

Who wins the book? Margin, customers and informed flow in Polymarket's esports markets

A sportsbook-style study of every on-chain esports fill settled in September 2026

Afonso Martins
afonsomartins.com · 4 October 2026

\$315M taker stakes on 15,255 markets	5.48M on-chain fills from 30,080 wallets	-0.50% book hold, 95% interval [-1.5%, +0.5%]	79% of stakes placed in-play, after the scheduled start
---	--	---	---

ABSTRACT

This study reads one month of Polymarket's esports markets the way a sportsbook trading team reads its own book. Polymarket has no house, so the "book" is every wallet that left a resting order. Using all 5.48 million on-chain fills settled in September 2026 (\$315M of taker stakes across 15,255 markets and 30,080 wallets), I measure margin, customer mix, liquidity supply, exposure and the information content of order flow. The book lost \$1.6M, a hold of -0.50%, but a bootstrap that resamples whole matches puts the 95% interval at [-1.52%, +0.45%], so the month is consistent with a book that roughly breaks even. Pre-match the book held +1.20%; 79% of stakes were placed in-play, where it gave back 0.94%. Money is highly concentrated: 277 wallets, the top 1% of takers, placed 63% of all stakes, and the largest stakers finished ahead while the smallest lost 11.9% of what they staked. The closing price is well calibrated (slope 1.01, s.e. 0.03), and even heavily one-sided pre-match flow won no more often than the closing price implied. What that flow did capture was the move before the close, a pre-match closing line value of +2.3% that the book paid for. Liquidity rewards on these markets came to \$0.22 for the month, so maker P&L is close to the makers' full economics. On 1,383 matches also listed on Kalshi, the two venues' prices sat about 1c apart, and Polymarket moved first on 71% of markets.

Keywords: prediction markets, Polymarket, esports betting, market making, hold, closing line value, favourite-longshot bias, adverse selection

CONTENTS

1	Introduction	6	Conclusion
2	Data and method		References
3	Results	A	Data quality checks
4	Discussion	B	Measure dictionary
5	Limitations	C	Reproducibility

1 Introduction

A sportsbook knows three things about its business at the end of every month: how much it took in stakes, how much of that it kept, and who it kept it from. Polymarket's esports markets now carry the volume of a mid-sized book, \$315M of stakes in September 2026 alone, but they are built differently. There is no house. Every price on the screen is a resting order from some wallet, and every fill pairs one of those wallets (the *maker*) with a wallet that crossed the spread (the *taker*). The aggregate of all makers plays the role of the bookmaker, and because the venue charged no taker fee on esports in September, what takers lost is almost exactly what makers won.

Polymarket settles on the Polygon blockchain, which makes this a rare dataset. Every fill names both wallets, the price and the size, and every market's payout is written on-chain. A sportsbook's customer table is private. Here it can be rebuilt from public data, wallet by wallet, and joined to an independent recording of the venue's order book. This study uses that to answer six questions a trading desk would ask about its own month:

1. **Size.** How large was the esports book, and which titles and market types carried it?
2. **Margin.** Did the book keep a margin, and where did it win or lose, by timing and by price?
3. **Customers.** Who were the takers, how concentrated were they, and which of them finished ahead?
4. **Supply.** How concentrated is the liquidity, and do the largest makers do better than the rest?
5. **Information.** When pre-match flow piled onto one side, did it know something the price did not?
6. **Two venues.** When the same match also trades on Kalshi, how far apart are the prices, and which venue moves first?

Each figure and table below is computed from the same warehouse that feeds the Power BI report *Polymarket Esports Desk*, so the paper and the dashboard agree to the cent. What the paper adds is inference: interval estimates that respect the fact that markets on one match settle together, a calibration test of the closing line, and a direct test of whether lopsided flow beat the price it traded at.

2 Data and method

2.1 Data

The sample is every Polymarket esports fill with a Polygon block time between 1 and 30 September 2026 (UTC), 5,482,207 fills on 15,255 markets. Four inputs feed it, all read from the research desk's content-addressed R2 lake, where every object is hash-checked on read:

- **On-chain fills.** `OrderFilled` events decoded by the desk's audited decoder from four independent builds (a published fill table, two HyperSync pulls for the windows it did not cover, and a VPS archive kept as a cross-check). Where builds overlap they agree fill for fill.
- **Settlement.** Conditional Tokens payout vectors from the chain, with the venue's own `market_resolved` message and the last traded price as second and third sources.

