Vera | RESEARCH

CRYPTO BRIEFING RESEARCH
Research Note No. 07 · July 2026

The Orphan Moves

When a prediction market repriced sharply, how often could our pipeline attach a news story to it? Over the four days our matching ran best, about *one in six*. What mostly separated the moves we could source was not the market. It was whether we were already watching.

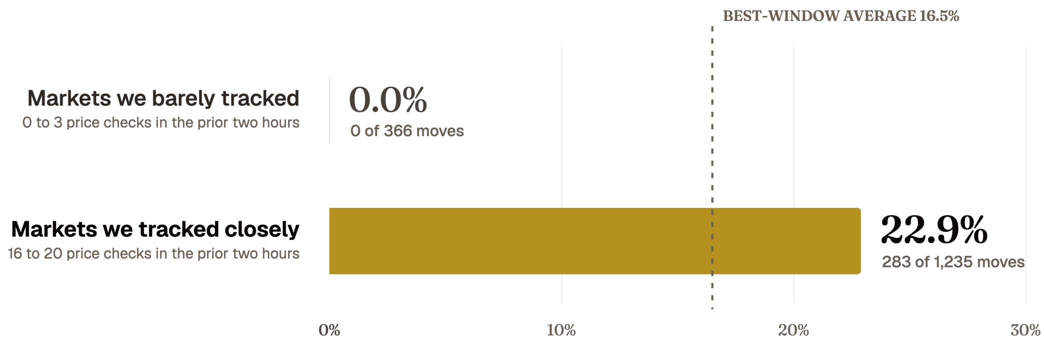

Share of sharp 3-cent moves our pipeline could attach news to, by how closely we were price-checking that market in the two hours before. June 13 to 16, 2026.

Attribution coverage, not the existence of news: the matcher only sees a 60-second window.
The matcher decayed across the wider June 13 to July 8 window and has been down since July 15, 2026.

Figure 1 · How often we could attach news to a sharp move, by how closely we were already watching. Peak-health window, June 13 to 16, 2026. In the markets our price feed barely sampled, we sourced 0 of 366 sharp moves; in the markets it sampled continuously, 283 of 1,235. The gap tracks our own coverage, not anything the market did. “Attach news” means our matcher found a likely cause within its 60-second window, so this is a floor on coverage, not a count of moves that happened without news.

THE FINDING

A market jumps and no headline explains it. The first thought is always that someone knew something. We went looking for how often that gap was real, and found something quieter: a map of where we were looking.

Over the four days our matching ran best, June 13 to 16, we recorded **1,761** sharp repricings of **3 cents or more** in one minute. For **1,470**, about **five in six**, our pipeline could attach no news within its **60-second** matching window. At its best, attribution reached about **one move in six**, and among moves of **10 cents or more**, **26 of 75**. That is a floor on what we could source, not a count of moves that happened in silence: news that led the price by minutes was invisible to the matcher.

The map of what we could explain was really a map of where we were looking.

WHAT THE MAP REALLY SHOWS

What separated a move we could source from one we could not was not the move. It was our own price tracking. We ranked each market by how many times our feed had sampled its price in the two hours before a move, then compared the extremes. In the markets our feed **barely sampled**, three checks or fewer, we attached news to **0 of 366** sharp moves. In the markets it **sampled continuously**, sixteen checks or more, we attached news to **283 of 1,235**. Same definition of a move, same window, opposite result.

IT WAS NOT THE MOVES, UNDER TEN CENTS

Over the same four peak days, **June 13 to 16**, the middle of each group was the same **3.5 cents**, and across the 3 to 10 cent range holding **96%** of the moves the averages matched too, **3.90** against **3.89 cents**. In that range a move's size told us nothing about whether we would find its story. The tail is the exception: counting every move, the sourced ones averaged **5.2 cents** and the unsourced **4.3**, and essentially all of that gap sits above **10 cents**, where we sourced **26 of 75**, about twice the rate of the bands below. (Across the full window, June 13 to July 8, the same comparison reads **4.9** against **4.6 cents**; the decaying matcher flattens it.) Below that tail, what told us was whether the market was already inside our active coverage: daytime, actively priced, the kind of contract our feed watched closely. An orphan move was not a market behaving strangely. It was a market we were barely watching.

THEY BEHAVED THE SAME AFTERWARD

If the unsourced moves were driven by something we missed, they might have snapped back once the dust settled. They did not. Across the same peak-health days, **June 13 to 16**, sourced or not, the moves behaved the same in the hour that followed: **37.8%** of the sourced moves gave back at least half their distance, against **37.0%** of the unsourced. Whether we had a story for a move carried no information about whether it would stick. Knowing why a market moved and knowing what it does next are separate problems.

WHY THIS WINDOW

Across the wider window, June 13 to July 8 and 5,551 moves, the unattached share climbs to **91.1%**. We do not lead with that number. Our matcher was decaying toward an outage across those weeks, and it has been **down since July 15**, so the later unattached moves increasingly mean the matcher was failing, not that no news existed. The four peak days are the only stretch where the rate reflects the tool working. We are reporting a measure of our own coverage, and we are reporting it from its best.

THE TAKEAWAY

THE ORPHAN MOVES · SUMMARY OF ANALYSIS

01 Over the four days our matching ran best (**June 13 to 16**), **1,470** of **1,761** sharp repricings of 3 cents or more, about **five in six**, had no news our pipeline could attach. At its best it sourced about **one move in six**.

02 That is a **floor on coverage, not a count of moves without news**: the matcher only ever looked **60 seconds** either side of a move, so news that led by minutes was invisible to it.

03 What separated the sourced moves was our own tracking, not the market. In barely-tracked markets we sourced **0 of 366** moves; in closely-tracked ones, **283 of 1,235**.

04 At peak health, sourced and unsourced moves were the **same size** below ten cents (**3.90c** vs **3.89c** across the 3 to 10 cent range that holds **96%** of moves; only the 10-cent-plus tail separates, at **26 of 75**), and they behaved the same afterward (**37.8%** vs **37.0%** gave back half within the hour). Having a story told us nothing about whether a move would stick.

05 Across the full window (5,551 moves) the unattached share is **91.1%**, but the matcher was decaying and is now down, so we lead with the peak-health window, not that contaminated figure.

Every figure records how far our own news attribution reached over a fixed past window, measured against what markets did afterward. It is a description of **past coverage and past market behavior, not a forecast** of any outcome and not a claim about the existence of news.

V VERA RESEARCH

Built on the Vera dataset: sharp prediction-market repricings on Polymarket detected from stored price history, matched against our news pipeline. Peak-health window June 13 to 16, 2026 (1,761 moves); full window June 13 to July 8 (5,551 moves). A sharp move is 3 cents or more in 60 seconds. “Attach news” means the pipeline found a likely cause inside a 60-second window; figures are a record of past attribution coverage and past market behavior, not a forecast.

vera.cryptobriefing.com