

Polymarket Momentum Strategy

Crypto Up/Down lead-lag taker: how it works, where the edge comes from, 5.5 months of paper results graded on real settlement, and whether to continue

Research paper 2 of 2 (paper 1 = Polymarket market overview)

Data: 2026-05-03 to 2026-10-17, 1,444 paper trades

Generated 2026-10-18

Internal research - not investment advice

0 Executive summary

Verdict: the edge was real, and it is now close to zero. Over 1,444 paper trades (2026-05-03 to 2026-10-17), Momentum made **+163.99 \$ per 1-share stake** when every trade is graded on the **real Polymarket outcome** (win rate 67.4% vs 50.9% break-even, ROI on capital staked 22.3%). But 81.1% of that came before 22 June, two days produced 53.1% of all profit, and in Sep-Oct the strategy earned only **+0.037 \$/trade on 2.4 trades/day**. A June fill-realism test showed that at our current 5-second reaction time we fill mostly losers. Polymarket also changed the rules in August: the 5m/15m/4h markets now settle on a Chainlink **TWAP**, which our model ignores.

Metric	Value	Comment
Trades / days	1,444 / 92 trading days of 168	paper, cron-started daily at 00:05
Win rate (real) vs break-even	67.4% vs 50.9%	break-even = average price paid
Net PnL, real outcome	+163.99 \$ per 1-share stake	paper (Binance-graded) +172.15
Gross PnL (no costs)	+245.56	cost drag 29.9%, costs = 1000 bps of capital staked
ROI on capital staked	22.3% real (23.4% paper)	per trade, fully collateralised
Sharpe (daily, annualised)	4.13 real / 3.89 paper	calendar days incl. idle days; tiny stakes, see section 7
Max drawdown	5.97 \$ (1-share)	2026-09-30 to 2026-10-06 = ~12 average stakes
Holding period	median 11 min, p90 2.4 h	held to settlement, no early exit
Trades per day	15.7 per active day, 8.6 per calendar day	Sep-Oct: 2.4 per calendar day
Best / worst day	2026-05-11 +47.36 / 2026-10-02 -2.32	top-2 days = 53.1% of PnL
Label flips (Binance vs real)	123 trades = 8.5%	32% on 5m markets

Recommendation: do not fund the current version. Stop the paper run, or keep it running only as a data feed, and rebuild it in the order given in section 12: (1) a fair-value model that knows about TWAP settlement, (2) a sub-second reaction time over WebSocket, (3) a tighter universe (BTC/ETH/SOL/XRP 15m, at least 5 min to expiry, edge of 0.15 or more). Retest against real outcomes and recorded order books before any money goes in. Even rebuilt, capacity is small (section 10): a few hundred dollars per signal.

1 Which markets it trades

Momentum trades Polymarket's **crypto Up/Down markets**. These are rolling binary markets asking whether BTC, ETH, SOL, XRP or DOGE will be higher at the end of a fixed window than at its start. A new market opens every window, and each outcome token pays \$1 if right and \$0 if wrong. Horizons listed: 5m, 15m, 1h, 4h, daily, weekly, monthly. Fills happened only on **5m (56), 15m (1,171) and 4h (217)**: the 1h/daily/weekly/monthly markets never passed the filters, because the minimum-time-left and price-move gates rarely line up with a large mispricing on long horizons.

Horizon	Markets live (10-18 snapshot)	Median 24h vol / market	Total 24h vol	Settles on
5m	200	\$15	\$36k	Chainlink TWAP (since Aug 2026)
15m	440	\$305	\$964k	Chainlink TWAP (since Aug 2026)
1h	181	\$115	\$503k	Chainlink price
4h	60	\$263	\$97k	Chainlink TWAP
daily	82	\$1,240	\$774k	Chainlink / Binance close
weekly	78	\$7,524	\$1.29M	strike ladders
monthly	138	\$1,589	\$691k	strike ladders

Source: our recorder's metadata snapshot of 1,179 crypto markets on 2026-10-18 (fee_base_rate 0.018 on all of them). Volume figures are Polymarket's own volume24hr field.

