Improvement of Daily Temperature Forecasts from a Prediction Market

Patrick T. Brown[1,2,3]

[1]Trend Analytics, LLC and Interactive Brokers, Greenwich, CT, United States
[2]Energy Policy and Climate Program, Johns Hopkins University, Washington, D.C., United States
[3]The Breakthrough Institute, Berkeley, CA, United States
ORCID 0000-0002-5058-1718

A non-peer-reviewed preprint submitted to EarthArXiv

Abstract

Prediction markets aggregate and distill the collective views of participants who have a financial stake in being correct. They have been shown to produce more accurate forecasts in realms like election results, which are not easily amenable to mechanistic modeling because they entail forecasting human behavior. It is far less clear that prediction markets should improve the accuracy of predictions of physical quantities where credible mechanistic models already exist, and their output is further routinely augmented with statistical post-processing and expert human oversight. Here I evaluate ForecastEx prediction markets on the daily maximum temperature at 24 US locations from 11 February 2026 to 1 October 2026 (5,405 location-days) against 20 operational forecast systems at every hourly lead from 36 hours before the end of the target day. The alternatives include state-of-the-art numerical weather prediction models with biases removed, model output statistics applied to those models, optimal model blends, an artificial intelligence ensemble, and human-supervised forecasts with access to all the aforementioned information. I converted prediction market probabilities to a central value for comparison with alternative deterministic systems that output only a single value, and I compared them using mean absolute error. Where alternative systems provided probabilistic information, I applied basic ensemble model output statistics corrections and compared them to the prediction market probabilities using the continuous ranked probability score (CRPS). I made substantial efforts to verify alternative systems against the observations that align with their forecasts when they differed from the observations that the prediction market targeted. Whether lead time was counted to the end of the target day or to the time the daily maximum was observed, the prediction market central forecast value and CRPS had the lowest error at all meaningful lead times. At 6:00 PM local time on the evening before the target day, when none of the day's temperatures had been observed, the prediction market's mean absolute error was 1.49°F, compared with 1.76°F for the closest alternative at that lead, the NBM station bulletin (17% [14%, 19%] lower), and its CRPS was 1.06°F compared with 1.19°F for the best EMOS-corrected probabilistic alternative (11% [8%, 15%] lower). The margins generally held or improved through the target day, when the market's responsiveness to live information becomes more of an advantage. The advantage was a genuine gain in accuracy rather than an artifact of forecast latency, since the market was more accurate than each alternative even when sampled only at the times when the alternative issued a new value. The prediction market's performance improvement over model output statistics deliberately designed to remove systematic biases, and over human-expert systems with access to the same information, indicates that the market combines information and potentially incorporates additional information more flexibly and dynamically than conventional alternatives.

1 Introduction

1.1 Prediction markets

A prediction market allows trading on predictions about the future. The most prominent contracts listed on contemporary prediction markets are state-price securities or Arrow-Debreu securities, or event contracts (under current CFTC nomenclature) that pay a fixed amount (typically $1 per contract) if a specified event occurs. If traders hold beliefs about that event and buy or sell until the price matches what they think the contract is worth, the price can be read as the market's probability of the event [1, 2]. For example, a contract that pays one dollar if today's high temperature exceeds 85°F and trades at 70 cents implies a 70% chance that it will. This reading requires assumptions about traders' risk preferences and the distribution of their beliefs [3], but it provides a simple way to turn dispersed judgments into a single number.

The idea that a group can estimate a quantity better than its individual members was most famously formalized by Galton [4], who collected 787 entries in a contest to guess the dressed weight of an ox and found that the median guess, 1,207 lb, was within 1% of the true weight of 1,198 lb. Surowiecki [5] argued that such collective estimates are accurate when individual judgments are diverse, independent, and decentralized and when a mechanism like prediction markets exists to combine them. Sunstein [6] also highlighted prediction markets as one such mechanism, with an advantage over group deliberation in that traders are rewarded for acting on what they privately know. Wagers serve a related purpose, because they turn a disagreement into a precise claim about the future that will be tested, as in the 1980 wager between Julian Simon and Paul Ehrlich on the prices of five metals [7]. A prediction market extends this idea from two parties to many [8, 9].

Overall, a prediction market has two features that a single forecaster lacks. It combines many sources of information [10, 11], and it continually reallocates influence toward participants whose past judgments have been accurate, because those participants accumulate the confidence and capital to trade further [12–14].

Empirical support for these arguments comes mostly from outcomes that depend on human choices. For example, prices on the Iowa Electronic Markets were closer to the eventual presidential vote share than contemporaneous polls in most comparisons across five elections [15], an advantage already visible in the first of those markets [16] and in the organized betting markets that forecast US presidential elections before scientific polling existed [17]. Surveys of this literature find market prices quick to incorporate new information and generally more accurate than polls and professional forecasters [18]. However, debiased and aggregated polls narrow or reverse this advantage in some analyses [19, 20], as do statistical models built without markets at all [21], and in a two-year geopolitical tournament, team-based probability polls with weighting and recalibration outperformed the market run alongside them, although the market was more accurate than the simple average of those same forecasters [22, 23]. Internal corporate markets have forecast sales and project milestones with good calibration [24, 25], and markets among researchers predicted which published psychology results would replicate more accurately than a survey of the same participants [26]. At the level of whole exchanges, prices are close to calibrated [14, 27], with a documented tendency to overprice low-probability outcomes [28, 29]. On the largest US-regulated exchange, calibration improves with trading volume, the number of traders, and proximity to resolution [14], and markets on macroeconomic releases matched or exceeded professional consensus forecasts for several indicators [30].

These settings share a feature that favors markets in that competing forecasts (i.e., polls, surveys, and analyst consensus) are built from human judgment, with incomplete information and cognitive biases, so a mechanism that averages across many people can fill in what any one lacks and offset idiosyncratic biases. It is less obvious that the same mechanism that adds value in those domains would also add value when the competing forecasts come from physics and statistics.

1.2 Weather forecasting as a benchmark

Daily weather forecasting is one of the most mature applications of physics and statistics to prediction. Operational numerical weather prediction integrates the governing equations of the atmosphere forward from an observed initial state, and its skill has increased by about one day of lead time per decade [31]. Model output statistics then correct each model's systematic errors at individual stations [32, 33]. Averages of several such products are more accurate than any one of them [34, 35] (an analogy to aggregation benefits of markets discussed above), and the National Blend of Models now provides a calibrated multi-model product that serves as the starting point for the official National Weather Service forecast [36, 37].

Human forecasters have also long served as the procurers of the final forecast [38, 39] (e.g., the computer-based tools above are often referred to as "guidance" in the meteorological jargon, which conveys that they are thought of as a menu of tools that can be used to guide a human forecaster). As automated "guidance" has improved, however, the value that human forecasters add to it has been questioned [40, 41].

Ensemble prediction systems handle uncertainty by running a model many times from perturbed initial states [42, 43]. Raw ensembles often contain mean biases and are underdispersive at individual stations, so statistical post-processing improves their forecasts [44, 45]. Most recently, artificial intelligence models trained on decades of reanalysis have matched or exceeded the leading physics-based global models on many headline scores [46–50].

Thus, a market that trades on a quantity these systems forecast faces a demanding benchmark. Participants would need to combine information or bring new information to the market in a way that a well-verified body of physics, statistics, and human expertise does not.

1.3 Markets and weather

Evidence on whether markets have already demonstrated added value in weather and climate domains remains limited. Roll [51] found that returns on frozen concentrated orange juice futures predicted errors in National Weather Service temperature forecasts for central Florida, and later work confirmed that these prices respond to freezing-temperature news [52]. Those futures price the orange crop, however, so their temperature information was indirect. Weather derivatives, which pay according to temperature indices at specified cities, were designed for hedging, and their prices likewise carry forecast information only indirectly [53–56]. Legal scholars, economists, and climate scientists have proposed prediction markets on climate outcomes [57–61], but few have operated, and an early market on the following year's global mean temperature anomaly, listed on the commercial exchange Intrade in 2011, attracted little trading [56, 62].

The most extensive study of expert prediction markets in this domain is that of Roulston and Kaivanto [56], who ran 24 markets for invited experts between 2018 and 2021 on UK monthly temperature and rainfall, the Niño-3.4 sea surface temperature anomaly, Atlantic hurricane counts, and UK wheat yield, at horizons of 2 to 12 months, with trades made against a subsidized automated market maker [63]. Their markets incorporated new information as each approached its close, carried more information than climatology, and produced forecasts consistent with reliability given the 18 to 24 forecasts available at each horizon. The same group has since run markets on seasonal hurricane activity and on joint outcomes [64, 65]. These markets were benchmarked against climatology, which is the natural reference at monthly to annual horizons, and each resolves only once, so evaluating them takes years [56]. Until this year, no study had compared market prices with operational numerical weather prediction and post-processed guidance for the same quantity using a common forecast sample.

The recent emergence of daily temperature prediction markets now makes such a comparison possible. They trade on US-regulated exchanges at dozens of stations and resolve every day, so a single season produces thousands of verifiable forecasts for a quantity that operational centers forecast explicitly. They also differ in design from the expert markets described above. Anyone who meets an exchange's account requirements can trade with their own capital, orders meet in a continuous order book, and liquidity comes from participants, including market makers, who trade for their own account.

The closest work to the present study is being conducted concurrently by Crosier [66], who compares Kalshi daily maximum temperature markets in seven US cities over five years with six public forecasts. They find the market's root-mean-square error is about 10% below the National Blend of Models at the end of the market's first hour of trading and 11% below at the Blend's final bulletin on the target morning, and 3% to 4% below an optimally weighted combination of the public products refitted daily on past data. That study and this one differ in exchange, settlement definition, period, location, and benchmark set, yet, as shown below, they reach similar conclusions.

1.4 Purpose of this study

This study examines whether a public prediction market can effectively aggregate information and generate more accurate forecasts than traditional systems for daily maximum temperatures.

There are potentially two bars to cross, representing different levels of achievement.

In the first bar, the market could outperform relatively generic, non-customized forecasts, such as those from numerical weather prediction models. In that case, the market would remove the biases those generic forecasts carry at each location as an emergent property of trading, which would still be a non-trivial achievement given the lack of barriers to entry and coordination among participants.

The second bar would entail outperforming the station MOS products and the human experts with access to the same underlying information, judged against the verification datasets they target and sampled at the time these systems issue new forecasts. In that case, the market would be finding a better way to combine and incorporate information to maximize forecast accuracy.

To test these possibilities, I compare ForecastEx daily maximum temperature market probabilities at 24 locations with 20 operational forecast systems over eight months. I compare forecasts with

observations under deterministic and probabilistic frameworks. The first compares central values, the median of the prediction market's price ladder against each alternative's forecast. The second compares full probability distributions, for the market and the four alternative systems that publish ensembles and thus probabilistic information. Each is evaluated on an absolute and a relative timeframe, the hours remaining before the end of the target day and the hours remaining before the maximum, respectively. I also toggle the verification observations for the alternatives and compare the prediction market to alternatives sampled only at the times when the alternatives had just released a new forecast.

