

# Same weather label, different payout

What cross-venue weather data tell us about the outcome a model should predict.

By Dejan Nešić, BroCode ([brocode.rs](https://brocode.rs))

30 September 2026

**Disclosure:** This research uses data from the [Assymetrix Data API](https://data.assymetrix.com) at [data.assymetrix.com](https://data.assymetrix.com). BroCode is an Assymetrix founding data partner, receives API access, and earns a share of revenue from customers who sign up through links in this article. [Referral link](#).

On August 3, 2026, a Miami temperature market labeled 90–91°F paid YES on Polymarket. The Kalshi contract with the same city, date and numerical range paid NO. In the neighboring 92–93°F range, the result reversed.

City, date and temperature range are enough to join two rows in a database. Establishing a hedge requires more. In the 90–91°F pair, one Polymarket YES plus one Kalshi NO would have paid \$2 under the observed outcomes; reversing the venue sides would have paid \$0. Opposite sides of the same binary payout condition would ordinarily total \$1. This illustrates the payoff difference; historical purchase prices are outside this comparison. [Verified Miami comparison](#) <sup>[1]</sup>


## Why one airport can produce two maxima

For Miami on August 3, the named airport aligns. Polymarket's rule points to Wunderground's **KMIA** daily history and the highest temperature across that date's readings. Kalshi's retained rule names **Miami International Airport** and the **National Weather Service (NWS) Climatological Report (Daily)**. These are different reporting products for the same named location. An archived pre-observation blockchain read corroborates the Wunderground instruction for the 90–91°F Polymarket contract. Kalshi's historical rule version remains unverified; its retained text and source link were retrieved in September. [Historical contract instructions](#) <sup>[2]</sup>, [native source comparison](#) <sup>[1]</sup>

The difference matters because a daily maximum can occur between published spot observations. Applying the documented five-minute averaging method to the available minute records reproduces **93°F at 16:37–16:38 UTC**. Sampling that reconstruction only at the retained Weather.com/Wunderground observation times produces a maximum of **91°F**. Miami has complete coverage of the NWS local-standard day. The peak also falls inside the local civil day, so changing the day boundary alone does not explain this case. The NWS report records 93°F, matching Kalshi's expiration value; the retained 91°F spot maximum falls in Polymarket's winning 90–91°F bucket. [Reconstruction and coverage](#) <sup>[3]</sup>, [dated climate report comparison](#) <sup>[4]</sup>

## A sampled series can miss the daily peak

Miami · five-minute reconstruction versus retained observation timestamps



Reconstruction reaches 93°F; its maximum at retained spot timestamps is 91°F.

Sources: IEM-processed NCEI one-minute averages, NWS CLI text, project Weather.com observations.

Retrospective mechanism check with complete Miami LST coverage; venue source-field selection unverified.

Figure 1. A brief peak between retained observations produces two different maxima in the same Miami station reconstruction: 93°F over the full climate day and 91°F at the sampled times. This supports the sampling mechanism; the exact resolver-used field is still unverified. Market discovery: [Assymetrix](#); analysis and chart: BroCode.

This establishes how sampling can produce the two numbers at one station. It does not establish which Wunderground field or page the resolver actually used. Our retained series is not a captured Wunderground daily-summary page, and the August 3 contract does not explicitly choose the Daily Observations table over the Day High & Low summary. The same observation payload also contains 93°F in a next-date max\_temp field; its exact historical window remains unresolved. Calling Wunderground simply a source that reported 91°F would therefore be too broad. **Different source products are established; sampling is supported; the final link to the resolver's chosen record is missing.** Field distinctions<sup>[3]</sup>, remaining determination evidence<sup>[2]</sup>

This makes the risk in our earlier [Resolution risk is the mispriced variable](#) observable: even a correct airport forecast can assign the wrong bucket probability when it targets a different statistic from the contract. The evidence does not justify declaring either venue wrong or applying a permanent temperature adjustment.

## One catalog, several kinds of evidence

We used [Assymetrix](#) to discover markets and retrieve historical trades for seven cities across August 1–30: Miami, Austin, Los Angeles, San Francisco, Seattle, Seoul and London. The five US cities supply the cross-venue payout comparisons. London contributes to the separate Polymarket quote study; Seoul lacks the local weather/book overlap used in that pilot. [City and station coverage](#)<sup>[5]</sup>

[Assymetrix](#)'s catalog and event groupings let us enumerate related markets across venues and revisit the same cohort. Venue-specific identifiers link that panel to native payouts and complete sets of temperature buckets. Our collector supplied forecasts, observations and Polymarket order books.

