

The Gap:  
Prediction-Market Activity as an Equity Earnings Signal  
Aryan Patel

## Introduction

Over the course of this past year, prediction markets have been a blazing trend. Volume has grown about 130x between 2024 to 2025 and over \$13 billion a month. The concept is unique – the price only moves when somebody risks real money on being right. This strategy takes a specific corner in Polymarket, the “will [company] beat its quarterly earnings?” contracts. Research on these contracts finds that these contracts are more accurate than analyst consensus, free of the biases that show up in analyst forecasts, and that they move before the stock's own price catches up [1]. However, none of that is a trading edge on its own. If Polymarket's price is simply more accurate, then anybody can just use it. The real hypothesis is whether the size of the gap between Polymarkets implied probability and a company's historical beat probability predicts how a stock moves.

## Economic Hypothesis

The strategy relies on one number: the Gap. It equals Polymarket's implied probability of a beat, minus the company's own historical beat rate over the last 8 real earnings reports. A large positive Gap means that Polymarket is far more confident in a beat/miss than the company's own recent record would suggest, and that this disagreement can predict how a stock moves the day after.

This framing comes from a finding in the raw data. Across 14,477 real earnings events spanning 20 years, a beat was followed by the stock rising only 60.2% of the time and a miss was followed by the stock falling 71.5% of the time (Figure 1). Bad news moves a stock more decisively than good news is rewarded. A regression of surprise size against reaction size gives  $r^2 = 0.05$ . Even knowing a company beat or missed leaves most of the reaction unexplained. The likely reason a Gap still predicts returns is that Polymarket's price is set by a small group of especially informed trades who move fast [2]. A large Gap may capture information the stock has not fully caught up to yet.

That trailing beat rate needs enough history to mean something. A minimum of 4 quarters of historical earnings reports is required; however, the strategy uses 8 as the working number. 4, 6, 8, 10, and 12 quarters were tested and discovered catch-up rates 55.6%, 67.7%, 61.9%, 56.5%, and 56.3% respectively. Catch-up means the percentage of Gap-signaled events where the stock actually moved in the direction the Gap predicted. The 6 quarter test looks the highest but comes from the smallest sample (99 events vs. 134 with 8). The floor of 4 rarely applies as well, 667 of 671 matched events already had 8 quarters. The 4 that didn't were from recent index additions (SanDisk & Reddit) without 2 years of history yet.

Before relying on Polymarket's price, using the 679 resolved contracts grouped by price, the group leaning towards a miss actually missed 80% of the time and the group leaning toward a beat actually beat 93% of the time (Figure 2).

## Implementation

The fixed universe is the full S&P 500 (503 companies). Only 202 of these companies have had a Polymarket earnings contract. And it's uneven. The companies that have had a Polymarket contract have a median value of \$72.2 billion, while for companies that don't have a median value of \$33.6 billion. This explains why the S&P was chosen, most of the broader market has no contracts at all, and the S&P is the smallest standard list that already contains nearly every name Polymarket covers. Randomizing for 500 would mostly return companies with nothing to trade.

For context, the numbers referenced throughout this section narrow in stages: 202 of the 503 companies have ever had a Polymarket earnings contract (671 usable events); after the Gap threshold, a fresh price re-check closer to the earnings date, and basic data-quality filters, that narrows to 66 companies and 94 actual trades (49 long, 45 short).

Data comes from yfinance (earnings and prices) and Polymarket's own Gamma and CLOB APIs (contract listings and prices respectively). The codebase is in Python (using pandas, numpy, matplotlib, and scipy).

The backtest runs from October 2025 through August 2026, short because Polymarket's earnings contracts are new, the initial rollout was on September 26, 2025.

The Gap threshold (minimum gap value) is 0.20 for a long position and 0.30 for a short position. At these points the long side has a 61.9% catch-up rate while the short side has a 63% catch-up rate. And the probability climbs up as the gap increases (Figure 3). The threshold is asymmetric on purpose, evidence isn't as strong for the short side at 0.20 as it was for long positioning.