2 Data and methods

2.1 Markets and settlement

ForecastEx is a designated contract market registered with the US Commodity Futures Trading Commission. It lists daily contracts on the maximum and minimum temperature of a location-local calendar day. For each location and day, the exchange lists a ladder of contracts at whole-degree Fahrenheit strikes, and a 'yes' contract on the maximum pays if the settlement value is strictly greater than its strike [67].

The settlement value is the location's daily maximum in whole degrees Fahrenheit from the METAR observations of its Automated Surface Observing System [68], as published by Weather Underground, taken once the first observation of the following day is available [67].

ForecastEx first listed daily high-temperature markets in November 2025 but changed the settlement protocol on 11 February 2026. This study examines forecast skill starting from that date, from 11 February 2026 to 1 October 2026. It does so across 24 US airport weather stations.

I focus on the daily maximum temperature. The minimum-temperature contracts account for 1.9% of all contracts traded in the record, with a median of 451 per location-day, compared with 12,813 for the maximum. When trading volume is low, prices reflect standing market-maker quotes more than prices driven by active trading among many participants, so the minimum market is a weak test of this study's research questions.

2.2 The prediction market's forecast

Each price is read as a probability, so a contract at the 85°F strike priced at 58 cents implies a 58% probability that the maximum exceeds 85°F (Fig 1). From 3 July 2026 the exchange's order books were captured every ten minutes, and the Yes price of a contract is the midpoint between the best Yes bid and one dollar minus the best No bid,

$$p = \frac{b_{\text{Yes}} + (1 - b_{\text{No}})}{2} \qquad (1)$$

with a missing side set to its limit, one cent for a Yes bid and 99 cents for a No bid, and a book holding only those limits treated as empty. At each hour, I use the last capture within that hour. Before that date, prices come from the exchange's published trade record, with each strike's last traded price carried forward hourly. Prices are not adjusted for fees. Over the window where both

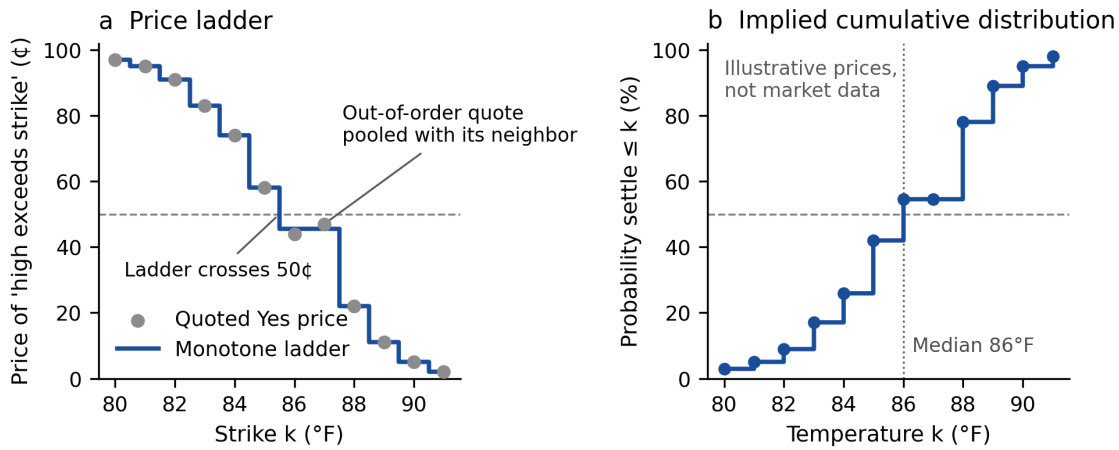

Fig. 1. An illustrative exceedance price ladder and the forecast it implies. (a) Prices of contracts that pay if the daily maximum exceeds each strike (grey), and the ladder after enforcing that prices fall as the strike rises (blue). The out-of-order quote at 87°F is pooled with its neighbor. (b) The cumulative distribution implied by the monotone ladder and its median, which is the market's central value.

sources exist, 3 July 2026 to 1 October 2026, 2,131 location-days on 91 dates, the two constructions give the same central value at 82% of instants, and the difference in the market's error at the six leads reported here spans zero at every lead. In the last six hours of the target day the trade record is the worse of the two, by 0.070°F to 0.096°F at every hour, so the market's scores at the shortest leads are conservative while most location-days are priced from it (3,241 from the trade record, 2,164 from the order-book capture).

Contracts at adjacent temperature values in the exceedance ladder (strikes) are priced by independent markets, so an exceedance ladder can briefly violate the ordering a probability distribution requires, as at the 87°F strike in Fig 1a. I enforce monotonicity by pooling adjacent violators and close the distribution at the lowest and highest listed strikes. The monotone ladder gives the probability that the settlement value falls at or below each whole degree k, $F(k) = 1 - P(k)$, where $P(k)$ is the monotone Yes price at strike k (Fig 1b). The market's central value is the median of this distribution, the temperature at which the ladder crosses 50 cents, rounded up to match whole-degree settlement.

2.3 Alternative forecasts

Figure 2 shows the 20 operational systems I compare against the market, grouped into five categories (Table 1).

One category consists of twelve raw numerical weather prediction outputs that forecast a grid-box average.

Another category includes five model output statistics (MOS) products that downscale raw numerical weather model output to point values [32, 33] at the same weather stations used for ForecastEx contract settlement.

Three systems represent their own categories. One is the National Blend of Models, which is a calibrated multi-model blend on a 2.5 km grid [36]. Another is the National Weather Service's

Table 1. Forecast systems. Update frequency is the provider's published value or, where run times were recorded, the observed value. Four systems publish ensemble members (Section 2.3). Target gives the observation each system's daily value is verified against, either the METAR settlement value (METAR), the National Weather Service climate report (Climate), or the Unrestricted Mesoscale Analysis (URMA), as Section 2.5 explains. Rows marked with b were reconstructed from an archive of previous runs before live capture began [69]. Record start is the first date scored for the daily maximum.

System	Description	Updates	Target	Record start
ForecastEx market	Daily contracts on each location's maximum [67]	10 min	METAR	12 Feb
Station MOS				
LAMP	Localized Aviation MOS Program, updated with the latest observations [33]	1 h	METAR	12 Feb
NBM station MOS (NBS)	NBM station bulletin, daytime maximum [36]	6 h	Climate, METAR	12 Feb
GFS MOS	Model output statistics from GFS, issues a daytime maximum [32]	6 h	Climate, METAR	12 Feb
NAM MOS	MOS from NAM, daytime maximum [32, 70, 71]	12 h	Climate, METAR	12 Feb
DWD MOSMIX [a]	DWD guidance from ICON and ECMWF IFS [32, 72]	6 h	METAR	18 Jul
Blend				
National Blend of Models	NOAA calibrated multi-model blend, 2.5 km grid [36]	1 h	METAR, URMA	27 Jun
Human-edited				
National Weather Service	Official forecast from the National Digital Forecast Database [37], daytime maximum [73]	About 12 h	Climate, URMA	14 Jul
Artificial intelligence				
ECMWF AIFS [b]	ECMWF artificial intelligence system, 51-member ensemble when live [74], 0.25°	6 h	METAR	12 Feb
Raw numerical weather prediction				
ECMWF IFS [b]	ECMWF global model, 0.25° [75, 76]	6 h	METAR	12 Feb
ECMWF IFS ensemble mean	Mean of the 51-member ECMWF ensemble [75, 76]	6 h	METAR	27 Jun
NOAA GFS [b]	Global Forecast System, 0.11° [77, 78]	6 h	METAR	12 Feb
NOAA GEFS	Global Ensemble Forecast System, 31 members [79]	6 h	METAR	16 Jun
DWD ICON [b]	Global model of the German weather service, 0.1° [80]	6 h	METAR	12 Feb
DWD ICON-EPS	ICON global ensemble, 40 members [80, 81]	12 h	METAR	16 Jun
ECCC GEM [b]	Canadian Global Environmental Multiscale [82]	6 h	METAR	12 Feb
ECCC GEPS	Global Ensemble Prediction System, 21 members [83, 84]	12 h	METAR	16 Jun
Met Office UM [b]	Unified Model, global configuration [85]	6 h	METAR	12 Feb
Météo-France ARPEGE [b]	Météo-France global model, 0.25° [86, 87]	6 h	METAR	12 Feb
JMA GSM [b]	Global Spectral Model, 0.5° [88, 89]	6 h	METAR	12 Feb
NOAA HRRR	High-Resolution Rapid Refresh, 3 km [90]	1 h	METAR	14 Jul

[a]MOSMIX has no station for Buckley Space Force Base (Denver), Houston Hobby, or Chicago Midway, so it is scored at 21 stations. [b]Reconstructed from an archive of previous runs before live capture began [69].

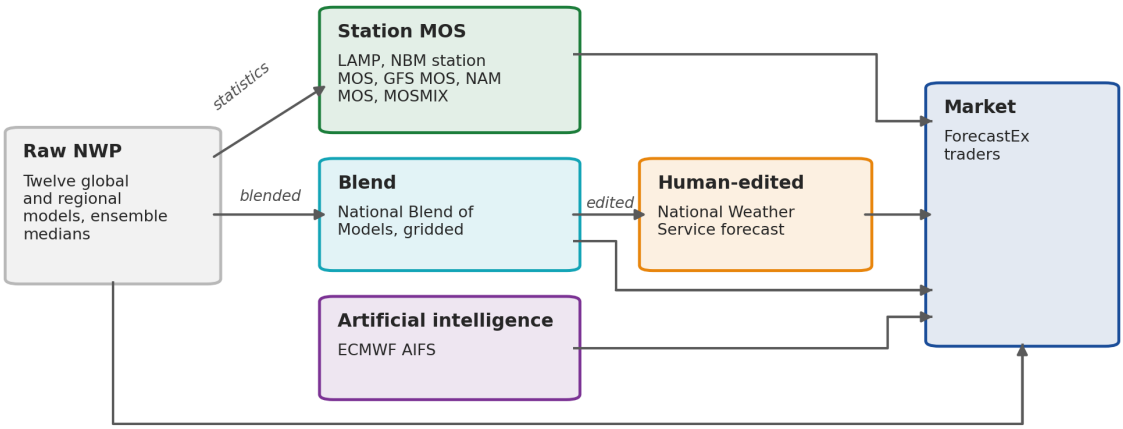

Fig. 2. An illustration of how the market sits downstream of the alternatives. Raw models feed the station MOS products and the National Blend of Models, forecasters edit the Blend into the National Weather Service forecast, and ForecastEx traders, hypothetically, have access to all of them.

human-edited forecast, taken from the National Digital Forecast Database on the same 2.5 km grid [37]. Forecasters edit that grid rather than issue a value for a particular weather station, so the series used here is the grid box containing the station. The last is the ECMWF artificial intelligence system AIFS, which is trained on decades of reanalysis and produces a forecast by learned pattern rather than by integrating the equations of motion at run time [49, 74].

Four alternative forecast systems investigated here, AIFS, the NOAA GEFS, the ECCC GEPS and the DWD ICON-EPS, publish ensemble members and thus probabilistic information. For these, I construct a full distribution (Section 2.6) and take that distribution's median as the system's central value.

The starting date for live capture varied by system (Table 1). For eight systems, I fill the record before live capture from an open-access service's archive of past model runs [69]. An archived run carries its initialization time but not the time it reached users, so crediting it at initialization would treat a forecast as available hours before public publication. Each reconstructed run is therefore assigned a publication time equal to the initialization time plus a system-specific delay (1.6 to 9.6 h).