The provider-listed panel contains **3,120 markets and 248,399 retained trade records**, reconstructed from 3,908 hashed raw captures. The hashes and collection controls are ours; these counts do not establish a complete exchange census. [Assymetrix](#) is correcting the audited trade records. Until the corrected records are checked, the affected fields do not support complete-volume or execution-return claims. Payout findings rely on separately verified native records. [Panel and audit](#) <sup>[6]</sup>, [provider acknowledgement and correction status](#) <sup>[7]</sup>

## Matching labels, different outcomes

The first payout investigation covers ten US city-days—each a city and observation date—on August 2–3. They were selected during an earlier data review, not at random. This original sample spans only two calendar dates and cannot estimate a market-wide disagreement rate.

Eight city-days supplied matching finite numerical ranges, producing **32 pairs**. Eight pairs had opposite YES payouts, spread across four city-days. Austin on August 2 and Seattle on August 3 supplied no matching finite ranges. Native checks retained the full event ladders and explicit final payouts. [Original sample and selection](#) <sup>[1]</sup>




Figure 2. Eight pair differences across four city-days in the original sample. Neighboring ranges can disagree together when the winning bucket shifts; the rows are not independent weather events. Market discovery: [Assymetrix](#); analysis and chart: BroCode.

We then fixed all five US cities on August 5, 10, 15, 20, 25 and 30 before retrieving their additional native payouts. Market metadata was already available, so this was a systematic descriptive expansion, not preregistration or random selection. No date or city could be replaced after seeing the result.

All **30 city-days** had complete verified ladders: 330 Polymarket and 180 Kalshi markets. Nineteen city-days supplied matching finite ranges, yielding **76 pairs**. Ten had opposite YES payouts, across six city-days. The other eleven city-days had shifted grids with no identical finite ranges and remain in the selection denominator. [Expanded comparison](#) <sup>[8]</sup>

## Checking beyond the original payout cases

Five US cities · six fixed August dates · all 30 city-days retained




Figure 3. The separate expansion retains all thirty selected city-days, including eleven without directly matching finite ranges. Ten pair differences occur on six city-days; this sample must not be pooled with the original sample to estimate a prevalence rate. Market discovery: [Assymetrix](#); analysis and chart: BroCode.

Combining neighboring buckets can sometimes align the grids. Under the explicit assumption that both venues use the same integer temperature in the same unit, the ladders produce **114 nonconstant common payoff cells**: 76 direct matches and 38 requiring a basket. Twelve have different realized payouts, comprising the ten direct differences and two additional tail-basket differences. The eleven remaining city-days have only a constant common payoff under exact equality across the full integer domain; this does not rule out combinations that guarantee a minimum payout. A basket aligns boundaries; it cannot establish a common weather value. [Construction and assumptions](#) <sup>[9]</sup>

Equal realized payouts are also limited evidence. The rule audit leaves all 76 direct pairs unqualified for historical replay because complete source policies and effective rule versions are missing. Matching numerical labels alone does not certify equivalent units or contract terms. These gaps prevent a qualified replay; they do not establish that zero arbitrage opportunities existed. [Contract qualification](#) <sup>[10]</sup>

## What the other cases tell us

The other three original disagreement city-days show the same Wunderground-versus-NWS pattern. Their NWS daily maxima match Kalshi's expiration values, while lower maxima computed from retained spot readings fit Polymarket's winning buckets. Available minute windows reproduce the NWS maxima and their reported times, but Seattle and both San Francisco traces contain gaps. Only Miami supplies the complete-day reconstruction. The individual values and limits are retained in [Appendix A. Source comparison](#) <sup>[4]</sup>, [minute-level checks](#) <sup>[3]</sup>

San Francisco on August 25 and Los Angeles on August 30 introduce a different pairing: Polymarket's inspected representative contracts specify NOAA's hourly Temp column; Kalshi's current source links point to the Weather Company Climate Data Portal. Portal records corroborate Kalshi's temperatures.