- **Order book.** The desk's own websocket capture of the venue's book and prints, 80% of on-chain turnover cross-checked against it and 704 hours of the month recorded.
- **Membership.** The venue's event crawl, which defines which markets are esports, their title, market type and scheduled start.
- **Kalshi (Section 3.8 only).** The desk's own capture of Kalshi's websocket for the same esports series, reduced to the venue's best bid and ask per second, with team names and start times from Kalshi's public events API.

The data are modelled with dbt on DuckDB into a star schema (39 models, 36 tests) around a flow fact at the grain of market, day, phase, price paid, ticket size, customer segment, wallet rating and maker tier. Wallet addresses stop at the intermediate layer; the marts carry pseudonyms (MM-01 for the largest maker, and so on).

2.2 Definitions

The study borrows sportsbook vocabulary. Table 1 maps each measure to the term a trader would use; the full dictionary with caveats is Appendix B.

Table 1. Measures used in the study and their sportsbook equivalents.

Measure	Sportsbook term	Definition
Turnover	Handle	What takers staked: contracts $\times$ the price of the outcome they bought
Maker P&L	GGR	Resting side's P&L on every fill, marked to the on-chain payout
Hold	Margin	Maker P&L $\div$ turnover on settled markets
Taker return	Customer yield	$-\text{Maker P\&L} \div \text{settled turnover}$, for a group of takers
Customer segment	Segment	Taker stake in the month: Micro $< \$100$, Small, Mid, Large, Whale $\$100k+$
Closing line value	CLV	Contracts $\times$ (closing price $-$ price paid) $\div$ stake, pre-match fills
Wallet rating	Sharp / square	CLV over pre-match stakes: Sharp $\geq +2\%$, Square $\leq -2\%$ (rated from \$500)
Liability at start	Liability	Larger of the book's two possible losses at the scheduled start
Flow imbalance	Weight of money	Net taker contracts on one side $\div$ all pre-match taker contracts

2.3 Inference

Results swing with outcomes, so every hold and return in the paper comes with a 95% interval. Markets are not independent draws: a match winner market and its map markets settle on the same games, and they are often traded by the same wallets. Intervals are therefore percentile intervals from a cluster bootstrap that resamples whole matches (2,834 clusters, 4,000 resamples) and recomputes each ratio (Cameron, Gelbach and Miller, 2008). Calibration of the closing line is tested with a logistic recalibration slope (Cox, 1958) and the Brier score (Brier, 1950). Whether one-sided flow beat the price is tested by comparing the share of markets the crowd's side won with the probability the closing price gave that side, with a normal approximation to the sum of Bernoulli variances.

2.4 Data quality

Every build of the warehouse runs 21 checks (Appendix A). All 14 of the 14 checks with a pass mark passed in this build. The two that matter most for the results: maker P&L and taker P&L net to zero to the cent, and the book's hold measured a second time from the venue's own prints, a tape recorded independently of the chain, comes out at -0.65% against -0.50% on chain. The on-

chain payout matched the venue's resolution message on every market where both exist, and 99.94% of turnover sits on markets with a settled label.

3 Results

3.1 Size and composition

Takers staked \$315M across 5,482,207 fills in September, about \$10.5M a day, with the busiest day on 12 Sep (\$18.3M). Three titles carried 97% of it (Table 2). League of Legends was the largest by stakes even though Counter-Strike listed more than twice as many markets: the LCK and LPL playoffs ran through the month and concentrated money on fewer, bigger matches. Winner markets (match and map) took 96% of stakes; handicaps, totals and props are numerous but small.

The distribution across markets is steep. The 712 markets that drew more than \$100k each carried 74% of all stakes, while the 9,367 markets under \$1k carried less than half a percent.




Figure 1. Daily taker stakes by title. Weekend peaks follow the playoff schedules of the LCK, LPL and the large Counter-Strike events.

Table 2. Turnover, markets and hold by title, September 2026.