2.4 Absolute and relative lead time

Figure 3 shows the two ways I measure lead. Absolute lead time is the number of hours remaining before location-local midnight at the end of the target day, so 30 h is 18 local time on the preceding day, 18 h is 06 local time on the target day, and 12 h is local noon. Relative lead time is the number of hours remaining before the maximum was observed. When multiple METAR reports represent the high, the last one is used as the time of maximum. I bin absolute lead time hourly from 36 h to 0 h and relative lead time from 30 h to 1 h.

A system's standing forecast at a lead is its most recent value at or before that time, so the comparison includes timeliness as well as accuracy. Figure 4 shows how often that value is refreshed. At 06 local time, the median time since a system's last update was 0.7 h for LAMP, 0.7 h for HRRR, 0.9 h for the National Weather Service forecast, and 7.7 h for NAM MOS, the slowest of the twenty.

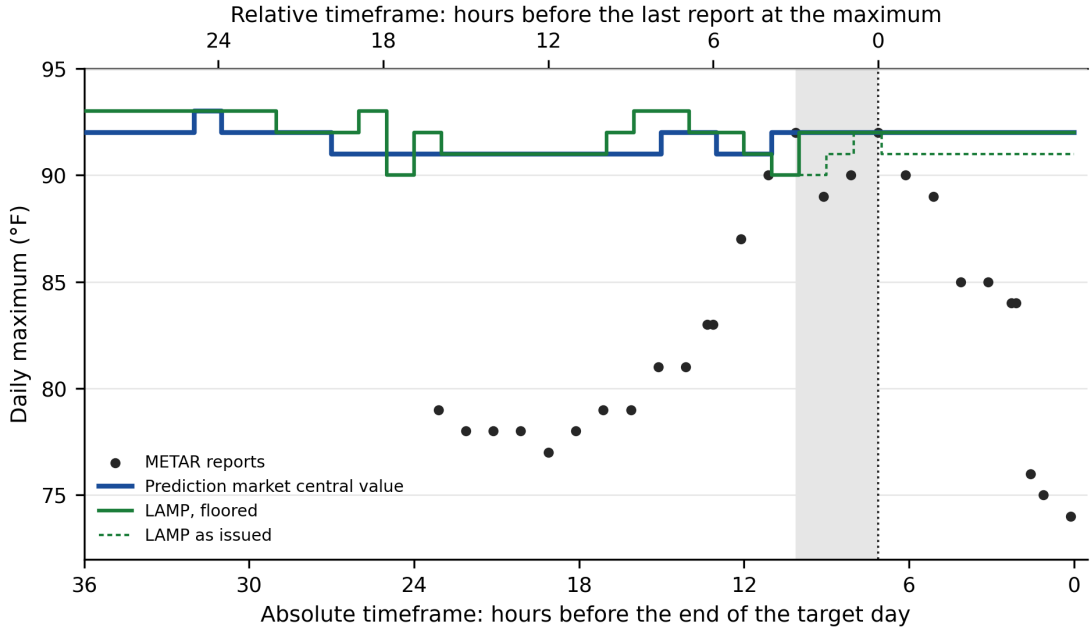

Fig. 3. The two timeframes and the floor, at Nashville on 26 July 2026. The maximum, 92°F, was first reached at 13:53 CDT and repeated again at 16:53 CDT (shaded). Absolute lead time (top axis) counts to the end of the target day and relative lead time (bottom axis) to the last report at the maximum. LAMP as issued falls to 90°F after the maximum, and flooring keeps it at the observed 92°F. The prediction market central value stays between 91°F and 93°F and settles at 92°F from 11 h.

On the target day, information about the daily high begins accumulating, and a forecast of 86°F after the temperature has reached 88°F has been falsified. The alternatives often contain stale, falsified forecasts which could be counted against them. However, in an effort to be generous to the alternatives, I floor each alternative's standing forecast at the running observed maximum, meaning they are not penalized for forecasts that have already been falsified by observations. I provide no such help to the prediction market forecast, which is left unaltered.

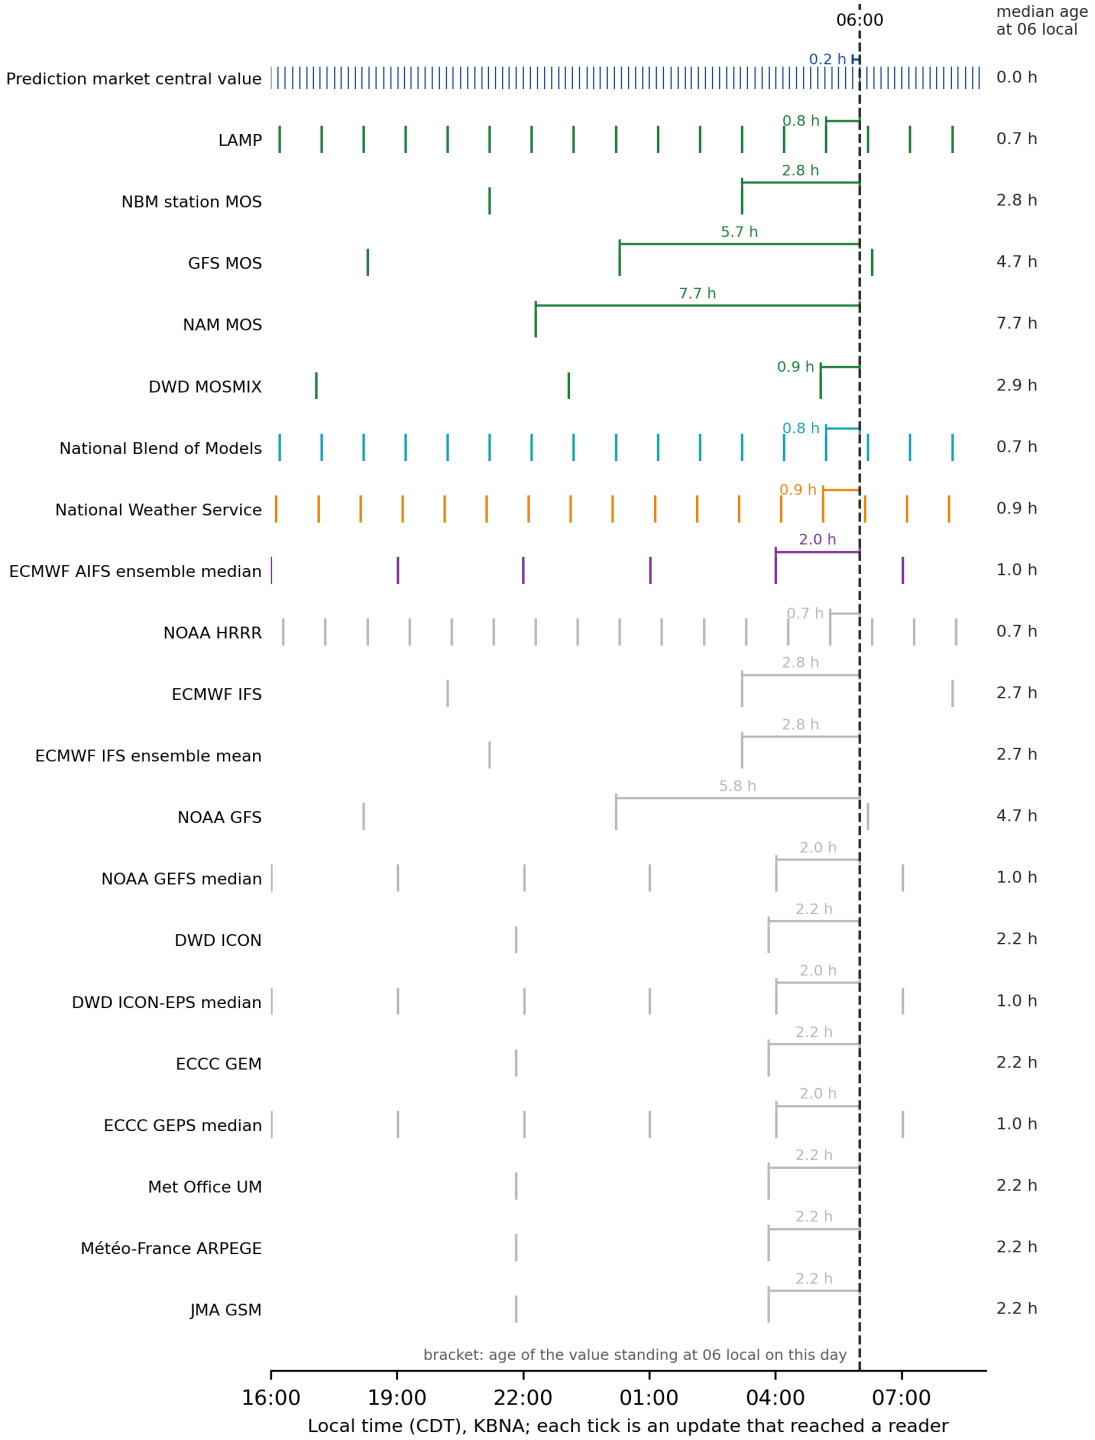

Fig. 4. Each system's standing forecast refresh cycle. Each tick is an update, and the line runs from the last update before 06 local time to 06 local, labeled with the age of the value standing at 06 local. The right column gives each system's median age at 06 local time over all location-days, measured over the period in which that system was captured live. The market is 'real time' but my capture interval was ten minutes.

2.5 Verification against various standards

The station MOS systems (Table 1) target the same weather stations as the prediction market. These systems produce forecasts at two temporal frequencies. One is hourly, matching the prediction market's temporal frequency. This is the case, for example, with LAMP, which provides the purest apples-to-apples comparison with the prediction market in this study.

All other comparisons between the prediction market and alternatives require more care.

Several of the station MOS systems also output a daily high temperature value which is not simply the highest hourly value of their own output, but instead most closely corresponds to the high defined by the National Weather Service daily climatological report, the maximum temperature at 1-minute resolution, with each minute representing a 5-minute rolling average, occurring between 07 and 19 local standard time [73, 91].

I therefore score those products twice. Their hourly values are scored against the METAR settlement value, which is the market's own target, and their issued daily maximum is scored against the climate report (Fig 5). Three products issue both: GFS MOS, NAM MOS and the NBM station bulletin. LAMP issues no daily maximum, and the MOSMIX daily maximum covers 06 to 18 UTC, which ends before the afternoon at US locations, so for those two the hourly values are used.

All other alternative systems produce discrete gridded output that represents areas larger than the point locations of the stations. The discrete nature of these grids limits their ability to forecast temperature at a particular weather station because they represent an average over a larger area. Nevertheless, using the single grid box that contains the station is the best representation that the system has for the station location. Noting that the system has a coarser spatial resolution than the station is a questionable defense because coarse spatial resolution is not a designed feature but a flaw that would be overcome if possible. These operational weather models are intended to forecast as granularly as possible, and their resolution is limited by computational constraints, not by design. Thus, it could be argued that these systems can be compared to station observations and the mismatch attributed to the system's lower resolution. However, I do apply basic bias correction to remove the systematic offset between model grids and station locations (described more below).