Rule change, August 2026. Market descriptions fetched for every traded market show that from **7 Aug 2026** the 5m, 15m and 4h markets resolve on "the time-weighted average price (TWAP) of [the coin], generated by Chainlink, of the time range specified in the title", compared with the price at the start of the window. Before that (last seen 5 Aug) they resolved on the Chainlink price at the end of the window. Momentum's model still assumes end-of-window settlement and grades its own paper trades on the Binance end price.

2 How the strategy works

Idea in one line: Binance prices move first; the Polymarket order book catches up seconds to minutes later. When the coin has moved enough since the window opened that the outcome is now likely, but the Polymarket ask has not yet repriced, buy the likely side and hold to settlement.

Step 1 - fair value. For each live market the strategy takes the move since the window opened, $r = (S_{\text{now}} - S_{\text{open}})/S_{\text{open}}$, and the time left τ . It treats the remaining move as Gaussian with annual volatility σ (fixed per coin: BTC 0.22, ETH 0.32, SOL 0.52, XRP 0.54, DOGE 1.05), giving **fair $P(\text{Up}) = \Phi(r / (\sigma * \sqrt{\tau}))$** , with no drift. $P(\text{Down}) = 1 - P(\text{Up})$.

Step 2 - gates (config_momentum_realistic.ini). A signal fires only if *all* of these hold:

- the move since open is at least 0.15%;
- the window is 25%-65% elapsed (not too early to know, not so late that the book has repriced);
- at least this much time is left: 75 s (5m), 150 s (15m), 400 s (1h), 2 h (4h);
- the gap between fair value and the Polymarket mid is at least 0.05 (lag detection);
- the last 30 s of Binance moves the same way (momentum confirmation);
- over the last 30 s fair value has moved at least 0.04 more than the Polymarket mid (lead-lag detection: Polymarket is behind);
- the market has at least \$2,000 of 24h volume;
- **edge = fair x (1 - fee) - (ask + slippage) of at least 0.06.**

Step 3 - execution model (paper). Buy one share at the best ask. The paper fill charges the ask plus 1.5c slippage, 1.5c adverse selection, 1.0c placement latency and another 1.0c if the signal is more than 30 s old, so roughly **4-5c per share**. Fills are scaled down on thin markets. A win pays 1 - 2% fee. No early exit: every position is held to settlement.

Step 4 - settlement. In the paper record, won/lost is decided from the Binance price at window end. In this report every trade has been **re-graded on the real Polymarket outcome**: 1,444 of 1,444 resolved, 827 from our resolution cache and 647 fetched from the Gamma events API.

Operations. A cron job at 00:05 starts strategy_momentum.py unless state/momentum.disabled exists. It polls the recorder's CSV files every 5 s, writes logs/paper_trades_momentum_YYYYMMDD.csv and sends a report to

Telegram. There is **no order-placing code**: everything in this report is paper.

3 Where the edge comes from

Source of the edge: slow repricing (an information lag). Up/Down tokens are quoted by a handful of market makers and many retail clicks. When BTC jumps 0.3% in 20 s, Binance reprices in milliseconds, but resting Polymarket asks on the now-likely side stay put until a maker pulls them or a faster taker lifts them. Momentum is a taker that tries to be that faster lifter. The lead-lag gate is the direct measure: fair value moved but the Polymarket mid did not.

Evidence the edge was real:

- Win rate graded on real outcomes is 67.4% against a 50.9% break-even. Over 1,444 trades, a gap that size is not luck ($z > 12$ for a binomial test against the break-even rate).
- Bigger modelled edge earns more, in order (section 6): edge 0.06-0.10 makes +0.037 \$/trade, edge 0.25 or more makes +0.285 \$/trade. A model that was only fitting noise would not rank trades this well.
- Grading on Binance instead of the real outcome changes only 8.5% of labels and 5% of PnL. The edge is not an artefact of settlement.

Why it has faded:

- **Competition.** Average entry price rose from 0.48 (May to mid-June) to 0.61 (Sep-Oct): by the time our 5-second loop sees the gap, others have taken the cheap asks. Signals per day fell from about 20 to 2.4 with the same gates.
- **Taker fees.** All 647 recent markets fetched carry `feesEnabled=true`, `feeType=crypto_fees_v2` (section 8).
- **TWAP settlement** (from Aug 2026). Once half the window has passed, half of the average is already fixed. The real probability is then *more* certain than our end-price model says, so the market's quotes are closer to the truth than our model thinks. The model's view of where the edge is no longer matches how the market settles.
- **The model is overconfident at the extremes** (section 6): trades with fair value of 0.9 or more win 86%, not 97%.