Title	Turnover	Share	Markets	Hold
LoL	\$138M	43.7%	3,504	+0.01%
CS2	\$105M	33.2%	7,768	-1.15%
Dota 2	\$63.0M	20.0%	1,572	-1.01%
VAL	\$9.2M	2.9%	793	+2.70%
Seven others	\$504k	0.2%	1,618	-1.68%
All titles	\$315M	100%	15,255	-0.50%

Hold is maker P&L over settled turnover. Positive means the book won.

3.2 Margin: where the book won and lost

Over the month the book lost \$1.6M on \$315M of settled stakes, a hold of -0.50%. The cluster bootstrap interval is [-1.52%, +0.45%]. A month of esports results is not enough to tell this apart

from zero, and the weekly figures show why: the book lost in 4 of 5 calendar weeks, with weekly holds from -1.79% to $+0.55\%$ (Table 3).

The split by timing is more informative (Figure 2). Fills placed before the scheduled start held $+1.20\%$ for the book (interval $[-2.3\%, +4.8\%]$), the pattern a sportsbook would recognise: a margin earned on pre-match prices. But 79% of all stakes were placed after the start, and in-play the book held -0.94% with an interval of $[-2.00\%, +0.02\%]$, only just touching zero. In-play esports prices move round by round and map by map. Makers who quote through those moves are exposed to takers who see a round result a few seconds before the book does, and the in-play figure is where that cost would appear.




Figure 2. Taker stakes (a) and the book's hold (b) by when the fill happened relative to the scheduled start. Whiskers are 95% intervals from a bootstrap over matches.

Table 3. Weekly results. Weeks run Monday to Sunday; the first and last are partial.

Week	Days	Turnover	Fills	Maker P&L	Hold
1 Sep to 6 Sep	6	\$74.2M	1,033,715	-\$328k	-0.44%
7 Sep to 13 Sep	7	\$79.3M	1,320,805	-\$914k	-1.15%
14 Sep to 20 Sep	7	\$62.7M	1,049,623	-\$292k	-0.47%
21 Sep to 27 Sep	7	\$73.6M	1,509,288	\$403k	+0.55%
28 Sep to 30 Sep	3	\$25.2M	568,776	-\$447k	-1.79%

3.3 Price paid

Betting markets have a long record of the favourite-longshot bias: outcomes at long odds are overbet and return less than favourites (Griffith, 1949; Thaler and Ziemba, 1988; Snowberg and Wolfers, 2010). September's flow points the same way (Figure 3). Takers who paid under 20c got back -8.4% of what they staked, and those who paid 60c or more got back $+2.7\%$. The under-20c band is small (\$5.4M of stakes) and its interval, $[-23.8\%, +9.0\%]$, is wide, so on its own it proves little. The middle bands are better measured: takers paying 20-39c lost 11.9% (interval $[-23.0\%, -0.6\%]$), and takers paying 60-79c made $+5.2\%$ ($[+0.7\%, +9.9\%]$). Both intervals exclude zero. The underdog side of a two-way esports market is where takers gave money to the book, and the favourite side is where they took it back.




Figure 3. Taker return on settled stakes by the price paid for the outcome bought. Whiskers are 95% cluster-bootstrap intervals.

3.4 Customers

27,656 wallets took liquidity in September. Grouped by what each staked in the month (Table 4), they look like the customer base of a book with a large VIP desk. The 485 whales, each staking \$100k or more, placed 72% of all stakes; the 13,462 micro accounts, nearly half of all takers, placed 0.1%. The top 1% of takers (277 wallets) account for 63% of stakes, and the Gini coefficient of stakes across takers is 0.95.

A retail sportsbook limits the small group of customers who win. Polymarket has no limits, and in September size and return went together (Figure 4): whales finished +0.8% on settled stakes and 57% of them were ahead for the month, while the smaller segments lost more the smaller they were, down to -11.9% for micro wallets. A wallet is an account, not a person, and many of the largest are likely automated, so "whale" here means large flow rather than a large gambler.





Figure 4. Customer segments by monthly taker stake (wallet counts in brackets): share of stakes (a) and return on settled stakes (b).

Table 4. Customer segments, September 2026.