Also, the producers of the gridded systems for the National Blend of Models and the human-edited National Digital Forecast Database verify against the National Weather Service's Unrestricted Mesoscale Analysis (URMA). An analysis is itself a model product that assimilates the same observations and inherits the model's own terrain and land surface, so scoring a forecast against it measures agreement with a later state of the same modeling system rather than the agreement directly with a true observation. Nevertheless, for these two systems I show results scored against URMA analysis. In every comparison that follows, the market's error is measured against the METAR settlement value and each alternative's against its own standard (Eq 3).

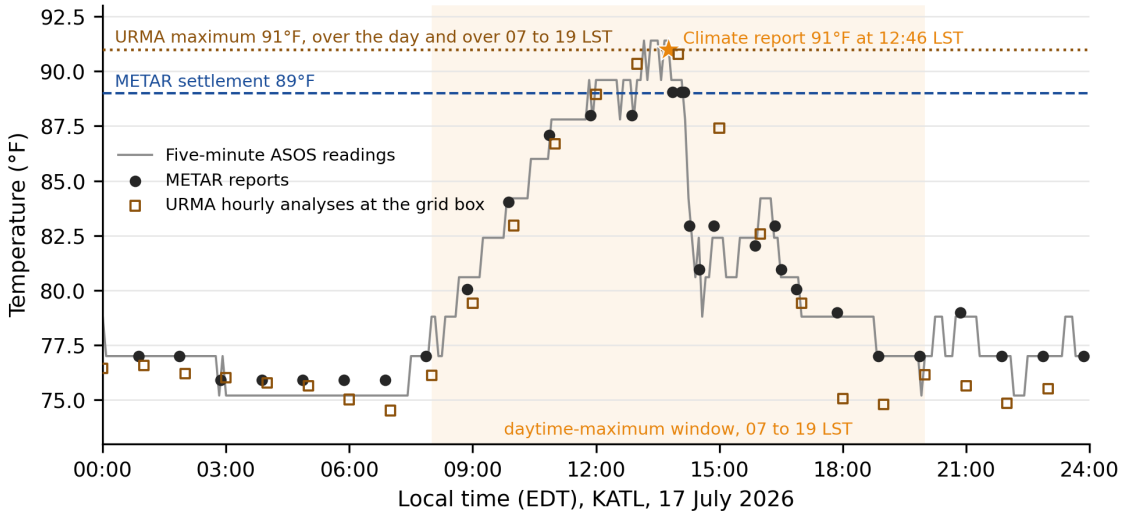

Fig. 5. The three verification targets (METAR, NWS CLI and URMA) at Atlanta on 17 July 2026, with five-minute ASOS readings (line). The METAR maximum was 89.06°F, so the prediction market contract settled at 89°F. The five-minute readings peaked at 91.4°F, and the climate report gave 91°F at 12:46 local standard time, inside the daytime-maximum window (shaded). The National Weather Service forecast at 06 local time was 91°F, matching its own target. Had it been scored against the prediction market standard, it would have counted as a 2°F error.

2.6 Scoring

2.6.1 Central values

I score deterministic values with the mean absolute error of system s at lead h,

$$\mathrm{MAE}_s(h) = \frac{1}{N_h} \sum_{i=1}^{N_h} \left| f_{s,i}(h) - o_i \right| \qquad (2)$$

over the N_h location-days with a value, where f_s is the system's value and o_s the observation that system is verified against, which is the METAR settlement value, the climate report or URMA as Section 2.5 sets out. The mean error, the signed average of $f_s - o_s$, measures level bias.

Every comparison therefore holds each system to its own standard. The market is always scored against the METAR settlement value its contracts resolve on, and each alternative against the readily available observation that most closely conforms to what that forecast system targets, so the market's percent improvement over system s on the location-days both cover is

$$\mathrm{PI}_s = 100 \times \left(1 - \frac{\mathrm{MAE}\,(\mathrm{FX\ vs\ METAR})}{\mathrm{MAE}\,(s\ \mathrm{vs\ its\ own\ standard})} \right) \qquad (3)$$

A positive value means the market is closer to the observation it is judged on than system s is to the observation it is judged on. Where a system is scored twice, each series carries its own denominator.

For the location maps (Section 3.2), each location-day's value is a system's mean over one window per timeframe, 30 to 25 h and 18 to 13 h before the end of the target day, and 18 to 13 h and 6 to 1 h before the maximum.

The gridded systems carry the bias correction promised in Section 2.5. Specifically, at each location and lead, I subtract the mean error over the trailing 30 target dates, requiring at least ten. This correction removes grid-to-point mismatch related to local elevation, land type, proximity to water, and similar fixed features, and because it uses only 30 days, it is seasonally aware, so a correction appropriate for winter is not smeared into summer. What it cannot remove is representativeness error, the part of the mismatch that varies from day to day with wind direction and cloud cover, since that is not a fixed offset.

2.6.2 Probability distributions

For the four systems that publish ensemble members I build a distribution for the daily maximum from each hour's ensemble mean and member spread, treating each hour as a normal distribution. The daily maximum exceeds a strike when at least one of the hours left in the day exceeds it, so the probability depends on how the hourly errors are related. Two extremes bound it. If the hourly errors were independent, many near misses would combine into a high chance that some hour exceeds the strike. If they are perfectly correlated, a warm or cold day moves every hour together, and the probability of exceeding the strike is the probability for the single hour least likely to exceed it. Hourly temperature errors within a day are strongly correlated, since they are set by the same air mass and the same model error, so I use the perfectly correlated case, which is the Fréchet bound for the maximum of dependent variables [92]. It uses every remaining hour rather than the single hour with the largest mean, and it is the more conservative of the two extremes for the market, since independence would make the ensembles look falsely sharp.

For these probabilistic ensembles, I score with the continuous ranked probability score (CRPS), which generalizes absolute error to a full distribution [93, 94],

$$\mathrm{CRPS}_s(h) = \frac{1}{N_h} \sum_{i=1}^{N_h} \sum_{k} \left[F_{s,i,h}(k) - \mathbf{1}(o_i \leq k) \right]^2 \qquad (4)$$

where $F_{s,i,h}(k)$ is the system's probability that location-day i settles at or below k. The CRPS is the integral over thresholds of the Brier score used for probabilistic binary forecast verification [94].

The four ensembles are gridded products, so I apply bias corrections before scoring them against station-level observations. Specifically, I calibrate these probabilistic systems with basic ensemble model output statistics (EMOS) [44]. For each ensemble, location and lead window, a Gaussian whose mean is a linear function of the ensemble mean and whose variance is a linear function of the ensemble variance is fitted by minimizing CRPS on the preceding 30 days. The fit is made on the forecast as the comparison scores it, held at the maximum already observed, so a training instant after the day's high counts only through the chance it gives the remaining hours of beating it. The market's ladder is used as quoted, with no calibration.

Additionally, I test calibration with reliability diagrams [95, 96]. I pool exceedance probabilities in 10% bins in one window per timeframe, 24 to 12 h before the end of the target day and 18 to 6 h

before the maximum, and plot the mean probability in each bin against the outcome frequency. I summarize calibration by the count-weighted mean absolute difference between forecast probability and occurrence frequency, and by the slope of a logistic regression of outcomes on the log-odds of the price, which equals one for calibrated prices [97].

Errors at different locations on the same date share large-scale weather, which travels with the synoptic waves that set daily temperature anomalies [98], and errors on consecutive dates are correlated for the same reason [99]. I therefore estimate 95% intervals from 1,000 bootstrap resamples of 7-day blocks of dates [100–102], resampling the market and each alternative under the same draws on the location-days both cover. I describe the market as better than an alternative at a lead when their paired difference is below zero at the 5% level after Holm adjustment [103] across the 25 series. A system counts as beaten only when every series of it is beaten, so the counts reported below are over the 20 systems.

2.7 Data completeness

Table 2 documents where data is incomplete or irregular. I make no exclusions for thin trading or location-days with a gap in the observation (METAR) record. Each system is scored on the location-days its own record covers, and a lead bin appears only if it includes at least 30 location-days. Because the exchange lists each ladder at a single clock time irrespective of the local time zone, only some location-days reach absolute leads above 30 h, which are shaded in the figures below.

Table 2. Data completeness. The data-quality factors that influence the results, along with their measured effects on the results. The records are grouped by date, and Table 1 shows the start dates for each system.

Item	Count	Effect on the results
Listed location-days	5,458	
No prediction market central value at any lead	45	no market to compare; left out
Recorded settlement value carries an ingest error (Buckley)	8	left out; listed with the exchange's value in settle_exclusions_t16.csv
Scored location-days	5,405	
Location-days with an hour missing from the observation record	270	kept; without them the market's MAE at 18 h is 1.29 against 1.29, and no verdict changes at any lead
Captures that hit the feed's 5,000-row cap, replaced by the full-day value	70,917	9,553 daily maxima changed, by +1.01 °F on average; the six affected series move by at most 0.002 °F at 18 h
Bias-correction warm-up, fewer than ten training dates (rows)	158,494	left out; scored as issued instead, the mean improvement at 18 h is 35.8% against 35.3%
Record start, 11 February 2026	16 series	the record's first day; the market is priced from the trade record to 2 July 2026 and from the order-book capture after. LAMP, NBM station MOS, issued max, NBM station MOS, hourly rows, GFS MOS, issued max, GFS MOS, hourly rows, NAM MOS, issued max, NAM MOS, hourly rows, ECMWF AIFS ens. median, ECMWF IFS, NOAA GFS, DWD ICON, ECCC GEM, Met Office UM, Météo-France ARPEGE, JMA GSM, prediction market
Record start, 14 July 2026	3 series	its capture began on 14 July 2026 and no archive reaches earlier. National Weather Service, climate report, National Weather Service, URMA, NOAA HRRR
Record start, 16 June 2026	3 series	the ensemble capture began on 16 June 2026. NOAA GEFS median, DWD ICON-EPS median, ECCC GEPS median
Record start, 18 July 2026	1 series	its capture began on 18 July 2026 and no archive reaches earlier. DWD MOSMIX
Record start, 27 June 2026	3 series	its capture began on 27 June 2026 and no archive reaches earlier. National Blend of Models, METAR, National Blend of Models, URMA, ECMWF IFS ens. mean
Priced instants whose ladder never crosses one half (no central value), at 30, 18 and 12 h	521, 392, 151	left out for every series, since a comparison needs the market's value; 13%, 7%, 3% of the priced instants

2.8 Robustness tests

Four tests repeat the comparisons under changes that could plausibly favor the market. The first restricts the sample to the location-days that every system covers at every absolute lead from 30 h to 0 h. The second scores every alternative as issued, without the station bias correction. The third calibrates the ensembles by a mean shift alone instead of EMOS, which separates the effect of correcting their center from the effect of correcting their spread. The fourth scores the market from the trade record and from the order-book capture over the window where both exist, which tests whether the change of price source affects results.

3 Results

3.1 Central-value forecasts

Figure 6 shows the mean absolute error, averaged across the 24 locations, of the central forecast value of all systems on both the absolute and relative timeframes.

On absolute lead time, the ForecastEx prediction market error was statistically significantly lower than all 20 systems at 33 of 37 hourly leads, including every lead from 36 h down to 5 h. The exceptions were between 1 h and 4 h before the end of the day, when the maximum had passed most of the time and errors clustered near zero (helped by the flooring I impose on the alternative systems).

