\$100	\$100.00	\$604.25	+504.2%	-24.95%

Worst absolute dip over the run:  $-\$151.00$  (peak-to-trough of cumulative working-capital P&L — identical dollars for every tier). The percentage columns are a function of the cushion, not of more skill or money: the absolute profit is identical; the tier is a risk dial.



Touchline NAV/share by tier. The same settled cumulative profit-and-loss, expressed as NAV/share under each tier's denominator: one underlying curve at three vertical scales — and therefore three different percentage returns and drawdowns from identical dollars.

Mind the denominators. The max-drawdown column is the realized NAV/share drop off its own peak. As a fraction of initial AUM the  $-\$151$  dip is ~1.5% (TL-1), ~15% (TL-2), and ~151% (TL-3) — but a \$100 seed cannot literally lose 151%: on TL-3 that dip only materializes late in the run, after the share price has already climbed, which is why its realized drawdown is -24.95%. Read the

percentage correctly — more cushion (TL-1) means a smaller percentage and far smaller drawdown; less cushion (TL-3) means a bigger percentage and bigger drawdown.

## 6 Capacity Limits & Maximum Vault Size

### 6.1 The Working-Capital Cap

The engine has a hard cap on money deployed per window: a maximum of **\$600 of working capital per fifteen-minute window**. Each tier's committed basis points are fixed, so a tier takes deposits — working capital scaling with AUM to hold the bps constant — until working capital hits the \$600/window cap. That defines the tier's maximum vault size:  $\text{max AUM} = \$600 / (\text{committed bps} / 10,000)$ .

Tier	Committed	Max vault (AUM)	Working cap	Full-capacity profit (2.25d)
TL-1	60 bps	\$100,000	\$600/window	≈ +\$5,042
TL-2	600 bps	\$10,000	\$600/window	≈ +\$5,042
TL-3	6,000 bps	\$1,000	\$600/window	≈ +\$5,042

### 6.2 Equal Dollars at Full Capacity

If the \$600/window cap proves fillable, all three tiers deploy the same \$600/window at full capacity and earn the same absolute dollars — roughly **10× the base run, ≈ +\$5,042 over 2.25 days**. They differ only in the AUM that profit sits on — exactly the percentage column. Each vault hard-closes to deposits at its maximum AUM, since further deposits would only dilute the return.

### 6.3 Measured Capacity

A separate size-ladder runner reports deployed capital per window by cap. This is very early data (only a handful of windows, noisy):

Per-window cap	Deployed / window
\$100	~\$62
\$200	~\$113
\$500	~\$182
\$1,000	~\$368
\$3,000	~\$629

*Deployment is still rising at the \$3,000 cap (≈ \$629/window), so the books may absorb at or above the \$600/window design cap — but this is not yet established. This is far too few windows to pin the real ceiling; the \$600/window cap is a design assumption, and per-dollar edge can erode as the engine takes more of the book (adverse selection, market impact). The maximum-AUM figures move with the validated cap.*

## 7 Operating Notes

- **The cushion must earn the risk-free rate.** TL-1 is ~99% idle cash; held in T-bills or staked stables (~4–5%), that yield is most of a high-cushion depositor's return. Left idle, it is pure drag.
- **Deposits and redemptions settle at window boundaries** (epoch / lock-up), never mid-position, so flows do not whipsaw the maker or let anyone arbitrage in-flight profit-and-loss.
- **A high-water mark plus entry/exit fees** preserve fairness when capital scales the book up or down.

## 8 Risk Factors & Limitations

### 8.1 Significance is borderline

The sample is 2.25 days with a window-clustered  $t \approx 1.94$  (the conventional bar is  $\geq 2.0$ ). The realized profit-and-loss is settled, but a longer sample is required to confirm the effect rather than rule out part of it as sampling variation.

### 8.2 Drawdown is understated by the short sample

The  $-\$151$  worst dip is the deepest in 2.25 days; a longer run will likely print deeper. Size the cushion for two-to-three times that (roughly  $-\$302$  to  $-\$453$ ).

### 8.3 Capacity caps the working unit, not the cushion

Adding cushion (moving from TL-3 toward TL-1) is free and safe; growing the working unit is capped at  $\$600$ /window — a design assumption that the early, noisy ladder data does not yet confirm and may sit below true saturation. The real ceiling and any edge erosion at size must be measured before the maximum-AUM figures can be relied upon.

### 8.4 Do not full-reinvest (volatility drag)

Compounding the whole bankroll into each window turns the profit into a loss, because per-window volatility (~21% on the  $\$60$  unit) exceeds the per-window edge. The constant-working-size design of Section 3 is deliberate.

### 8.5 Live execution will differ

Queue priority, market impact, latency, and partial fills all affect realized results, and per-dollar edge can erode as the engine takes more of the book.

## 9 Governance

Itô Markets is the research sponsor and calculation authority for Touchline. All methodology decisions — NAV construction, tier definitions, capacity caps, and fee parameters — are the responsibility of the nowcast\_v2 Research Committee.

Changes to this methodology require a written proposal, committee review, consensus approval, and re-publication of the updated specification. The NAV series and every figure in this document are re-derivable from the run's settled event log, and the methodology is maintained on the grounds of measurement integrity alone.

## Changelog

Version	Date	Changes
1.0	Q2 2026	Initial publication. Touchline NAV methodology and the three unlevered tier variations (TL-1 / TL-2 / TL-3), with capacity limits and the measured size-ladder. Consolidates and supersedes the standalone fifteen-minute prediction-market making methodology.

## Approvals

Role	Name	Signature	Date
Head of Research			
Chief Executive Officer			