RESEARCH PERSPECTIVE

Beyond Sports: Weather and Crypto Event Contracts

A broader research universe requires new models of information and resolution

In our worked example, moving the effective boundary of an “81°F or higher” contract by half a degree changes the probability that it pays by about ten percentage points. Weather and crypto event contracts settle on a measurement, such as a station’s official report or an index average at a stated time, so a forecast that is right about the weather or the asset can still be wrong about the contract. We set out what would have to be rebuilt before a market-making approach could extend to these categories.

By Chris Park, Chief Investment Officer

Current edition: October 27, 2026 · 6 min read

First published October 18, 2026; revised October 27, 2026 to add key takeaways, Exhibit 1, current settlement conventions and published category evidence.

KEY TAKEAWAYS

01 Price the settlement variable, not the phenomenon. Station, reporting day, rounding, averaging window and the treatment of equality are part of the payoff.

02 Near a threshold, conventions can matter as much as forecasting skill: in Exhibit 1, moving the effective boundary by half a degree changes the probability by about ten percentage points.

03 Published studies find favorite–longshot pricing patterns in weather and crypto contracts, but their size depends on venue and measurement method, and in weather even the sign depends on how returns are weighted. Each category needs its own evidence before any expansion.

> Weather and crypto event contracts are future research areas, outside the current live sports scope.

▼ **Method, sources and limitations**

Study type	Future research agenda
Method	Primary-document examples, published category evidence and original conceptual analysis, with hypothetical arithmetic. No category-specific strategy result is reported.
Data period	Documentation reviewed October 27, 2026; no empirical trading dataset.
Source categories	Official weather-market and resolution documentation; Published prediction-market research
Limitations	Weather, crypto event contracts and other non-sports categories are future research, not current live trading. An event-contract thesis does not establish forecasting skill, profitable execution or portfolio diversification.

What can transfer, and what must be rebuilt

The basic market-making problem transfers across event types: offer prices, observe fills, manage inventory and reconcile the outcome. Contract understanding, executable size and capital use remain essential. But the source of uncertainty changes, so a quoting model built for one category cannot simply be relabeled for another.

A sports contract responds to a sequence of plays. A temperature contract depends on a specific station's official daily report. A reference-price contract depends on an index value or average at an exact time. Each has a different information clock and a different way for a reasonable forecast to disagree with the eventual payout.

The research advantage we intend would come from understanding those distinctions well enough to price and manage liquidity responsibly. Whether that produces attractive net economics is a separate empirical question.

Weather markets settle on a measurement

A question about a city's high temperature can look simple while containing several material definitions. Which station or data source? Which reporting day? Fahrenheit or Celsius? What rounding convention? Is the relevant value a preliminary observation or a final published report? Those details define the contract.