On relative lead time, the prediction market had statistically significantly lower errors than all 20 alternatives at 29 of 30 leads, the exception being 30 h before the maximum, when the number of location-days holding a market value is small.

The prediction market's error was 1.49°F the evening before the target day, 1.29°F at 06 local time, and 1.06°F at noon.

Table 3 gives every series at the six leads. The five station MOS products were the strongest alternative group. LAMP was at 1.97°F, 1.63°F and 1.40°F, and the NBM station bulletin on its issued maximum at 1.76°F, 1.65°F and 1.55°F, so the leading station MOS product was the NBM station bulletin at 30 h and LAMP at 18 h and noon.

Below, I will often highlight results against the National Weather Service forecast, which represents expert human judgment with access to the same information as the prediction market participants, as well as the results against LAMP, the alternative whose resolution, design, and observational target is closest to the market's.

The market's error was 24% [20%, 29%], 21% [18%, 24%] and 24% [22%, 26%] below LAMP's (Eq 3), and 27% [23%, 31%], 32% [28%, 36%] and 38% [34%, 41%] below the National Weather Service forecast, at 30, 18 and 12 hours before the end of the target day, respectively.

On relative lead time the margins were similar. The market's error was 23% [20%, 26%], 21% [19%, 24%] and 23% [21%, 26%] below LAMP's, and 29% [25%, 33%], 31% [28%, 35%] and 35% [31%, 39%] below the National Weather Service forecast, at 18, 12 and 6 hours before the high was observed, respectively.

Figure 7 ranks every series at 30 h, the evening before the target day, when none of the day's temperatures had been observed. The market was at 1.49°F and had statistically significantly lower error than all 25 series after Holm adjustment. The closest were the NBM station bulletin on its issued maximum at 1.76°F, the AIFS ensemble median at 1.77°F, the National Weather Service forecast at 1.83°F and the Blend at 1.84°F. The furthest was NAM MOS read from its hourly values.

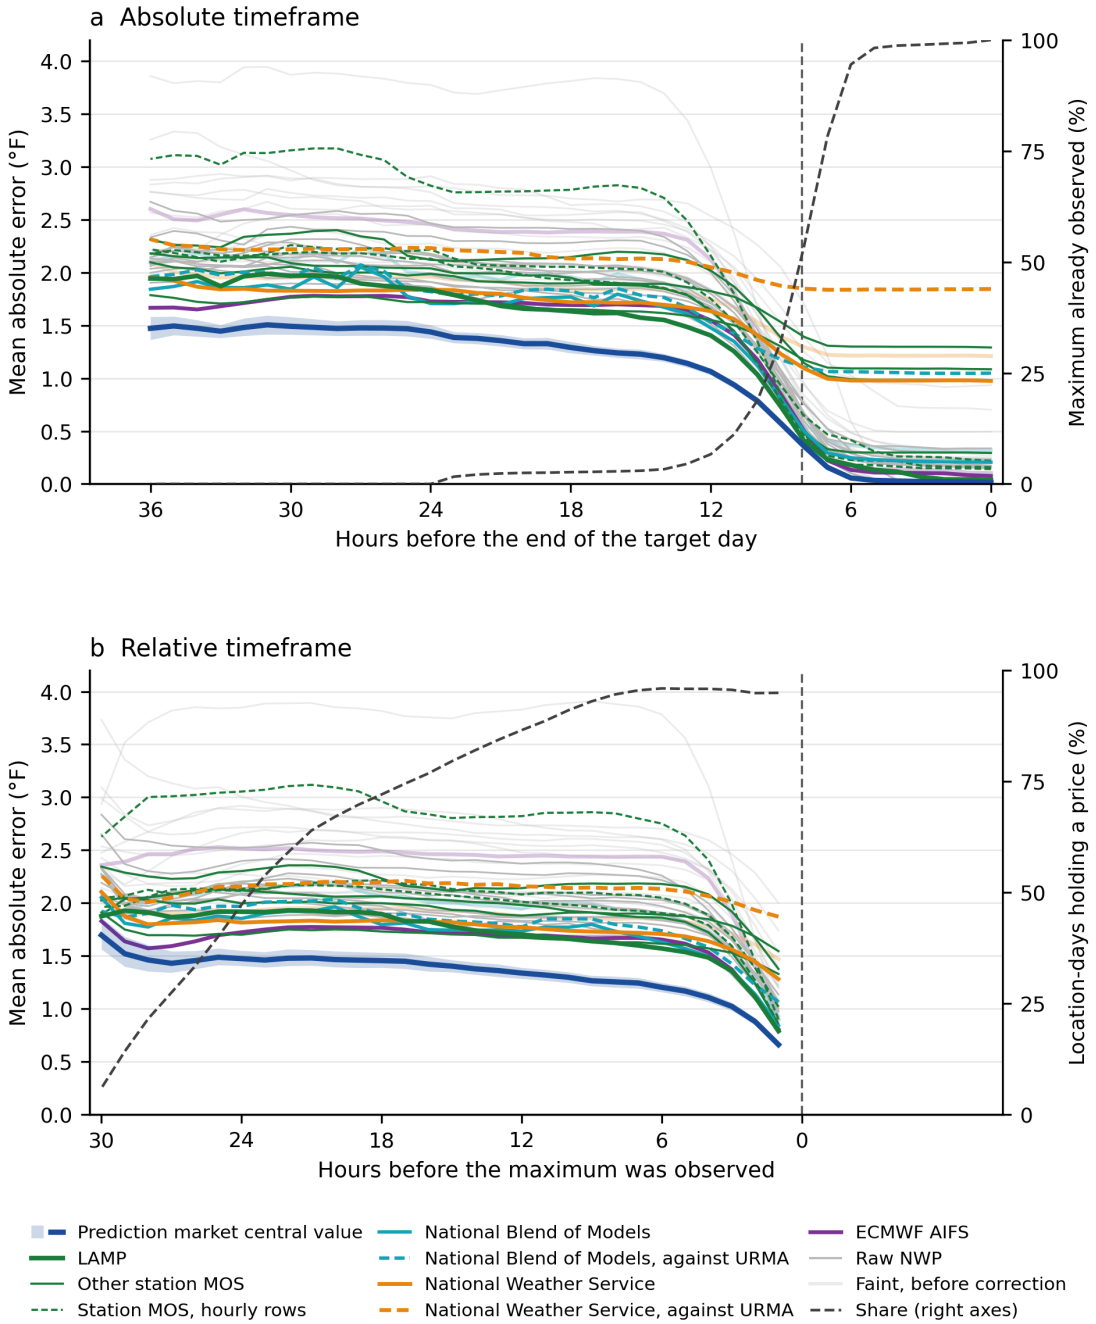

Fig. 6. The prediction market central value had the lowest error of the 20 alternatives in both timeframes, with each system scored against its own standard and each grid-box system corrected for its station offset. Faint lines are the same systems as issued without correction. Mean absolute error (°F) of the daily maximum against (a) absolute lead time and (b) relative lead time. The market is blue with its 95% interval shaded. Alternatives are colored by group, station MOS green with LAMP drawn heavier, the National Blend of Models teal, the National Weather Service forecast orange, AIFS purple, and the raw models grey. The five systems evaluated against two observational verification targets are drawn twice, with the second series dashed (the three station MOS products on their hourly values and the National Weather Service forecast and the Blend against URMA). The dashed black lines, on the right axes, give in (a) the share of location-days whose maximum had already been observed at that lead and in (b) the share where the market price was available at that lead. Alternatives are floored at the running observed maximum and the market probability is left unaltered.

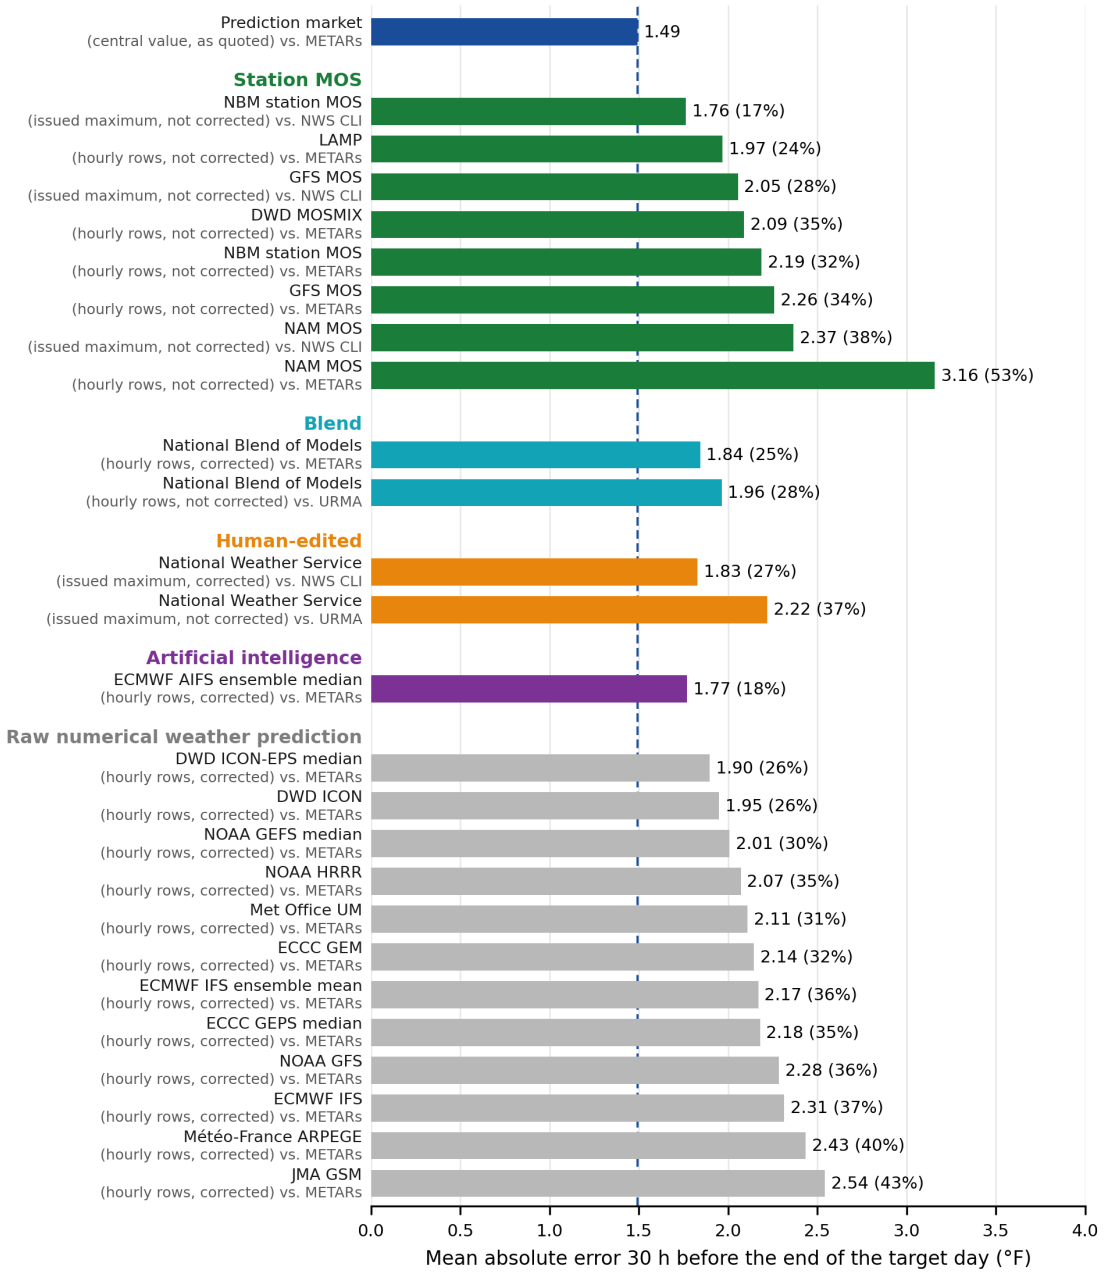

Each bar carries its mean absolute error in °F and, in brackets, the prediction market's improvement over it. Dashed, the market's error. Issued maximum, the daytime maximum the system issues; hourly rows, the largest of its hourly values for the day; corrected, its trailing station offset removed. The values are Table 3's.