Fig 3.1 - Calibration: realised win rate vs model fair value. Below the diagonal = overconfident.

4 Backtest / paper results

Units: one share per trade, so PnL is in \$ per \$1 of payout. ROI = PnL / capital staked (the ask). All figures are net of the paper cost model unless marked gross. "Real" = re-graded on the Polymarket outcome.




Fig 4.1 - Equity curve. Grey dashed = gross; navy = paper net (Binance-graded); red = net graded on real outcome. Orange dotted lines = period boundaries (22 Jun, 1 Sep); red dotted line = TWAP switch (6 Aug).




Fig 4.2 - Daily net PnL. 11 May (+47) and 12 May carry most of the record.

Period	Trades	Trades/cal day	WR real	Break-even	Net PnL real	\$/trade real	ROI real	Label flips
P1 May-Jun22	1053	20.6	64.6%	47.6%	+132.96	+0.126	26.5%	10.4%
P2 Jun23-Aug	278	4.0	76.6%	59.0%	+26.85	+0.097	16.4%	4.0%
P3 Sep-Oct	113	2.4	70.8%	61.2%	+4.18	+0.037	6.1%	1.8%




Fig 4.3 - Edge decay: the gap between win rate and break-even is shrinking, and so is \$/trade.

Before and after the TWAP switch.

Settlement regime	Trades	WR paper	WR real	Break-even	PnL paper	PnL real	Flips
pre-TWAP (to Aug 5)	1307	68.1%	67.3%	50.0%	+171.63	+160.18	9.1%
TWAP era (Aug 7+)	137	65.7%	68.6%	59.5%	+0.52	+3.81	2.9%

Since 7 Aug the strategy has averaged +0.028 \$/trade on real outcomes, about 1.9 trades a day. That is within noise of zero once the fill-realism penalty (section 9) is applied.




Fig 4.4 - Paper vs real settlement by month, and how often Binance grading disagrees with the real outcome, by coin. DOGE (46%) and XRP (28%) flips are high: Binance and Chainlink prices for small coins diverge most at the boundary.




Fig 4.5 - Trades per day (log scale). Activity has fallen about eight-fold since May.

5 Performance by horizon, coin, month, week and hour

Performance by market horizon




Fig 5.1 - By market horizon (5m / 15m / 4h). No trades fired on 1h, daily, weekly or monthly markets.

Horizon	Trades	WR paper	WR real	Break-even	\$/trade paper	PnL real	Flips	Verdict
5m	56	73.2%	55.4%	49.9%	+0.143	-1.68	32.1%	paper-only edge; real PnL < 0
15m	1171	69.7%	69.9%	50.6%	+0.143	+167.47	8.9%	the whole edge lives here
4h	217	56.7%	57.1%	52.7%	-0.013	-1.80	0.5%	dead: below break-even after costs

5m markets look best on paper (73% WR) but a third of the labels flip on real settlement: the move needed is tiny, and Binance vs Chainlink (and now TWAP vs end price) decides the outcome. **The 15m market is where the edge actually is: 1,171 trades, +167 real PnL, flips only 9%.**

Where the edge is: net PnL per trade, symbol x horizon




Fig 5.2 - Where the edge is: net \$/trade by coin x horizon (paper). BTC/ETH/SOL/XRP 15m are the profitable cells; 4h is red or flat.

Performance by underlying




Fig 5.3 - By coin. BTC: highest win rate, also the most expensive entries; alts win less often but enter cheaper.

Coin	Trades	WR real	Break-even	PnL paper	PnL real	Flips
BTC	683	75.0%	56.7%	+90.27	+83.72	2.5%
DOGE	67	56.7%	43.1%	+3.10	+5.51	46.3%
ETH	302	62.9%	50.5%	+33.64	+23.47	5.0%
SOL	202	56.4%	41.6%	+23.74	+22.69	3.5%
XRP	190	62.6%	43.1%	+21.40	+28.61	27.9%

Performance by month




Fig 5.4 - Monthly. May: many cheap trades. Jun-Jul: fewer, pricier, still positive. Aug and Oct: negative on paper.