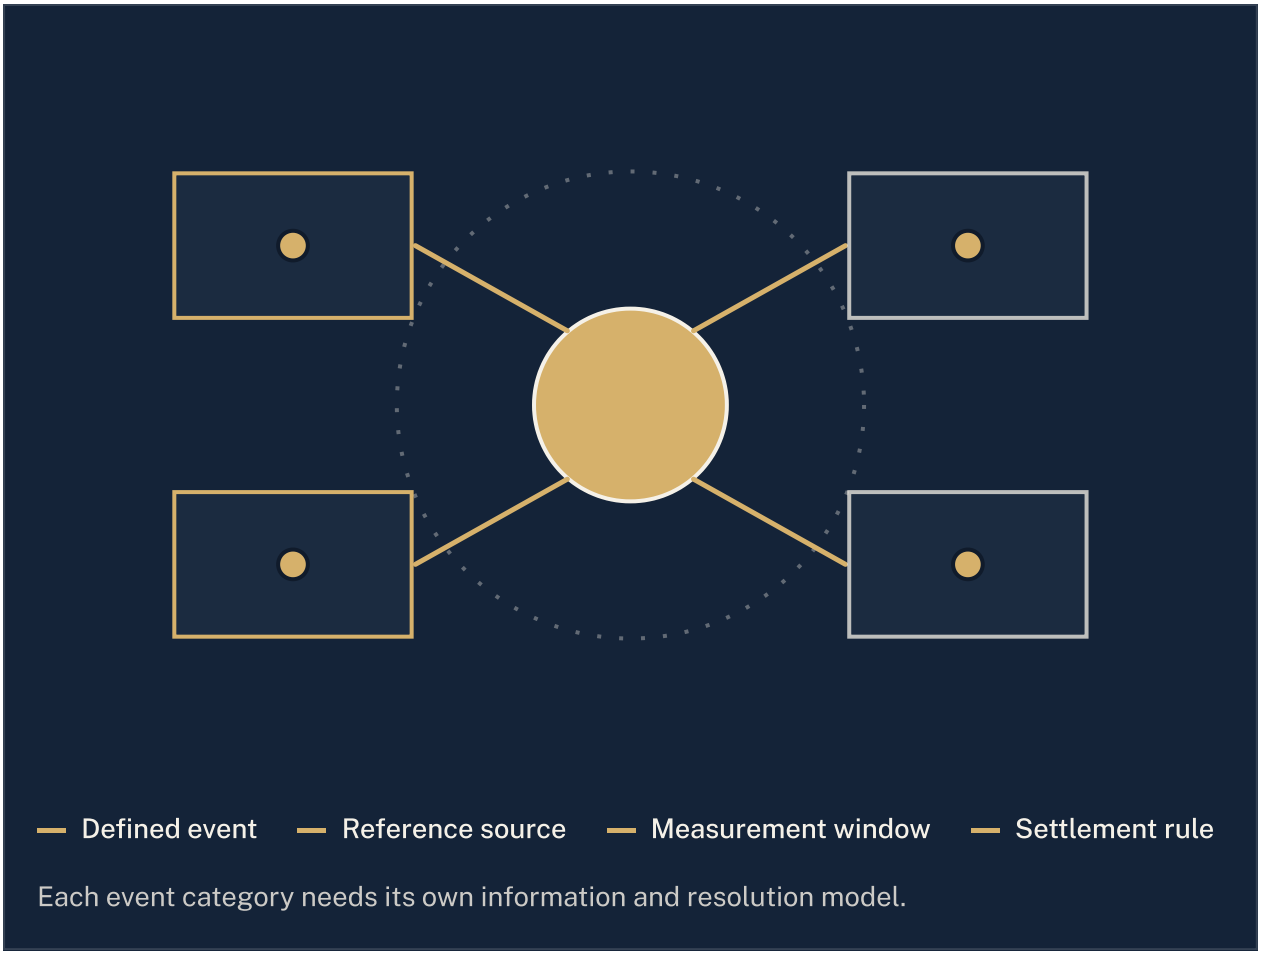

Each event category needs its own information and resolution model.

Under Kalshi's weather guidance, daily temperature markets settle on the final National Weather Service Daily Climate Report for the listed station, typically the next morning. Hourly markets settle on The Weather Company's reading at named station coordinates, about 25 to 35 minutes after close. Because NWS climate reports use local standard time, the reporting day during daylight saving time runs from 1:00 a.m. to 12:59 a.m. the following day. Other weather contracts can differ, which is why each market's terms must be read.

A consumer weather app may be useful context without being the settlement authority, and a nearby station may record a different temperature. An apparently better forecast can still predict the wrong object.

A weather forecast needs a distribution around the threshold

A point forecast of 80 degrees does not determine the probability that a contract settles above 81, at 80 or within a range. Near a threshold, small differences in

rounding or observation method are economically significant. Exhibit 1 holds one hypothetical forecast fixed and moves only the effective boundary.

EXHIBIT 1

One forecast, three boundary conventions

Effective boundary for "81°F or higher"	Probability the contract pays
True high of 80.5°F or more	47.8%
True high of 81.0°F or more	36.9%
True high of 81.5°F or more	27.1%

Forecast distribution: normal, mean 80.4°F, standard deviation 1.8°F.

Source: Predict Here Capital calculation. All inputs hypothetical. The relevant boundary is whatever the settlement source's measurement and reporting conventions imply.

In this hypothetical distribution, shifting the effective boundary by half a degree changes the probability by about ten percentage points. Getting the convention right is part of pricing.

Forecast uncertainty also changes as the reporting window progresses, so models should be tested with the forecasts and observations available at decision time, including revisions and publication delays. Inventory has structure as well: adjacent temperature ranges can be mutually exclusive or overlapping, and one weather system can move several cities together. More locations do not automatically mean independent risks.