Segment	Monthly stake	Wallets	Staked	Share	Median stake	Return	Ahead
Whale	\$100k+	485	\$227M	72.2%	\$236,760	+0.8%	57%
Large	\$10k-100k	2,185	\$67.5M	21.4%	\$23,102	+0.4%	51%
Mid	\$1k-10k	4,870	\$17.4M	5.5%	\$2,773	-1.9%	45%
Small	\$100-1k	6,654	\$2.5M	0.8%	\$293	-5.9%	43%
Micro	<\$100	13,462	\$322k	0.1%	\$13	-11.9%	52%

Return is taken P&L over settled stake for the whole segment. Ahead is the share of wallets with positive P&L for the month.

Retention. Of the wallets first seen in the week of 7 Sep, 38% traded again the next week and 28% two weeks later (Figure 5). The first week's cohort is larger and stickier because it contains every wallet that was already active before September; "first seen" means first seen in this edition.

**Figure 5.** Weekly retention: share of each week's first-seen wallets that traded again one, two and three weeks later. Whole weeks only.

Sharp and square. Wallets with at least \$500 of pre-match stakes were rated by their closing line value (Table 5). The 973 sharp wallets beat the close by 7.1% on average and made +5.3% on settled stakes, with 60% of them ahead. Square wallets lost 4.3% to the close but still finished slightly positive, so the rating separates the best wallets better than it identifies the worst. Read it as a description of the month rather than a forecast; CLV is an unbiased ruler for a single bet but a noisy predictor of next month's results.

Table 5. Wallet ratings by closing line value on pre-match stakes.

Rating	Wallets	Pre-match stake	CLV	Return	Ahead
Sharp	973	\$9.1M	+7.1%	+5.3%	60%
Neutral	2,277	\$48.1M	-0.1%	-1.6%	48%
Square	1,135	\$7.4M	-4.3%	+0.7%	43%
Not rated	25,695	\$986k	-0.3%	-0.5%	48%

Not rated: under \$500 of pre-match stakes in the month. Return covers all of a wallet's settled stakes, pre-match and in-play.

3.5 Who is the book

13,236 wallets made at least one fill as the resting side. Liquidity is spread more widely than the customer money: the ten largest makers made 15% of all volume and it takes 110 makers to reach half (Figure 6). The Herfindahl-Hirschman index of volume made is 44, far below the 1,500 that competition authorities treat as the start of concentration. Most makers also take: 10,812 wallets did both in the month, and 2,424 only ever rested orders.

The largest makers did slightly better than everyone else. Fills where a top-10 maker was the resting side held +0.10%, against -0.60% for all other makers combined. Individual results vary much more (Table 6): among the eight largest, holds run from -3.7% to +2.5%. Several of the largest makers rest almost all of their volume (the last column), which is the profile of a dedicated market-making operation rather than a trader who sometimes quotes.

Rewards. Polymarket pays makers a daily liquidity reward on some markets, outside the on-chain fills, so maker P&L could understate what makers earned. It does not here. In September the venue's event data listed a reward programme on 188 of the 15,255 esports markets, each at \$0.001 a day, which comes to \$0.22 for the whole month. Maker P&L is effectively the makers' full economics.




Figure 6. Concentration of liquidity supply: cumulative share of volume made by makers ranked from largest to smallest.

Table 6. The eight largest makers by volume made.

Maker	Volume made	Share	Markets	P&L	Hold	Maker share
MM-01	\$7.82M	2.5%	187	-\$173k	-2.21%	46%
MM-02	\$7.64M	2.4%	3,402	\$38k	+0.49%	97%
MM-03	\$6.24M	2.0%	859	\$158k	+2.54%	99%
MM-04	\$6.10M	1.9%	1,041	\$74k	+1.21%	61%
MM-05	\$3.58M	1.1%	816	\$27k	+0.77%	100%
MM-06	\$3.42M	1.1%	2,085	\$18k	+0.53%	47%
MM-07	\$3.39M	1.1%	6,479	\$4k	+0.11%	100%
MM-08	\$3.07M	1.0%	109	-\$113k	-3.68%	85%

Maker share: the share of the wallet's own activity in which it was the resting side. Liquidity rewards on these markets were nominal; see the text.