Fig. 7. Every alternative was less accurate than the prediction market the evening before the target day. Mean absolute error (°F) at 30 h before the end of the target day for the market and the 25 series, with the market's percent improvement over each in brackets, grouped and coloured as in Table 3 and ordered by error within each group. The market's error is the dashed line, and improvement is Eq 3, the market against the METAR settlement value over the system against its own standard. Each label gives the system, the daily value it is scored on (the daytime maximum it issues, or the largest of its hourly values), whether its station offset is corrected, and the observation it is verified against.

Table 3. Central-value scores. Mean absolute error (°F) at 30, 18, and 12 h before the end of the target day and at 18, 12, and 6 h before the maximum was observed, with the market's percent improvement over each series in parentheses (Eq 3, the market against the METAR settlement value over the series against its own standard), computed on the location-days both cover. ME is the mean error at 18 h before the end of the day. Fifteen series carry a trailing station bias correction and the rest are scored as issued (Section 2.5), as the last column marks. Five systems contribute two series each, which gives 25 series across the 20 systems. Each group's row gives the mean over its systems, and groups and the series within them are ordered by error at 18 h before the end of the day.

	Hours before the end of the day			Hours before the maximum				
System	30	18	12	18	12	6	ME	Offset
Prediction market	1.49	1.29	1.06	1.45	1.34	1.20	+0.13	
Artificial intelligence	1.77	1.69	1.55	1.76	1.69	1.65	0.00	
ECMWF AIFS ens. median	1.77 (18%)	1.69 (24%)	1.55 (31%)	1.76 (19%)	1.69 (22%)	1.65 (28%)	0.00	corrected
Human-edited	1.83	1.72	1.63	1.84	1.77	1.71	−0.03	
National Weather Service, climate report	1.83 (27%)	1.72 (32%)	1.63 (38%)	1.84 (29%)	1.77 (31%)	1.71 (35%)	−0.03	corrected
National Weather Service, URMA	2.22 (37%)	2.13 (44%)	2.05 (49%)	2.20 (39%)	2.16 (42%)	2.13 (46%)	+1.05	as issued
Blend	1.84	1.77	1.48	1.80	1.73	1.63	+0.08	
National Blend of Models, METAR	1.84 (25%)	1.77 (31%)	1.48 (30%)	1.80 (25%)	1.73 (27%)	1.63 (29%)	+0.08	corrected
National Blend of Models, URMA	1.96 (28%)	1.83 (33%)	1.53 (32%)	1.89 (28%)	1.77 (28%)	1.74 (34%)	−0.76	as issued
Station MOS	2.05	1.86	1.69	1.98	1.87	1.82	−0.46	
LAMP	1.97 (24%)	1.63 (21%)	1.40 (24%)	1.90 (23%)	1.70 (21%)	1.57 (23%)	−0.45	as issued
NBM station MOS, issued max	1.76 (17%)	1.65 (22%)	1.55 (31%)	1.73 (17%)	1.67 (21%)	1.61 (27%)	+0.07	as issued
GFS MOS, issued max	2.05 (28%)	1.90 (32%)	1.82 (41%)	2.00 (28%)	1.88 (30%)	1.89 (37%)	+0.25	as issued
NBM station MOS, hourly rows	2.19 (32%)	1.93 (33%)	1.65 (35%)	2.13 (32%)	2.01 (34%)	1.90 (37%)	−1.06	as issued
DWD MOSMIX	2.09 (35%)	1.95 (40%)	1.64 (39%)	2.06 (37%)	1.98 (38%)	1.86 (40%)	−1.16	as issued
GFS MOS, hourly rows	2.26 (34%)	2.06 (37%)	1.75 (39%)	2.22 (34%)	2.08 (36%)	2.05 (41%)	−0.88	as issued
NAM MOS, issued max	2.37 (38%)	2.15 (40%)	2.02 (47%)	2.23 (35%)	2.14 (38%)	2.18 (46%)	−1.03	as issued
NAM MOS, hourly rows	3.16 (53%)	2.78 (54%)	2.16 (51%)	2.96 (51%)	2.82 (53%)	2.75 (56%)	−2.24	as issued
Raw numerical weather prediction	2.17	1.98	1.74	2.10	1.99	1.93	0.00	
DWD ICON-EPS median	1.90 (26%)	1.69 (28%)	1.51 (31%)	1.86 (27%)	1.70 (25%)	1.66 (31%)	−0.03	corrected
NOAA HRRR	2.07 (35%)	1.79 (34%)	1.48 (31%)	1.97 (34%)	1.82 (32%)	1.66 (32%)	−0.24	corrected
NOAA GEFS median	2.01 (30%)	1.81 (33%)	1.64 (37%)	1.94 (30%)	1.83 (31%)	1.79 (36%)	0.00	corrected
DWD ICON	1.95 (26%)	1.84 (31%)	1.67 (36%)	1.90 (25%)	1.86 (29%)	1.84 (35%)	+0.04	corrected
Met Office UM	2.11 (31%)	1.87 (32%)	1.68 (37%)	1.99 (29%)	1.87 (30%)	1.83 (35%)	+0.02	corrected
ECCC GEPS median	2.18 (35%)	1.93 (37%)	1.72 (39%)	2.17 (37%)	2.02 (37%)	1.90 (39%)	−0.02	corrected
ECCC GEM	2.14 (32%)	2.00 (36%)	1.80 (41%)	2.06 (31%)	2.01 (34%)	1.97 (39%)	+0.05	corrected
ECMWF IFS ens. mean	2.17 (36%)	2.01 (40%)	1.71 (39%)	2.10 (36%)	1.99 (36%)	1.91 (39%)	−0.07	corrected
ECMWF IFS	2.31 (37%)	2.08 (39%)	1.83 (42%)	2.21 (35%)	2.09 (37%)	2.03 (41%)	+0.08	corrected
NOAA GFS	2.28 (36%)	2.08 (38%)	1.83 (42%)	2.15 (34%)	2.09 (37%)	2.04 (42%)	+0.13	corrected
Météo-France ARPEGE	2.43 (40%)	2.28 (44%)	2.03 (48%)	2.33 (39%)	2.27 (42%)	2.22 (46%)	+0.03	corrected
JMA GSM	2.54 (43%)	2.41 (47%)	2.00 (47%)	2.50 (43%)	2.37 (44%)	2.26 (47%)	+0.03	corrected

Scoring the three products that issue a daily maximum on both of their outputs shows the effect of the verification target on their scored errors. Each of them was weaker on its hourly values than on its issued maximum at almost every lead, which is what the issued maximum exists for. Counting those hourly series, the station MOS group falls to last at 18 h, artificial intelligence 1.69, blend 1.80, human-edited 1.93, raw numerical weather prediction 1.98, station MOS 2.01, while counting one series per system it is artificial intelligence 1.69, human-edited 1.72, blend 1.77, station MOS 1.86, raw numerical weather prediction 1.98.

Removing each grid-box system's station offset changed which group was closest to the market. As issued, the groups at 18 h ran human-edited 1.84, station MOS 1.86, blend 1.96, artificial intelligence 2.39, raw numerical weather prediction 2.49. After correction they ran artificial intelligence 1.69, human-edited 1.72, blend 1.77, station MOS 1.86, raw numerical weather prediction 1.98. Systems enter on their own records, which start in different months, so the groups are better compared on the 1,042 location-days where every series holds every lead. There the ordering at 18 h was artificial intelligence 1.44, human-edited 1.63, blend 1.63, station MOS 1.71, raw numerical weather prediction 1.78. Once the station offset is removed, a raw artificial intelligence ensemble is about as accurate at a location as MOS products fitted to that location.

3.2 Geography

Figure 8 shows the market against the National Weather Service forecast, which represents expert human judgment with access to the same information as market participants, and against LAMP, the alternative whose target is closest to the market's, at each of the 24 locations and at two leads in both timeframes. The prediction market advantage held across the geographic diversity of locations in both timeframes.

On the absolute timeframe, the median improvement over the National Weather Service forecast was 25% [19%, 32%] at 30 h and 29% [25%, 37%] at 18 h, and over LAMP 22% [17%, 29%] and 17% [16%, 24%]. On the relative timeframe it was 29% [21%, 34%] and 28% [25%, 37%] at 18 and 12 h before the maximum, and over LAMP 19% [16%, 25%] and 20% [17%, 24%].

Pooling every series at every location produces essentially the same result (Fig 8i, j). Of 1,186 comparisons across the two absolute leads the market was ahead in 1,178, and of 1,186 across the two relative leads in 1,182, with median improvements of 32% and 31%.

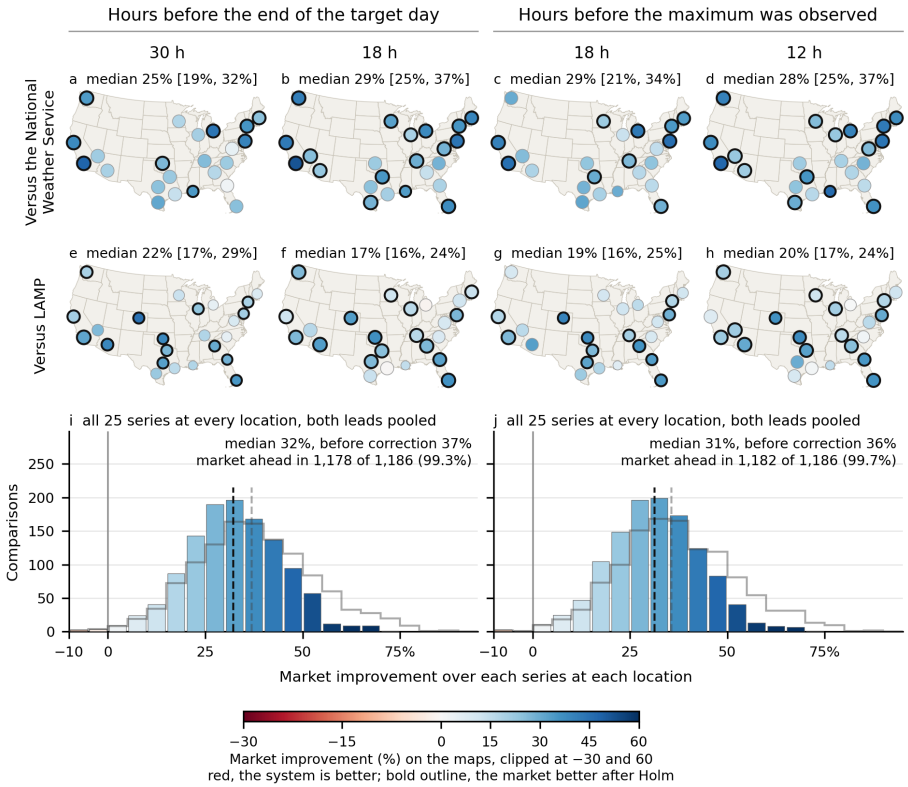

Fig. 8. The market had lower error than the National Weather Service forecast and LAMP at nearly every location, and lead, in both timeframes. Displayed is the percent improvement of Eq 3, the market scored against the METAR settlement value over each system scored against its own standard, for the National Weather Service forecast (a to d) and LAMP (e to h), at 30 h and 18 h before the end of the target day and at 18 h and 12 h before the maximum. Each circle is a location, colored by the improvement and outlined in bold where the market is better after Holm adjustment, with the median and its interval above each panel. Panels i and j pool every alternative at every location across the two leads of that timeframe, with the distribution before the bias correction drawn as a grey outline. The National Weather Service panels have 23 locations, because Buckley Space Force Base has no climate report.