Weekly net PnL (bars) and trade count (line)




Fig 5.5 - Weekly net PnL and trade count.

Performance by weekday (UTC)




Fig 5.6 - By weekday (UTC). Mon/Tue carry most trades, the weekend is thin.




Fig 5.7 - By hour of day (UTC).

6 What makes a good trade: entry conditions

Performance by time-to-expiry at entry




Fig 6.1 - By time-to-expiry at entry.

2-5 min left looks best on paper (+0.183/trade) but has a 21% flip rate on real settlement, and real PnL halves (23.8 to 11.3). 5-15 min left is the reliable core (1,052 trades, +144 real). More than 1 h left (4h markets) is negative.

Performance by entry price




Fig 6.2 - By entry price. 0.30-0.60 is the sweet spot; above 0.75 there is too little upside to cover costs.

Performance by modelled edge at entry




Fig 6.3 - By modelled edge. Monotonic: the model's ranking works, so the minimum edge (0.06) is set too low.

What-if: tightening the filters (in-sample, so treat it as a hypothesis to test forward, not a result).

Rule set (real outcomes)	Trades	WR	Break-ev en	PnL	\$/trade	ROI	Sep-Oct n	Sep-Oct \$/trade
A all trades	1444	67.4%	50.9%	+163.99	+0.114	22.3%	113	+0.037
B 15m only	1171	69.9%	50.6%	+167.47	+0.143	28.3%	43	+0.096
C B + secs_left>=300	1052	70.4%	51.4%	+144.27	+0.137	26.7%	43	+0.096
D C + edge>=0.15	517	77.4%	50.1%	+107.05	+0.207	41.3%	22	+0.144

Rule set (real outcomes)	Trades	WR	Break-even	PnL	\$/trade	ROI	Sep-Oct n	Sep-Oct \$/trade
E D + ask 0.30-0.75	458	81.0%	53.3%	+95.39	+0.208	39.1%	20	+0.142

Rule D (15m only, at least 5 min left, edge of 0.15 or more) doubles \$/trade to +0.207 and lifts ROI to 41%, but keeps only 36% of trades. In Sep-Oct it would have taken 22 trades in 47 days. Better trades, but very few of them.

7 Risk, holding period, capital use

Metric	Value	Note
Sharpe, daily, calendar days, x sqrt(365)	4.13 (real) / 3.89 (paper)	idle days count as 0
Sharpe, active days only	5.34	overstates: ignores 76 idle days
Sharpe ex 11-12 May (real)	8.49	PnL drops to 82, but Sharpe RISES: the outliers add variance. A high Sharpe here reflects many small, steady days on tiny stakes, not scalable risk-adjusted return
Max drawdown (real, 1-share)	5.97 \$	peak 2026-09-30, trough 2026-10-06
Max drawdown, \$100-per-trade sizing	~\$1,173	vs ~\$32k total real PnL at the same sizing
Longest losing streak	7 trades	
Profitable days	75.0% of active days	
Holding period	median 11.0 min, mean 33 min, p90 2.4 h	4h markets make the tail
Max concurrent positions	18	sets the capital needed
Average stake per share	\$0.509	fully collateralised; no leverage or margin calls




Fig 7.1 - Holding period (entry to settlement), positions under 4 h.

Risk profile. Each position's worst case is losing the stake (the ask, about \$0.51 per share). There is no tail risk beyond the stake, no funding and no liquidation. The real risks are model risk (settlement-rule changes like the TWAP switch), execution risk (adverse fills, section 9), and platform risk: Polygon USDC custody, smart-contract risk, and regulatory access for an India-based team, which needs legal review.

8 Fees and cost analysis

Current schedule (Fee V2, since 30 Mar 2026; official docs): only takers pay; makers pay nothing. **Fee = shares x feeRate x p x (1 - p)**, so it is largest at $p = 0.5$ and shrinks toward 0 and 1. feeRate by category: **crypto 0.07** (15m since 6 Jan 2026, all crypto horizons since 6 Mar), sports 0.05 (raised from 0.03 around Jul 2026), economics/culture/weather 0.05, politics/finance/tech 0.04, geopolitics 0. 20% of crypto fees are paid back to makers. Trading is gasless; USDC deposits and withdrawals are free. There are **no volume tiers** on the international venue; Polymarket US adds maker volume rebates above \$250k/month of taker volume. The Gamma API reports feesEnabled=true, feeType=crypto_fees_v2 on all 647 markets we traded after 21 June.