3.6 Exposure and informed flow

At each market's scheduled start the book carries a net position, and its liability is the larger of the two losses it could take. Summed over the 14,953 markets scheduled in September that is \$23.2M; the largest single liability was \$401,808 on Spirit vs MOUZ (BO5), and 466 markets carried more than \$10k. Risk is spread thin: the ten largest positions hold 10.9% of the total, and no single maker carries any of them alone because the book is many wallets netted together. Counter-Strike carried 60% of liability at the start, well above its 33% share of stakes, because more of its

money arrives before the start: 36% of Counter-Strike stakes were pre-match, against 14% for League of Legends and 8% for Dota 2.

Table 7. Largest liabilities at the scheduled start, markets scheduled in September.

Match	Title	Start	Liability	Winner	Book P&L
Spirit vs MOUZ (BO5)	CS2	6 Sep	\$401,808	Spirit	\$87k
TYLOO vs G2 Esports (BO3)	VAL	24 Sep	\$344,852	G2 Esports	\$185k
LOUD vs EDward Gaming (BO3)	VAL	26 Sep	\$258,570	LOUD	\$113k
Invictus Gaming vs LGD Gaming (BO5)	LoL	8 Sep	\$257,605	Invictus Gaming	\$208k
KT Rolster vs Dplus KIA (BO5)	LoL	4 Sep	\$257,341	Dplus KIA	\$161k
1WIN vs B8 (BO3)	CS2	10 Sep	\$229,954	1WIN	\$209k

Book P&L is the market's final maker P&L, including in-play trading after the start.

The more interesting question is whether lopsided flow is informed. Markets were grouped by how one-sided their pre-match taker flow was, from balanced (under 20% net) to very one-sided (80% or more of contracts on one side). In the very one-sided group the side the crowd favoured won 55.5% of 2,597 markets, which looks like an edge until it is set against the price. The closing price gave that same side 55.6% ($z = -0.09$), and no group differs from its closing price by more than 0.7 standard errors (Figure 7a). One-sided flow did not know the result better than the market's final price.

It did know the direction of the next move. Takers in very one-sided markets beat the close by +2.3% on their pre-match stakes, against -0.5% in balanced markets, and the book's hold falls steadily from +0.18% to -1.86% across the groups (Figure 7b). The price ends up right, but makers who quoted before it moved sold to takers at the old price. This is the adverse selection of Glosten and Milgrom (1985): the book loses not because the crowd outguesses the final price but because part of the crowd arrives before it.




Figure 7. One-sided flow against the closing price. (a) Share of markets won by the side pre-match takers favoured, with the probability the closing price gave that side. (b) Taker closing line value on pre-match stakes and the book's hold, by the same groups.

The closing price itself is well calibrated across the 7,116 two-way markets that had one (Figure 8). The recalibration slope is 1.007 with a standard error of 0.034, consistent with 1, and the Brier score of 0.208 compares with 0.250 for a coin flip. The only visible departure is in the tails, where outcomes priced under 10c won 12% of the time across 119 markets, a sample too small to carry much weight.




Figure 8. Calibration of the closing line. Markets binned by the closing price of outcome A; dot area shows the number of markets, whiskers are 95% binomial intervals, and the dashed line is perfect calibration.

3.7 Liquidity

The desk's own book recorder shows the venue's spread hour by hour (Figure 9). On Counter-Strike winner markets the median quoted spread is 1.8c a day before the start and narrows to 1.2c in the final hours, as volume and makers arrive. Once play begins it widens to 2.2c, because every round changes the fair price and makers widen to protect themselves. League of Legends and Dota 2 winner markets quote about 1.0c and 1.0c in the hours before the start, which is the 1c tick on most prices, so the typical pre-match book on the big titles is as tight as the venue allows. The in-play widening is the liquidity side of the in-play hold in Section 3.2: makers are paid a wider spread exactly where they are most exposed to faster information.




Figure 9. Median quoted spread on match and map winner markets by hours from the scheduled start, from the research desk's own websocket capture of the venue's book.

3.8 Two venues, one match

Many of these matches also trade on Kalshi, a US exchange regulated by the CFTC, with its own order book and no on-chain settlement. The desk records Kalshi's websocket as well, so the same match can be priced on both venues at the same minute. Of 2,394 Kalshi match-winner events dated in September, 1,854 paired with exactly one Polymarket market on both team names, the title and a start within six hours; 1 matched more than one candidate and was refused, and the rest had no counterpart. On 1,383 matches both books were recorded and two-sided at the same time, which gives 1,411,312 market-minutes to compare (Table 8).