3.3 Full probability distribution

Figure 9 shows the market's full distribution against the four ensembles after EMOS calibration. The market had statistically significantly lower errors than all four at 29 of 37 absolute leads and 28 of 30 relative leads, the exceptions falling at the ends of the two timeframes, from 0 h to 5 h and at 7 h and 8 h before the end of the target day, when all systems approach zero, and at 29 h and 30 h before the maximum, where few location-days hold a price. Table 4 gives its CRPS at the six leads. It was 1.06°F at 30 h, 0.91°F at 18 h, and 0.78°F at noon. The best calibrated ensemble at 18 h was the ECMWF AIFS ensemble at 1.09°F, so the market's error was 16% [13%, 20%] lower.

Market prices lay closer to the diagonal of the reliability diagrams than any ensemble's probabilities even after EMOS calibration (Fig 9c, d). The count-weighted mean absolute difference between price and payout frequency was 0.027 in the absolute window and 0.027 in the relative one, so a price typically differed from the observed frequency by about 3 percentage points. The departures that remained were systematic. The logistic slope was 1.23 [1.19, 1.27], so outcomes were more extreme than prices implied and prices leaned slightly toward 50 cents, which is the direction of

Table 4. Scores for the full distribution. Continuous ranked probability score (°F) at the leads of Table 3, with the market's percent improvement over each ensemble in parentheses, on the contracts and location-days all five systems share, each ensemble scored against the METAR settlement value as the market is. The ensembles are calibrated, for this study, by ensemble model output statistics on the preceding 30 days (Section 2.6), are ordered by score at 18 h before the end of the target day, and the market's ladder is used as quoted.

	Hours before the end of the day			Hours before the maximum		
System	30	18	12	18	12	6
Prediction market	1.06	0.91	0.78	1.02	0.94	0.86
ECMWF AIFS ensemble	1.19 (11%)	1.09 (16%)	1.10 (30%)	1.18 (13%)	1.12 (15%)	1.09 (21%)
DWD ICON-EPS	1.32 (20%)	1.19 (23%)	1.15 (33%)	1.29 (21%)	1.21 (22%)	1.18 (27%)
NOAA GEFS	1.39 (23%)	1.29 (29%)	1.27 (39%)	1.33 (23%)	1.30 (28%)	1.29 (34%)
ECCC GEPS	1.51 (30%)	1.40 (35%)	1.37 (43%)	1.50 (32%)	1.43 (34%)	1.40 (39%)

the favorite–longshot pattern [27–29]. The calibrated ensembles departed from the diagonal by 0.027 to 0.034, with slopes below one, so they remained overconfident.

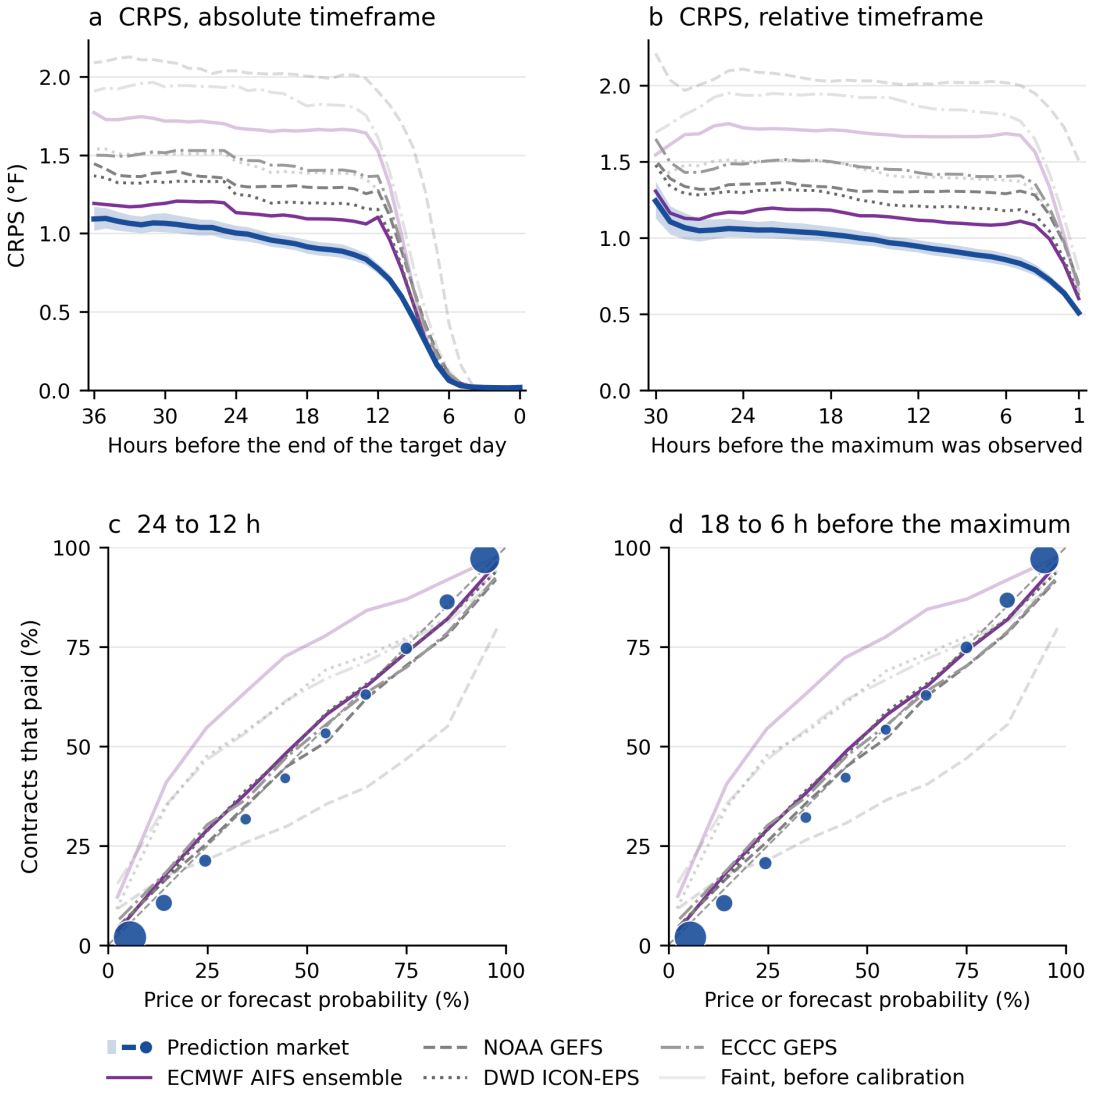

Fig. 9. The market's exceedance probability ladder had lower absolute error than the four ensembles and lay closer to the diagonal, in both timeframes, even after the ensembles were calibrated. Faint lines are the same ensembles as issued prior to calibration. (a, b) CRPS (°F) against absolute and relative lead time. The market is blue with its 95% interval shaded, AIFS purple, and GEFS, GEPS and ICON-EPS grey with different line styles. A CRPS of 1°F is comparable to a central-value forecast that misses by 1°F. (c, d) Reliability diagrams for contracts priced 24 to 12 h before the end of the target day and 18 to 6 h before the maximum. Market circles have area proportional to the number of contracts, and points above the diagonal indicate contracts that paid more often than their price implied. Contracts on the same location-day share one outcome, so the block-bootstrap intervals are about three times as wide as binomial intervals [104] and per-bin intervals are omitted. One legend serves all four panels.

3.4 Latency and accuracy

A forecast can be better because it is strictly more accurate or because it is more recent (even sampled at discrete times) and thus more representative of fresher information. The prediction market incorporates information dynamically in real time, and I sample it here at 10-minute intervals while the alternatives update on hourly to multi-hourly scales (Table 1, Figure 4). Figure 10 strives to separate pure accuracy from latency by comparing the market's improvement over the alternative forecasts at fixed time with the market's improvement over the alternative forecasts calculated only at times when the alternative forecasts were just refreshed. Pooled over the twenty systems, the market improvement was 32% [30%, 34%] at fixed leads and 32% [30%, 34%] with every alternative system just refreshed, meaning staleness accounted for none of the observed market advantage on average.

The prediction market's advantage over some of the best alternative systems was slightly narrowed at the refresh times, for example, by 0.8%, 2.4%, and 4.5% for LAMP, the corrected ECMWF AIFS, and the NBM station MOS (issued maximum against the NWS CLI report), respectively. This indicates that the best alternative systems also incorporate new information well and are most accurate when they are freshest (though they remain substantially less accurate than the market, Fig. 10). Part of the disadvantage of the alternative systems is the hours of upstream processing required before their publication. So even when they are just released, their forecasts already embed plenty of outdated information. That is combined with the idea that the market theoretically has all this information in hand and needs only to respond to additional pertinent information (Fig. 2).