The exact NOAA records used by Polymarket remain missing, and we do not know whether Los Angeles invoked its conditional Wunderground fallback. Historical rule checks cover one representative contract per case, not every sibling. [Historical rules](#)<sup>[2]</sup>, [portal records](#)<sup>[11]</sup>

The remaining expansion disagreements—Miami on August 5 and Austin on August 5, 10 and 30—also involve different named source products. Their weather records have not been reconstructed sufficiently to assign a cause. Source differences identify what to investigate; they do not by themselves explain a payout. [Expanded cases](#)<sup>[8]</sup>

London illustrates a different problem: Polymarket names London City Airport, **EGLC**, while the inspected Kalshi contracts name Heathrow, **EGLL**. This is an airport mismatch and was excluded from the US payout comparison. Within the original US cases, a different-airport explanation has not been established; matching names still leave historical sensor, window and rule equivalence to verify. [Station coverage](#)<sup>[5]</sup>, [station and time identity](#)<sup>[4]</sup>

Agreement can hide the source issue. In Seattle on August 4, a retained spot high of 80°F and an NWS daily high of 81°F both fit Polymarket's winning 80–81°F bucket. All four selected direct pairs paid NO on both venues. Those payouts cannot distinguish the underlying readings or certify equal source rules. [Seattle comparison](#)<sup>[12]</sup>

## Settlement matching is only the first test

Matching payouts does not establish a profitable trade. A separate Polymarket pilot using our local weather and order books illustrates the execution gap: at ten shares and a sixty-minute horizon, only **190 of 983** entries selected after limiting overlap had usable exits; **793** did not. Positive results after assumed fees are conditional on observed quotes, and cannot establish executable cross-venue arbitrage or a strategy's success rate. [Appendix B](#) retains the selection method, full count table, figure, fee assumptions and receipt checks. A cross-venue execution study still needs paired historical depth, effective fees and verified payoff equivalence. [Research limits](#)<sup>[13]</sup>

## From weather forecasts to contract predictions

Translate forecast distributions through the contract's observation window, sampling, rounding and bucket boundaries. Keep weather observations and native payouts as separate labels for their corresponding predictions; conflicting payouts should not be averaged into a shared outcome.

The next experiment should compare weather-only probabilities, a precisely timestamped market-price benchmark, and a model combining both. A model that accounts for the specified source should also be compared with a simpler daily-high baseline. Evaluate complete event distributions with proper probability scores, report missing inputs, and test on later dates excluded from fitting, calibration and threshold selection, with those choices frozen before evaluation. No improvement in prediction accuracy or trading performance has yet been demonstrated.

Closing the causal gap requires Polymarket's exact Wunderground or NOAA record and Kalshi's dated report version, linked to each determination. [Assymetrix's](#) catalog and historical retrieval structured this investigation. Verifying corrected trade records and obtaining qualified historical depth would support an execution study. [Determination evidence still needed](#) <sup>[2]</sup>, [historical data requirements](#) <sup>[13]</sup>

**A weather model must predict the temperature statistic selected by the contract's station, source, observation window and settlement rules.**

### About the data partner



**Company-supplied description:** [Assymetrix](#) is building a new kind of prediction market exchange, backed by a unified data layer across every major venue and an AI intelligence engine.

## Appendix A. Additional settlement records

| Case                     | Sources and recorded outcomes                                                                                                                            | Explanation supported by the evidence                                                                                                   |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| Miami, August 3          | Polymarket: Wunderground KMIA; retained spot high 91°F. Kalshi: NWS daily climate; recorded high 93°F.                                                   | Named airport aligns. Complete reconstruction shows a brief peak between sampled observations; the resolver-used field is unverified.   |
| Seattle, August 2        | Polymarket: Wunderground KSEA; retained spot high 69°F. Kalshi: NWS daily climate; recorded high 71°F.                                                   | Available minute windows reproduce 71°F; sampling at matched spot times gives 69°F. The trace has gaps and resolver linkage is missing. |
| San Francisco, August 2  | Polymarket: Wunderground KSFO; retained spot high 79°F. Kalshi: NWS daily climate; recorded high 80°F.                                                   | Available minute windows reproduce 80°F; sampling at matched spot times gives 79°F. The trace has gaps and resolver linkage is missing. |
| San Francisco, August 3  | Polymarket: Wunderground KSFO; retained spot high 77°F. Kalshi: NWS daily climate; recorded high 78°F.                                                   | Available minute windows reproduce 78°F; sampling at matched spot times gives 77°F. The trace has gaps and resolver linkage is missing. |
| San Francisco, August 25 | Polymarket: NOAA KSFO hourly Temp; winning bucket 66–67°F. Kalshi: Weather Company CLISFO; recorded high 68°F.                                           | Different hourly and daily products. The climate portal corroborates 68°F; the exact NOAA record used by Polymarket is missing.         |
| Los Angeles, August 30   | Polymarket: NOAA KLAX hourly Temp, with a conditional Wunderground fallback; winning bucket 82–83°F. Kalshi: Weather Company CLILAX; recorded high 84°F. | The climate portal corroborates 84°F. The exact Polymarket source record and whether its fallback was invoked remain unknown.           |

Table 1. Four original disagreement cases and two selected expansion cases; these add no new sample observations. Spot highs are computed from our retained series, not recovered resolver readings. Kalshi rule descriptions were retrieved in September. Source and station labels do not certify complete historical contract equivalence. [Source records](#) <sup>[4]</sup>, [contract instructions](#) <sup>[2]</sup>, [climate portal records](#) <sup>[11]</sup>, [expanded payouts](#) <sup>[8]</sup>.