The two prices sit close together. The median gap between the venues' mid prices is 1.0c before the start and 1.0c in-play, and it is above 2c in 21% of pre-match minutes. A gap can only be traded when one venue's bid sits above the other's ask. On minutes where both books updated, that happened in 16% of pre-match minutes, usually by a single cent, and in 6% once Kalshi's taker fee is paid. One-minute snapshots overstate this, because the two quotes can be seconds apart inside the minute and the size at the touch may be small, so these shares are an upper bound.

The more useful result is which venue moves first. Minute-to-minute price changes on the two venues correlate at 0.90 in the same minute, and the leftover correlation one minute apart is not symmetric: a Polymarket move is followed by the same move on Kalshi a minute later (0.067) more strongly than the reverse (0.027). Market by market, Polymarket leads on 71% of 2,590 markets, with a 95% interval of 69% to 73% from resampling matches (Figure 10a). On the paired matches Polymarket took \$129M of stakes against \$64.4M on Kalshi, about 2.0 times as much, which is consistent with price discovery happening where most of the trading is (Hasbrouck, 1995).




Figure 10. The same match on Kalshi and Polymarket. (a) Per market, the correlation of a Polymarket price change with Kalshi's change one minute later, minus the reverse; right of zero means Polymarket moved first. (b) Gap between the two venues' mid prices, pre-match and in-play.

Table 8. Price gaps between the venues on paired matches, one-minute snapshots.

Phase	Minutes	Median gap	Gap over 2c	Crossed	Crossed after fee
Pre-match	1,131,369	1.0c	21%	16%	6%
In-play	279,943	1.0c	18%	8%	4%

Both books two-sided with spreads of 10c or less. Crossed: one venue's bid above the other's ask, on minutes where both books updated. After fee: net of Kalshi's taker fee, $0.07 \times P \times (1 - P)$ per contract.

4 Discussion

Read as a sportsbook's month, September was a break-even book with a profitable pre-match business and a costly in-play one. The pre-match margin of +1.2% is thin next to the several points a retail book builds into its esports prices, and it was earned on only 21% of the money. The in-play share is much higher than a traditional book's, and on this venue in-play is where the book gives the margin back.

The customer picture is not the one a sportsbook would expect. The largest wallets are the most profitable and the smallest lose the most, so a venue that wants healthy liquidity cannot rely on the usual sportsbook move of limiting winners. The makers' answer, visible in the spreads, is to price in-play risk into a wider quote rather than refuse the business.

The flow result is the one most worth carrying forward. Heavy one-sided money in the run-up to a match is not evidence that the price is wrong at the start; the closing price had already absorbed it, and was well calibrated. The cost to makers came earlier, from quoting before the move. For a market maker this points to the timing of quotes in the hours before a start, not to the final price, as the place where pre-match risk is managed.

Across venues, Kalshi's esports prices track Polymarket's closely and tend to follow them. For a desk quoting on Kalshi, or a feed that aggregates both, Polymarket's book is the reference price, and the minute after a Polymarket move is when a Kalshi quote is most likely to be stale.

5 Limitations

- **One month.** Hold swings with results, and several intervals in this paper include zero. The direction of the pre-match and in-play margins should be confirmed on further months before it is relied on.
- **Rewards come from the event data.** Reward programmes were read from the venue's event crawl. A programme set up only outside it, or a one-off promotion, would not be counted.
- **Two venues at one-minute resolution.** The cross-venue section compares one-minute snapshots, so a lead shorter than a minute is invisible and a crossed quote may not have been tradable at size.
- **Wallets are not people.** One trader can run several wallets and one wallet can be a bot serving many people. Segments describe accounts.
- **First seen is relative to the edition.** Wallets active before September count as new in their first September week, which inflates the first cohort.
- **Book coverage.** The desk's recorder subscribes to markets as they list and has gaps of its own; spread and depth cover the 704 hours it captured. The on-chain fills do not depend on it.

- **Scheduled starts.** Phase is measured against the venue's listed start. 99.7% of turnover has a reliable start; matches that began late will put some pre-match flow on the in-play side of the line.