Entry timing also matters. The strategy reads Polymarket's price at the closest moment before the earnings release. This is because of the 226 events that cleared a threshold on a price up to 72 hours old, 130 no longer cleared that threshold once checked at the most recent price before earnings drop. Entry price is the last close before the earnings report drops.

Every trade uses the same fixed size: \$100. Regardless of Gap size like a real deployment which would size by percent of assets instead. A bigger Gap does predict a slightly bigger move ( $r^2 = 0.0145$ ,  $p = 0.002$  across 671 events), but that still leaves over 98% of the variation unexplained.

Two exit rules are tested in parallel. Jump buys/shorts at the close right before earnings and exits at the open right after, using standard Market-on-Close/Market-on-Open orders. Move buys/shorts also at the close right before earnings but holds until the price falls back to entry (breakeven) or 5% below its peak since entry (trailing stop). The trailing stop was tested across intervals between 3% to 15% with 5% performing the best (Figure 4). 15% also performed well but the win-rate dropped to 15.6%. A \$1,500 capital base and a 12-position cap are also set from the data. Most trades never open on the same day, but the highest amount happened to be 9 active trades (\$900 needed at once), so \$1,500 covers that with room to spare.

## Risk Assessment

Like every strategy, this strategy carries risks. Sample size is the biggest: 94 trades is a small base for every number in this paper.

Alpha decay is a second; prediction markets are growing, and signals built on informed trading moving first is exactly the kind of edge that shrinks as more capital notices it. Alpha Decay is hard to eliminate, but it can be monitored. The catch-up rate, the refinement gap between a stale and fresh PM reading, Polymarket's contract volume, and qualifying trades per quarter, with any of these moving the wrong direction an early warning to reduce size.

Third is downturn risk: the strategy was never properly tested with real Polymarket data through an actual downturn, since these contracts didn't exist during one. Across three different downturns 2008 (SPY dropped ~38%), 2018 (~6%), 2022 (~19%), the beat rate moves around, but the reaction side holds up fairly (Figure 5).

Fourth, no per-sector exposure limit exists, since no sector data is currently plumbed through this project.

And fifth, a stock split can look identical to a real move if not caught. The split risk is caught directly: any single-day move over 50% is dropped as a near-certain data artifact, both at the Jump exit and day-over-day throughout a Move hold.

Another potential limitation is that the S&P 500 list used reflects today's membership, not each trade's membership at the time. The index did change during this exact window: in June 2026, Conagra Brands, Pool Corporation, Campbell's, and EPAM Systems were removed while Honeywell Aerospace, Marvell Technology, Flex, and FedEx Freight were added, and in August 2026 Ferguson Enterprises replaced Electronic Arts. Given how short this window already is, that effect is probably small, but it is not zero. And it could potentially grow if this backtest is eventually extended across more years and passes through more index changes.

Three mechanisms doing the most work against these risks. The Move's exit's breakeven floor cuts a trade flat the moment price falls back to entry, capping how much a position that never worked can give back; its 5% trailing stop then locks in gains once a position has moved favorably, so a real winner can't fully drop back to a loss. At the portfolio level, the 12-position cap and \$1,500 capital base carry real headroom over the highest levels ever actually observed (9 positions, \$900 deployed) – a single company also can't build outsized exposure on its own, since it reports earnings only once a quarter. Whether these three actually hold up together is testable. A Monte Carlo simulation resamples the 94 real trades' returns into 5,000 new random orderings, and both exits' real realized paths sit comfortably inside that simulated range, not a lucky or unlucky outlier (Figure 6). A few supplementary checks sit on top of those three. Sizing is a flat \$100 trade, so no single trade can dominate.

## Liquidity and Capital Implications

How often this strategy finds something to trade matters as much as how well it performs when it does. The strategy performed 94 trades: 49 long and 45 short. This was across 66 companies over about 10 months. This would roughly mean 58 long and 53 short trades a year across the entire S&P 500. A real but modest opportunity set, not a high-frequency trading strategy.