For the later San Francisco and Los Angeles cases, pre-observation blockchain reads corroborate the NOAA hourly instructions for one representative Polymarket contract per case: San Francisco's 66–67°F contract and Los Angeles's 84–85°F contract. The Los Angeles rule allows Wunderground Daily Observations only if NOAA remains unavailable by 11:59 PM ET the next day. Kalshi's current source links point to the Weather Company Climate Data Portal, where dated records marked official give 68°F and 84°F respectively. The station identifiers align in those retained case records, but missing report versions, NOAA snapshots and determination records prevent a complete causal account. A fallback clause alone does not establish that it was used. [Historical rules](#) <sup>[2]</sup>, [portal evidence and limits](#) <sup>[11]</sup>

Source instructions also vary across dates on the same venue. Later Wunderground contracts explicitly prioritize Daily Observations over Day High & Low; the inspected NOAA contracts specify the hourly Temp column. Those are differences in the statistic a contract asks for. The retained text does not establish a universal amendment date, and later instructions cannot be applied retroactively to Miami on August 3. [Rule-text inventory](#) <sup>[3]</sup>

# Appendix B. Execution and fee evidence

## Quote coverage and overlap

This connects settlement risk to [Anatomy of a thin-book loss](#). A price comparison needs credible entry and exit depth before it can say much about trading economics.

A separate Polymarket pilot joins retained weather updates to order books in the five US cities and London on August 1–4. Weather enters the comparison only when the collector had received it. Books must already exist, be at most 120 seconds old, identify the correct outcome and support the requested size. The market must remain open with the same contract identity at the proposed exit. This analysis uses our local weather and Polymarket book records; trade prints do not supply historical depth. [Weather/book joins](#) <sup>[14]</sup>

Repeated updates can revisit the same token during an earlier comparison's holding period. We therefore added a retrospective sensitivity check: retain the first eligible entry for each event, token, size and horizon, then wait its entire planned horizon before admitting another. A missing exit occupies the same interval. Future price changes and payouts cannot choose which entries survive this rule. All 24 city-days and eight size/horizon combinations remain in the reports. This reduces one source of dependence; neighboring buckets and shared weather updates remain related. [Selection method](#) <sup>[15]</sup>

The ten-share, sixty-minute comparison shows why both the cost and coverage denominators matter:

| Ten shares, sixty-minute horizon    | Original entries | After limiting overlap |
|-------------------------------------|------------------|------------------------|
| Eligible entries                    | 1,679            | 983                    |
| Observed exit comparisons           | 363              | 190                    |
| Entries without a usable exit       | 1,316            | 793                    |
| Gross-positive comparisons          | 68               | 40                     |
| Positive under fee coefficient 0.05 | 55               | 32                     |
| Positive under fee coefficient 0.10 | 50               | 29                     |

Table 2. Conditional quote comparisons from the same four-date Polymarket pilot. Each fee scenario charges the consumed entry and exit levels. Missing exits receive no return value. [Independent reproduction and full results](#) <sup>[16]</sup>


The resulting 32 out of 40 is **80.0% of already gross-positive comparisons**, close to the original 55 out of 68, or 80.9%. Those 32 comparisons are also 16.8% of the 190 observed exits. Most selected entries—793 of 983—have no usable exit at the planned horizon. These conditional fractions cannot estimate executable cross-venue arbitrage survival or a strategy's success rate.

London contributes 262 selected entries at this horizon: 67 observed exits, fourteen gross-positive comparisons and eleven positive under coefficient 0.05. This is an international contribution to the conditional Polymarket study, without evidence that London offers a stronger signal or a matched Kalshi hedge. The full size, horizon and city results include the less favorable cases. [All comparison groups](#) <sup>[17]</sup>

## Limit repeated entries, keep the missing exits

Polymarket weather pilot · Six cities, August 1–4, 2026 · 10 shares, 60-minute horizon

### Exit evidence among eligible entries



### Gross-positive moves surviving the fee scenario

**55/68 = 80.9% → 32/40 = 80.0%**

A conditional repricing count; no arbitrage estimate.