Price p	Fee / share (crypto, 0.07)	Fee as % of premium	Fee as % of max payout
0.10	0.63c	6.30%	0.63%
0.20	1.12c	5.60%	1.12%
0.30	1.47c	4.90%	1.47%
0.40	1.68c	4.20%	1.68%
0.50	1.75c	3.50%	1.75%
0.60	1.68c	2.80%	1.68%
0.70	1.47c	2.10%	1.47%
0.80	1.12c	1.40%	1.12%
0.90	0.63c	0.70%	0.63%

Fees in our paper model vs the real schedule. The paper fill charges 2% of the payout on wins (about 1c/share at $p = 0.5$) plus about 4c/share of slippage, latency and adverse-selection penalties. Total modelled cost = gross 245.6 - net 172.2 = **73.4 \$**, a 29.9% cost drag, 1000 bps of capital staked. Applying the real schedule to all 1,444 trades costs **20.4 \$** (1.41c per trade on average) against **19.6 \$** under the paper model's 2%-of-payout-on-wins rule. **The paper fee assumption is accurate.** Exchange fees are only 8% of gross PnL; the remaining ~54 \$ of modelled cost is slippage, latency and adverse-selection penalties.




Fig 8.1 - Left: total gross PnL after a taker fee of $C \times p \times (1 - p)$ per share. Right: after a flat slippage per share. Slippage hurts far more than fees: each 1c/share costs about 14 \$ (6% of gross).

Conclusion: exchange fees are not what kills this strategy. Execution quality is: the gap between the ask you see and the ask you get.

9 Latency requirements

The edge is the delay before Polymarket reprices, so reaction time decides whether we earn it or pay it. A money-free fill simulator (fill_sim.py, June) replayed signals against the **recorded full order book** at delay D after the signal, filled only if the ask was still there, and graded on the real outcome:

Sample	D = 0 s	D = 1 s	D = 2 s	D = 5 s (current)	Reading
May batch (38 signals) - WR if filled	42%	-	-	19%	adverse selection +58 pp
May batch - fill rate winners vs losers				33% vs 91%	we fill the losers

Sample	D = 0 s	D = 1 s	D = 2 s	D = 5 s (current)	Reading
Jun batch (58 signals) - WR if filled	42%	-	-	37.8%	adverse selection +9 pp
20260512 focus (35) - WR if filled	-	-	45.8%	47.6%	below 47.8% break-even
Adverse-selection crossover	negative	negative	turns positive	positive	kicks in after 1-2 s

What this means:

- When we are slow, the ask is still there only when the move has reversed (the losers). On real winners the ask is gone. Paper fills at the observed ask hide this.
- Current loop: poll recorder CSV files every 5 s, so about 5-7 s from Binance tick to decision. **FAIL.**
- Needed: Binance WebSocket and Polymarket CLOB WebSocket feeding one event loop, with a pre-signed order sent in **under 500 ms** from Binance tick to order (ideally about 100 ms). Polymarket matching is off-chain, so blockchain block time does not matter for fills, only for settlement. A VPS close to Polymarket's API servers would help, but this is a 100 ms game, not a microsecond HFT game.
- The TWAP switch also changes the game: the outcome builds up over the window, so a TWAP-aware model has a few more seconds of predictable information, a possibly less crowded edge, but it still needs a sub-second loop.

10 Capacity: how much capital can be deployed

All positions are fully collateralised in USDC; there is no margin. Two things limit capacity: **(a) depth**, meaning how many dollars sit at the ask before the price moves, and **(b) signals per day**.




Fig 10.1 - Median ask-side depth in \$, sampled from today's recorder order books (prices 0.15-0.85).

Market	Best ask \$ (median)	5 levels \$ (median)	5 levels \$ (p25)	Median spread	Lifetime vol of traded mkts	Practical size / signal
BTC 15m	71	1,374	732	1c	\$34.8k	\$150-300
BTC 5m	78	496	284	1c	-	\$75-150
ETH 15m	17	282	141	1c	\$8.4k	\$50-100
SOL 15m	9	141	71	2c	\$4.1k	\$25-50
BTC 4h	7	294	152	3c	\$20.0k	drop (no edge)
XRP 15m / DOGE	<10	<100	<60	2-7c	\$0.4-1.8k	\$10-25

Lifetime volume = Polymarket volumeNum for the markets we traded after 21 June (median). Depth was sampled on one day only (2026-10-18, 05:46 UTC onwards); treat it as an order of magnitude.