## Analytical Evidence

Putting this together into a full backtest produced the results below, with SPY as the benchmark. Total return and CAGR (compound annual growth rate) are ordinary percent returns, CAGR is rescaled to a full year. The Sharpe ratio divided average returns by its volatility. Maximum drawdown is how far the portfolio ever fell below its last running peak.

|                                         | Jump    | Move    | SPY     |
|-----------------------------------------|---------|---------|---------|
| <b>Returns</b>                          |         |         |         |
| Final portfolio value                   | \$1,630 | \$1,624 | \$1,737 |
| Total return                            | +8.67%  | +8.29%  | +15.82% |
| Annualized (CAGR)                       | +10.29% | +9.84%  | +18.89% |
| <b>Risk-Adjusted Performance</b>        |         |         |         |
| Sharpe ratio                            | 2.08    | 1.46    | 1.39    |
| <b>Risk</b>                             |         |         |         |
| Maximum drawdown                        | −2.00%  | −3.46%  | −8.88%  |
| <b>Activity</b>                         |         |         |         |
| % of trading days with an open position | 31.8%   | 93.0%   | 100%    |

Both exits' portfolio value and running drawdown are charted in (Figure 7). Move leads for most of the window; however Jump closes the gap in the final weeks to finish near tied. Move's worst stretch comes late in August (-3.46%); Jump's has smaller but consistent peaks, otherwise stays relatively shallow. SPY returned +15.82% over the same stretch, more than either exit strategy. Additionally both strategies realized a beta of -0.005 for Jump and +0.038 for Move.

It's worth stating plainly: move wins on only about 33% of its trades, versus Jump's 62%. Of the 94 trades, 63 exit at break even and only 31 via the trailing stop which carry all of Move's gain, a "fewer but bigger wins" shape closer to a lottery ticket than a coin flip. Jump wins by majority – steadier, with the higher Sharpe, cleaner beta, and smaller drawdown throughout.

## Conclusion

This strategy set out to answer a question about whether the size of Polymarket's disagreement with a company's own earnings history predicts how the stock actually moves. Across the formal significance test, a date-based stability check, a complete backtest, a trading-cost model, and a Monte Carlo simulation, the answer held up. Both long and short versions of the signal predict the stock's actual reaction, at their own justifiable thresholds.

The transparency in presenting results (even with moderate returns, realistic drawdowns, and even negative periods) reflects confidence in methodology over outcomes. None of this is offered as settled — the sample size is still modest, the mechanism behind the effect remains a well-supported guess rather than a proven fact, and the strategy has never been tested through a real economic downturn. Still, although this is not the next market-beating miracle, it is a practical demonstration that a market phenomenon with real, if not yet fully proven, supporting evidence, applied systematically, can generate reliable returns in real-world trading environments.

## References

- [1] Gomez-Cram, R., Guo, Y., Jensen, T. I., & Kung, H. (2025). Financial Prediction Markets: A New Measure of Earnings Expectations. SSRN.
- [2] Gomez-Cram, R., Guo, Y., Jensen, T. I., & Kung, H. (2026). Prediction Market Accuracy: Crowd Wisdom or Informed Minority? SSRN.

## Appendix




**Figure 1: Beat/Miss vs. Stock Reaction** – Across 14,477 earnings, a beat moved the stock up only 60.2% of the time and a miss moved it down 71.5%. Beating or missing alone is a weak signal.




**Figure 2: Polymarket Calibration** – Contracts grouped by price; the miss-leaning group missed 80% of the time, the beat-leaning group beat 93%.




**Figure 3: Gap Catch-Up Rate** – Percentage of events that moved in Gap's predicted direction as threshold increases.




**Figure 5: Downturn Stress Test** – Beat and reaction rates recomputed across three real downturns (2008, 2018, 2022).




**Figure 4: Trailing-Stop Sensitivity** – Move exit's Sharpe recomputed at five trailing-stop values (3%-15%).




**Figure 6: Monte Carlo Simulation** – 5,000 reshufflings of the 94 trade returns tracing a range of portfolio paths.




**Figure 7: Equity Curve and Drawdown** – Portfolio value on a \$1,500 base for both exit rules, October 2025–August 2026, with running drawdown below.