6 Conclusion

Polymarket's esports book took \$315M in September 2026 and kept almost none of it: a hold of -0.50% with an interval that spans zero. Pre-match the book earned a margin, in-play it gave it back, and 79% of the money was in-play. Customer money is extremely concentrated and the largest customers finished ahead, while liquidity supply is broad and unconcentrated. The closing price is well calibrated, and one-sided flow did not beat it; what that flow beat was the earlier price, which is where the book's pre-match losses on lopsided markets came from. The dataset makes these questions answerable from public data for the first time at this level of detail, and the same pipeline will produce the October edition. On the matches listed on both venues the prices sat about a cent apart, and Polymarket moved first on 71% of markets.

References

- Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1), 1-3.
- Cameron, A. C., Gelbach, J. B. and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3), 414-427.
- Cox, D. R. (1958). Two further applications of a model for binary regression. *Biometrika*, 45(3/4), 562-565.
- Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. Chapman and Hall.
- Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71-100.
- Hasbrouck, J. (1995). One security, many markets: Determining the contributions to price discovery. *Journal of Finance*, 50(4), 1175-1199.
- Griffith, R. M. (1949). Odds adjustments by American horse-race bettors. *American Journal of Psychology*, 62(2), 290-294.
- Snowberg, E. and Wolfers, J. (2010). Explaining the favorite-long shot bias: Is it risk-love or misperceptions? *Journal of Political Economy*, 118(4), 723-746.
- Thaler, R. H. and Ziemba, W. T. (1988). Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2), 161-174.
- Wolfers, J. and Zitzewitz, E. (2004). Prediction markets. *Journal of Economic Perspectives*, 18(2), 107-126.

A Data quality checks

The checks every build of the warehouse runs, as written by this build. Info rows are measured and shown but have no pass mark.

Table 9. Data-quality scorecard for the September 2026 build.

#	Area	Check	Observed	Target	Status
1	Integrity	Fill keys unique after merging the builds	0	= 0	Pass
2	Integrity	Overlapping builds agree fact for fact	0	= 0	Pass
3	Completeness	Foundation = VPS archive, blocks both proved (12-13 Sep)	100.00%	≥ 99.9%	Pass
4	Completeness	Foundation = independent re-pull, blocks both proved (1 Sep)	100.00%	≥ 99.9%	Pass
5	Completeness	Foundation = month-end harvest, blocks both proved (28 Sep)	100.00%	≥ 99.9%	Pass
6	Completeness	Days of the edition with fills	30	= 30	Pass
7	Validity	Fills priced strictly between 0 and 1	100.00%	≥ 100.0%	Pass
8	Validity	Fills on markets in the esports membership	100.00%	≥ 100.0%	Pass
9	Reconciliation	Maker P&L + taker P&L = 0	\$0.00	≤ \$0.01	Pass
10	Reconciliation	Flow fact turnover = fill-level turnover	\$0.00	≤ \$0.01	Pass
11	Labels	Turnover on markets with a settled label	99.94%	≥ 95.0%	Pass
12	Labels	On-chain payout = venue resolution message	100.00%	≥ 99.9%	Pass
13	Labels	On-chain payout = last traded price (≥95c / ≤5c)	99.86%		Info
14	Schedule	Turnover with a reliable scheduled start	99.74%	≥ 95.0%	Pass
15	Schedule	Pre-match turnover with a closing line	99.53%		Info
16	Capture	Chain turnover cross-checked against the desk's own recorder	79.77%		Info
17	Capture	Recorder prints confirmed by a fill on chain	99.13%	≥ 99.0%	Pass
18	Capture	Hours of the month the desk's recorder captured	704		Info
19	Integrity	Self-trades (maker wallet = taker wallet)	0		Info
20	Reconciliation	Hold on chain (maker P&L / settled turnover)	-0.50%		Info
21	Reconciliation	Hold on the venue's own prints (second tape)	-0.65%		Info

B Measure dictionary

Table 10. Every measure, its sportsbook name, definition and caveat.