The first entry reserves its planned horizon even when its exit is missing. Other tokens and weather updates remain dependent.

Fee coefficient  $r=0.05$  in  $q \times r \times p \times (1-p)$ , charged on both legs; historical policy unverified. Missing exits are unknown.

Sources: project weather/orderbook archive and reproducible diagnostic; Assymetrix supplied the wider market catalog and history panel.

### Positive after assumed fees / usable exits




Figure 4. The overlap sensitivity retains 983 entries, including 793 missing exits. Of 40 observed gross-positive comparisons, 32 remain positive under the 0.05 fee assumption. The cost result is conditional on the observed subset. Market discovery: Assymetrix; analysis and chart: BroCode.

## Dated fees and provider checks

The scenarios extend [Why uncertainty makes market making more expensive](#). Our archived SDK configuration used coefficient 0.10 and exponent one; September market-specific checks returned 0.05. A later configuration cannot certify what applied to every August fill, so both remain explicit assumptions in this comparison. [Fee evidence](#)<sup>[18]</sup>

For these scenarios, the fee at each consumed level is  $\text{shares} \times r \times \text{price} \times (1 - \text{price})$ , rounded to five decimal places. The coefficient  $r$  shapes the fee curve; it is not a flat percentage of notional. This models costs on displayed depth, without reproducing an actual order's match sequence or account rebates. [Fee-scenario method](#)<sup>[15]</sup>

Two historical receipts give narrower, firmer evidence. A SELL of 13.46 YES shares at 0.999 paid a verified fee of 0.000670 collateral. A separate BUY of 6.96 NO shares at 0.43 cost 2.9928 collateral plus 0.08529 in fees, with all 6.96 shares delivered. Those reconciliations establish the cash and token flows for the two matches. They do not certify a market-wide coefficient, rounding policy or effective interval. [SELL reconciliation](#)<sup>[19]</sup>, [BUY reconciliation](#)<sup>[20]</sup>

The saved trade audit found fewer provider trades in two Kalshi windows, price and side differences for one Polymarket SELL, and an absent independently verified BUY. The September 9 recheck reproduced the September 7 array of 160 trade rows. These observations describe those saved captures. [Dated provider recheck](#)<sup>[21]</sup>

Corrected trade records from [Assymetrix](#) will follow; we will repeat these comparisons and record any changes in a dated update, preserving the original captures. [Provider acknowledgement and correction status](#)<sup>[7]</sup>

# Evidence notes

Supporting tables and method notes are included in the optional evidence bundle. The numbered notes identify the evidence cited in the article.

1. **Native payout comparison, August 2–3.**  
Verified Miami comparison · native source comparison · Original sample and selection
2. **Venue determination records and rule evidence.**  
Historical contract instructions · remaining determination evidence · Historical rules · Determination evidence still needed · contract instructions
3. **Minute reconstruction, sampling and source-field audit.**  
Reconstruction and coverage · Field distinctions · minute-level checks · Rule-text inventory
4. **Daily climate and observation source comparison.**  
dated climate report comparison · Source comparison · station and time identity · Source records
5. **City and station coverage method note.**  
City and station coverage · Station coverage
6. **Provider panel and trade audit summary.**  
Panel and audit
7. **Provider response and remaining verification — September 24.**  
provider acknowledgement and correction status · Provider acknowledgement and correction status
8. **Systematic expansion: selection and payout summary.**  
Expanded comparison · Expanded cases · expanded payouts
9. **Common payoff cells and basket construction summary.**  
Construction and assumptions
10. **Contract qualification and evidence readiness summary.**  
Contract qualification
11. **Weather Company daily climate portal evidence.**  
portal records · climate portal records · portal evidence and limits
12. **Seattle records and rule requirements method note.**  
Seattle comparison
13. **Historical data access and remaining research requirements.**  
Research limits · historical data requirements
14. **Weather and orderbook join method note.**  
Weather/book joins
15. **Pilot overlap sensitivity and fee-scenario method.**  
Selection method · Fee-scenario method
16. **Independent pilot reproduction and full results.**  
Independent reproduction and full results
17. **Pilot comparison groups by city, size and horizon.**  
All comparison groups
18. **Dated fee evidence method note.**  
Fee evidence
19. **Historical SELL receipt reconciliation.**  
SELL reconciliation
20. **Historical BUY receipt reconciliation.**  
BUY reconciliation
21. **September 9 provider trade recheck method note.**  
Dated provider recheck