Capacity scenarios (BTC-weighted average size about \$150 per signal; capital = max concurrent positions x size):

Scenario	Signals/day	ROI/trade	Notional/day	PnL/month (gross of adverse sel.)	Capital needed
May-June regime (history)	20	26%	\$3,000	~\$23,000	~\$3k (18 x \$150) - \$5k
Sep-Oct regime (now)	2.4	6%	\$360	~\$650	~\$3k
Rule D, now	0.5	24%	\$70	~\$500	~\$1k
Rebuilt, sub-second, TWAP model (target, unproven)	5-10	10-15%	\$750-1,500	\$2k-7k	\$5k

Opportunity vs capacity. The crypto Up/Down segment trades about \$1M a day in 15m markets alone (our 10-18 snapshot), but one signal can only take a few hundred dollars before it moves the price. **Even in the best historical regime this is a few-thousand-dollar bankroll earning low five figures a month at most.** Now it earns about the cost of a server. It works as a learning vehicle and a source of data; it is not a business on its own.

11 Pros and cons

Pros	Cons
Clear, explainable source of edge (Polymarket reprices after Binance)	Edge decayed ~6x in five months; Sep-Oct close to zero
Large, statistically solid sample (1,444 trades, real outcomes)	53% of PnL from two days (11-12 May): fragile, regime-dependent
Losses capped at the stake; no leverage or liquidation	At the current 5 s loop we fill mostly losers (fill-sim FAIL)
Model's edge score ranks trades well (monotonic)	Model ignores TWAP settlement since Aug 2026
Short holds (median 11 min): fast feedback and capital turnover	Capacity of a few hundred \$ per signal; tiny next to a team's cost
Full data stack already built: recorder, fill sim, re-grader	Taker fees on all crypto markets now; costs are 30% of gross
15m BTC/ETH/SOL/XRP are deep, 24/7, continuous markets	4h and 5m markets carry no real edge
Settlement verified at 100% coverage via the public API	Platform/regulatory risk for an India-based team (custody, access)

12 Next steps

#	Step	Cost	Pass / kill test
1	Pause the current paper Momentum (state/momentum.disabled) or keep it only for data; add a daily automatic re-grade against the Gamma API	0.5 day	-
2	TWAP-aware fair value: $P(\text{TWAP_window} \geq S_{\text{open}})$ given the elapsed path; Chainlink TWAP stream as input	3-5 days	beats current model on calibration (Brier) for Aug-Oct markets
3	Event-driven WebSocket rewrite: Binance + CLOB books, pre-signed orders, measure tick-to-order time	1-2 weeks	p95 under 500 ms
4	Re-run fill_sim at D of 0.25/0.5/1 s on Aug-Oct recorded books with rules D + TWAP model	2-3 days	WR if filled > break-even + 5 pp; adverse selection under 5 pp
5	Maker variant: rest bids below fair value instead of lifting asks (earns any maker rebate, avoids paying the ask)	1 week	queue-aware sim positive
6	Micro-live (only with explicit approval): \$2/order, \$50/day cap, funded Polygon wallet + CLOB API keys	2 weeks	real WR if filled within 3 pp of sim
7	Go / no-go: scale to \$150-300 per signal only if step 6 passes for 4 weeks	-	monthly PnL > infra + time cost

Method notes. Data: logs/paper_trades_momentum_*.csv (153 files, 2026-05-03 to 2026-10-18). Real outcomes: CLOB tokens[].winner (827 cached) + Gamma /events?slog=outcomePrices (647 fetched 2026-10-18). A flipped trade's real PnL is recomputed as $(0.98 - \text{ask} - 0.015)$ or $-(\text{ask} + 0.015)$. Sharpe uses daily PnL $\times \sqrt{365}$ (crypto trades 24/7). Fill-realism figures are from the June 2026 fill-sim study. Paper trading only: no real orders were placed.