

Lagisalpha

Litepaper - v1.1

The lead-lag edge in prediction markets. A prediction market sets its price by trading, so it lags the sharp, vig-free line that already holds the true probability. When it falls below fair, the cheap side is underpriced; across sixteen settled World Cup matches it travelled back to fair about 82% of the time, on every call the detector fired, none filtered out. Built on the TxLINE World Cup data layer by Onenept Studios.

1. The claim

A prediction market moves its price only when someone trades. News moves faster. So in the seconds after a goal the market sits behind the true probability, and one side is cheap. That lag is the edge: measurable, repeatable, and gone the moment the market catches up.

2. Why the edge exists

Two facts. The market is slow: it reprices by trading, not by knowing. The reference is fast: TxLINE strips the bookmaker margin from a live odds feed, so its price is the true probability, and it moves the instant news lands. The gap between them is how far the market is behind, and which way it is about to move.

3. The signal: a divergence

We work in probability space. TxLINE's de-vig 1X2 gives the fair probability a team wins. The market's moneyline gives its own probability of the same event. When the fair sits above the market price by more than a threshold, the cheap side is underpriced, and we mark an entry: which side, how far off fair, and how much size you could later exit into at fair. One dislocation is one entry, not a burst.

4. The proof: does it close, does it pay

Two tests, on sixteen settled matches, on the real fills. Reach: from the entry, does the market price travel to the fair before the match ends. It does about 82% of the time, and the move often takes minutes, so a short holding window hides it. Reach does not depend on who eventually wins, so it is the firmer number.

Return: the trade is to buy the cheap side and take profit at fair when the market catches up. Sized by Kelly on the gap, $f = \text{gap} / (1 - \text{price})$, capped at 30% of the balance per call, and compounded across every call with nothing excluded, that stands at about plus 1366% at a 5 point gap and plus 1777% at 10. The same bets held to the final result instead returned about minus 64% and plus 14%: whichever exit you pick, the convergence leg is where the money is, and holding to the outcome does far worse. The cap is what earns that: full Kelly, uncapped, once staked 81% of the balance on a single call and gave back 76% of it, so capping any one bet at 30% bounds the damage while keeping every call in the record. The compounded number is still concentrated, a couple of high-volume matches carry it, so reach is the firmer read; the return is published as-is and moves as each match settles.

5. The data, verifiable both sides

The fair is TxLINE's World Cup feed: odds and scores anchored on Solana, access minted by a real on-chain subscribe transaction. The market side is real fills read straight from Polygon, decoded to a price and a size per trade. Both legs are public: open any fill as a Polygon transaction, settle any outcome on TxLINE's on-chain scores, and recompute the edge yourself. Nothing here is asserted.

TxLINE endpoints used (server-held token: guest JWT plus an on-chain Solana subscribe transaction gives an apiToken, sent as Authorization Bearer and X-API-Token):

GET /api/odds/snapshot/{fixtureId}: live de-margined 1X2 fair, the core reference.

GET /api/scores/snapshot/{fixtureId}: final goals for outcome settlement.

GET /api/scores/stat-validation: validateStat Merkle proof vs the on-chain daily-scores root.

GET /api/fixtures/snapshot: live fixtures, team names, kickoff times.

Polygon OrderFilled logs: the prediction market fills, decoded on-chain (the market side).

6. How to trade it

Catch the divergence live, take the cheap side at the market price, and take profit at TxLINE fair when the market catches up. Size each bet by Kelly on the gap, so a bigger dislocation gets a bigger bet and you never over-bet into ruin. Holding to the final result instead is a losing trade on this data, so the play is the take-profit, not the settlement.

The size we show is the liquidity you could have exited into at fair or better, counted only when the price actually reached fair; when it never does, the size is zero, because you could never have exited there. How much you take, and any price you move by taking it, is your own execution cost. It is not part of the signal.

7. What we do not claim

The edge is validated on sixteen matches, so the return is a pilot, not a promise. The confidence interval still spans zero at this sample, and the return leans on a few high-volume matches; the reach rate is the firmer read, and both tighten as matches accrue. This measures a delay between two markets. It is not a trading strategy, it is not financial advice, and any sizing or slippage is your own.

8. What we found

The obvious idea is a sharp-movement detector: flag significant TxLINE odds shifts and track whether they call the result. We built it and killed it twice. Graded early, a significant fair shift by the 45th minute called the winner about 58% of the time, a coin flip. Graded whenever a big shift fires (10pp inside a minute, any time in the match) it looks strong at 83%, but that number is hollow: a rule that ignores the shift entirely and just backs whichever team the fair currently rates higher makes the identical call in 12 of 12 matches, with the identical score. Nine of twelve shifts sit within two minutes of a goal, so the shift is the goal being repriced; all of the predictive content is in the level of the fair, none is in the move itself. The forecast has no alpha.

The edge is not the line moving; it is the market being slow to follow it. A goal is new information: TxLINE reprices it instantly, but a prediction market only moves when someone trades, so for a window the cheap side sits below fair. That lead-lag converges about 82% of the time, and it is our strongest, most proven signal. The line move carries no forecast; the lag in the market's reaction to it is the entire product.

And the record rolls on its own. Every divergence the detector fires is published and scored: either side, any size, any minute of the match, each side named by its team. There is no exclusion filter and no curated subset; sizing is the only risk control, and it is Kelly on the gap capped at 30% of the balance per call. An earlier version instead cut two classes of buy-NO call; we retired that filter and cap the sizing instead, so the mechanism is shown whole, with the calls that hurt it left in and bounded rather than removed.

Separately, a TxLINE high-danger possession makes a goal by that team about four times more likely within two minutes, and a divergence it flags converges to fair about 84% of the time versus 75% without. All of this is on sixteen settled matches, in-sample; it is a promising pilot, not a settled result.

9. Roadmap: 75 leagues, forecasting done honestly

Scale the measurement. TxLINE carries the demargined consensus for around 75 leagues, and the lag is structural: every prediction market reprices by trading, so it sits behind a fast, vig-free reference in any league that has one. The same divergence detection runs per league and per market, and the first deliverable is a lag profile for each: how wide the average dislocation opens, how long it takes to converge, and how much size sits in the window. That map, which leagues lag longest, is itself the product, and every league gets the same public calibration ledger this pilot has.

Model the outcome as unknown. We killed the sharp-movement detector, but the flaw was the grading, not the ambition. Scoring a shift by "did that team win" assumes the winner is knowable from history and lets late, goal-driven shifts grade themselves. The honest frame is Bayesian: the pre-match de-vig consensus is the prior, each early odds shift updates a posterior over the result, and the model is scored on calibration (Brier and log score) across thousands of matches in 75 leagues, never on hit-rate over twelve. This is the approach of Robberechts, Van Haaren and Davis, "A Bayesian Approach to In-Game Win Probability in Soccer" (KDD 2021): treat the in-play outcome as uncertain throughout, seed the prior from pre-match strength, update sequentially, and judge the posterior, with the dynamic-prior lineage going back to Rue and

Salvesen (2000).

Where the two meet. The test that matters is whether an early-shift posterior moves before the traded market does. If it does, the lead-lag gains a forward leg on top of the post-goal reactive one; if it does not, we will publish that, the same way we published the coin flip. Either answer sharpens the signal API: per-league divergence feeds, portfolio Kelly across simultaneous matches, and new markets (cards, corners) as TxLINE streams them de-vigged.

Lagisalpha, Onenept Studios, built on the TxLINE / TxODDS World Cup data layer. This document is informational, not an offer of securities or a solicitation to gamble.