Crypto event contracts are defined by a reference question

A crypto event contract might ask whether a named reference price is above a threshold at a stated time, or whether an end value is above a start value. The relevant exposure is the contract's stated payoff, and the research concerns the probability and execution of that payout, including how information about the reference price reaches the market.

The exact index, observation time, averaging window, missing-data policy and treatment of equality all matter, and conventions differ even within one brand. Polymarket US's resolution guidance gives the example of a contract settling on the average of CF Benchmarks' Bitcoin Real-Time Index over the final minute before the settlement time. Polymarket US's gold and silver markets use Chainlink spot values at the start and end times, not averages. Polymarket International's up/down markets compare Chainlink time-weighted averages at the start and end and resolve Up when the final value is equal to or greater than the starting one, so equality pays Up.

A chart from a different exchange cannot settle a contract merely because it tracks the same named asset, and these mechanisms should not be combined into one generic rule.

Near-expiry sensitivity changes the quoting problem

When little time remains, a small move in the reference price can produce a large change in the probability of a threshold outcome, but only while the outcome is still in doubt. Late-Game Sensitivity shows that the same mathematics governs a close game's final minutes: sensitivity peaks when the distance to the threshold is small relative to the uncertainty that remains.

A maker's research must therefore address timing and measurement as well as direction. How close is the reference to the threshold? How much of an averaging window has elapsed? Are the available data contemporaneous with the settlement source? Averaging reduces the effect of each late price move on the final average, but the uncertainty left in the average shrinks too, so a close outcome can stay highly sensitive until the window ends.

A hedge in another instrument, if contemplated, would need its own access, cost and risk analysis. Its payoff may not match the event contract, particularly near expiry or in disrupted markets. The existence of an underlying asset does not make the event position automatically hedgeable or neutral.

Other event markets introduce other definitions of finality

Economic releases distinguish initial estimates from later revisions. Political contracts can distinguish a projected outcome from certification or taking office. A company announcement can differ from implementation of what was announced. A contract should say which milestone governs the payout.

These examples point to a reusable process, and we claim no universal forecasting skill. Identify the official source, map the publication sequence, preserve information availability and examine how quotes respond, then test execution and inventory outcomes under the specific venue's rules. A longer event horizon also changes capital requirements, because a position can keep an attractive expected margin while consuming funding that could be used elsewhere.

What published evidence says about these categories

Independent studies now cover both categories. On Kalshi from 2021 to April 2025, Bürgi, Deng and Whelan (2026) reject the hypothesis that prices are unbiased forecasts in each category they examine, including crypto and the combined climate-and-weather category. Low-priced contracts won less often than their prices implied. On Polymarket International from November 2022 to March 2026, Cardozo and Rivero-Wildemauwe (2026) find the classic pattern of losing longshots and winning favorites in crypto under both of their main weightings. In weather, the longshot result changes sign depending on whether markets or dollars are weighted equally.

Both are working papers or preprints covering different venues and periods, and both report average returns under stated sampling and weighting methods rather than the returns of any trading strategy. They show that pricing patterns exist and that they are measurement-dependent. Neither establishes an edge for a market maker. Who Earns the Spread? reviews the evidence in detail.

Expansion requires a separate evidentiary case

Each new category needs a defined contract universe, information model, execution assumptions and inventory policy, supported by source rights and lawful access. We would reject or revise the thesis if it predicts a proxy instead of the settlement variable, relies on information that arrived too late or loses its apparent value after execution and capital costs.

Sources and further reading

Sources reviewed October 27, 2026. Venue rules, fees and programs change; the specific contract and account terms govern. Views are the author's as of the edition date shown above.

1. Kalshi Weather Markets. Kalshi Help Center.
2. Polymarket US Market Resolution. Polymarket US documentation.
3. Polymarket resolution. Polymarket documentation.
4. Bürgi, C., W. Deng and K. Whelan (2026). Makers and takers: the economics of the Kalshi prediction market. Working paper, University College Dublin, January 2026.
5. Cardozo, M. and J. I. Rivero-Wildemauwe (2026). The favorite–longshot bias in prediction markets: evidence from Polymarket. Preprint, arXiv:2609.12878, September 2026.