Measure	Sportsbook term	Definition and caveat
Turnover	Handle / stakes	What takers staked: contracts x the price of the outcome they bought. A taker who sells outcome A at p has bought outcome B at $1 - p$, so every fill has one buyer of one outcome. Notional (USDC exchanged for the resting order's token) is also in the model; it understates stakes on the cheaper side of a market and is not the headline.
Maker P&L	Gross Gaming Revenue (GGR)	P&L of the resting order on every fill, marked to the settled payout: contracts x (price paid - payout) from the taker's side, sign reversed. Positive = the book won. Polymarket has no house. The book is every wallet with a resting order, so this is the aggregate P&L of all liquidity providers, before rebates and rewards paid off-chain.
Hold %	Hold / margin	Maker P&L / Settled Turnover. The share of settled stakes the book kept. Numerator and denominator are restricted to the same settled markets so unsettled turnover never dilutes the ratio. Hold swings with results week to week.
Settled Turnover	Settled handle	Turnover on markets with a payout: the on-chain Conditional Tokens payout first, then the venue's own market_resolved message. 50/50 splits settle at 0.5 per contract. A voided market (zero payout) is left out.
Fees Collected	Venue revenue	Fees the chain recorded on the resting orders' fills. Every esports print captured in September carried a 0 bps taker fee, so the book's hold is almost all of what takers lost.
Taker CLV %	Closing line value	For pre-match fills: contracts x (closing price of the outcome bought - price paid) / stake. Positive = the taker bought before the price moved their way. Closing line = the venue's book mid in the hour before the scheduled start when the spread was 10c or less; else the VWAP of the last 30 pre-start minutes; else the last fill in the final 6 hours. Markets that closed at 2c or 98c were played before their listed start and are left out.
Active Wallets	Active customers	Distinct wallets that took or made at least one fill in the selection. A person can run several wallets and a bot can run one for many people; a wallet is an account, not a customer.
Customer size	Customer segment	What a wallet staked as a taker in the edition: Micro under \$100, Small \$100-1k, Mid \$1k-10k, Large \$10k-100k, Whale \$100k+. Fixed for the month so a wallet never changes segment between days.
Wallet rating	Sharp / square	Closing line value over the wallet's pre-match stakes in the month: Sharp +2% or better, Square -2% or worse, Neutral in between. Wallets with under \$500 of pre-match stakes are not rated. CLV ranks wallets without bias but predicts next-period returns only weakly; read it as a rating of the month, not a forecast.
Volume made	Liquidity provided	Turnover where the wallet's order was the resting side. Makers are ranked on the whole month; MM-01 is the largest. Labels are pseudonyms and no address reaches the report.
Maker HHI	Market concentration	Sum of squared shares of volume made x 10,000. One maker = 10,000; under 1,500 reads as unconcentrated.
Liability at start	Liability / exposure	Per market, from the resting side's net position at the scheduled start: the larger of the two losses (if A wins, if B wins), floored at zero. The book is many wallets netted together; no single maker carries the whole figure.
Flow imbalance	Weight of money	Net taker contracts on outcome A / all pre-match taker contracts. 100% = every taker bought A.
Quoted spread	Overround proxy	Best ask - best bid on the venue's own book for outcome A, time-weighted within each market-hour from the research desk's websocket capture, reported as the median market. A median because thin markets with 30c+ books dominate a mean. A one-sided book has no spread and is left out.

Measure	Sportsbook term	Definition and caveat
Depth within 5c	Liquidity	USD resting within 5c of the best bid and best offer, both sides, from the full-book snapshots the venue sends on subscribe and after every trade. Snapshots arrive with trading activity; between them depth is not observed.
Capture share	Feed completeness	Share of on-chain esports turnover whose transaction the research desk's recorder also printed. The recorder subscribes to markets as they are listed; markets it never subscribed to count against it.
Label source	Settlement feed	On-chain payout > venue market_resolved message. The last traded price (95c+ / 5c-) is used only as a cross-check. Agreement between the sources is measured on the Data Quality page.

C Reproducibility

Inputs are read from the research desk's R2 lake without modifying it; the lake's catalog digest at build time was `a8c03f76ba9bd0c796435673051dfbdf...`. Extraction uses the desk's own decoders for fills and payouts. The warehouse is 39 dbt models with 36 tests on DuckDB, written as parquet marts, which also feed a ten-page Power BI report with 117 measures. This paper is generated by `tools/study/build_study.py`, which formats every number in the text from the marts at build time; bootstrap intervals use 4,000 resamples of 2,834 matches with a fixed seed, so a rebuild on the same lake state reproduces the paper exactly.