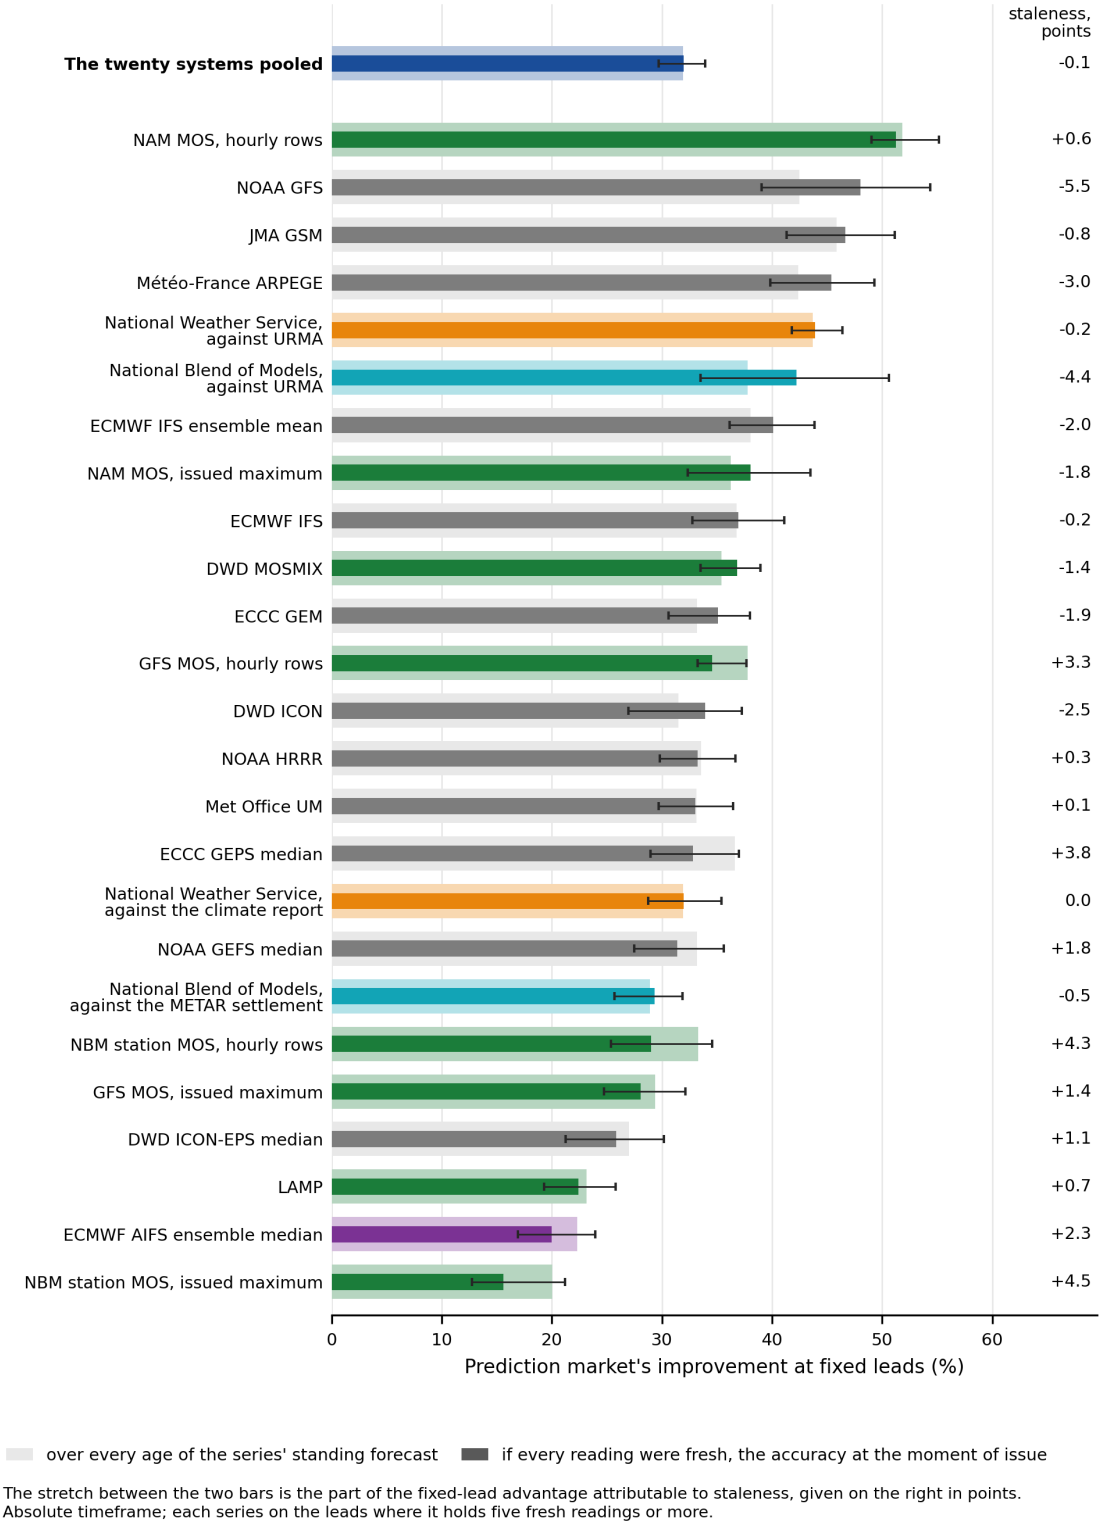

Fig. 10. The market's advantage over alternatives was primarily attributable to accuracy at the moment of issue, not necessarily a faster response to new information. For each series, the market's percent improvement in mean absolute error over that series at fixed leads and over every age of its standing forecast (pale), and the improvement if every reading were fresh (dark, with its 95% interval). The stretch between the two is the part of the advantage that the age of the standing forecast accounts for, printed on the right in points. The twenty systems pooled are at the top.

3.5 Robustness

No conclusion changed under the tests of Section 2.8 (Table 5). The market remained ahead of every series at the six leads on the balanced sample, with every alternative scored as issued and without the station bias correction, and with the ensembles calibrated by a mean shift instead of EMOS. Scoring the market from the trade record instead of the order-book capture, over the window where both exist, moved its error by less than the sampling noise at every lead.

Table 5. Robustness tests. Each test's modifications and whether the market still had the lowest error at 30, 18, and 12 hours before the target day's end, as well as at 18, 12, and 6 hours before the maximum. Series counts are provided across all 25.

			Hours before the end of the day			Hours before the maximum		
Test	Location-days	What is counted	30	18	12	18	12	6
Balanced sample	1042	series the market beats after Holm	25/25	25/25	25/25	25/25	25/25	25/25
All alternatives as issued, no bias correction		series the market beats after Holm	25/25	25/25	25/25	25/25	25/25	25/25
Ensembles with a mean shift instead of EMOS		ensembles the market beats on CRPS after Holm	4/4	4/4	4/4	4/4	4/4	4/4
Price-source overlap: the trade record against the order-book capture	2131	market MAE, trade record less capture (°F)	−0.001 [−0.020, +0.015]	+0.004 [−0.015, +0.022]	+0.056 [+0.002, +0.132]	−0.001 [−0.027, +0.027]	−0.004 [−0.028, +0.020]	+0.025 [−0.002, +0.060]

4 Summary and conclusions

While prediction markets have consistently demonstrated value in relatively subjective, less inherently quantitative realms of human prediction, it is far less clear that they should add value for predicting physical quantities compared with established mechanistic models with built-in statistical correction and/or human forecasters with access to the same tools.

Daily weather forecasting provides the grounds for such a test, as it is one of the most mature physical fields focused on prediction; conventional forecasting systems sit at the frontier of applied physics, statistics, and computer science, and highly trained human experts further supervise their output.

The analysis here showed that the ForecastEx prediction market central value for daily maximum temperature had statistically significantly lower error than each of the 20 alternatives investigated at 33 of 37 hourly leads counted to the end of the target day and 29 of 30 counted to the maximum, the exceptions falling in the hours after the maximum had already been observed on the absolute timeframe, and at 30 h before the maximum, where few location-days hold a price. Pooling every series at every location, the market was ahead in 1,178 of 1,186 (99.3%) comparisons. Its full price distribution also had lower error than the four alternative systems that produced probabilistic information in both the absolute and relative timeframes.

At 6:00 PM local time on the evening before the target day, the prediction market's mean absolute error was 1.49°F, compared with 1.76°F for the NBM station bulletin, the closest deterministic alternative at that lead (17% [14%, 19%] lower), and 1.97°F for LAMP, the closest product scored on hourly values (24% [20%, 29%] lower), and its CRPS was 1.06°F compared with 1.19°F for ECMWF AIFS, the best EMOS-corrected probabilistic alternative (11% [8%, 15%] lower). The margins in favor of the market generally held or got larger through the target day.

The introduction posed two bars for prediction market forecasts to potentially clear. The first bar was performing better than raw numerical weather prediction models, and the second was performing better than station MOS products designed to correct those models and the human-expert forecasters with access to the same tools (judged on their own verification framework). The ForecastEx prediction market cleared both bars. This result is corroborated by parallel work on a different exchange with a different settlement framework [66].

Station MOS is a sophisticated, deliberately designed product [32, 33], and the National Weather Service forecast is produced by trained meteorologists who see the same guidance and make use of additional information and experience. No participant in the prediction market receives any particular guidance, nor are they screened to ensure they are properly credentialed or have sufficient expertise.

What the market has instead is the mechanism described in Section 1.1. Galton's crowd and the market price share the property that no participant needs to be right for the aggregate to be, provided judgments are diverse and something combines them [4, 5]. Prices do the combining continuously, and the payoff attached to being correct rewards a participant for acting on what they privately know rather than on what is already published [6]. The market seems to accomplish better statistical correction than model output statistics do through fitted regressions, but as an emergent property without design.

That a price formed this way is more accurate than the operational systems it sits downstream of, over eight months, 24 locations and 20 alternatives, extends the documented successes of prediction markets from subjective realms into a mature field of physical prediction.

Data and code availability

ForecastEx publishes its daily price and trade records at https://forecastex.com/data. All alternative systems compared here are publicly available from their originating services. The station MOS products and the National Blend of Models are published by the NOAA Meteorological Development Laboratory at https://vlab.noaa.gov/web/mdl, with LAMP at https://vlab.noaa.gov/web/mdl/lamp, GFS MOS and NAM MOS at https://vlab.noaa.gov/web/mdl/mos and the Blend at https://vlab.noaa.gov/web/mdl/nbm. The National Weather Service forecast is published from the National Digital Forecast Database at https://digital.weather.gov. The GFS, GEFS, NAM and HRRR are published by the National Centers for Environmental Prediction at https://www.nco.ncep.noaa.gov/pmb/products and https://rapidrefresh.noaa.gov/hrrr. The ECMWF IFS, its ensemble and AIFS are published at https://www.ecmwf.int/en/forecasts/datasets/open-data and https://data.ecmwf.int/forecasts. ICON, ICON-EPS and MOSMIX are published by the Deutscher Wetterdienst at https://opendata.dwd.de. GEM and the GEPS are published by Environment and Climate

Change Canada at https://eccc-msc.github.io/open-data. The Unified Model is published by the Met Office at https://www.metoffice.gov.uk/services/data, ARPEGE by Météo-France at https://donneespubliques.meteofrance.fr, and the GSM by the Japan Meteorological Agency at https://www.jma.go.jp/jma/jma-eng/jma-center/nwp. The METAR observations, the daily climate reports and URMA are published by NOAA, and the MOS bulletins as issued are archived by the Iowa Environmental Mesonet [105] at https://mesonet.agron.iastate.edu, which is the source used here, each bulletin timed to its cycle plus the measured delay in its arrival.

A public archive will hold the data behind every figure of this manuscript and sample scripts that demonstrate processing. The verification statistics are also updated daily at https://weather.weatherclimatehumansystems.org/accuracy.html [106], which runs on the current record and will therefore differ from the frozen build reported here.

Competing interests

I am employed by Trend Analytics, LLC, a subsidiary of Interactive Brokers Group, which owns and operates ForecastEx, the exchange whose market prices are evaluated in this study. As with much academic research, certain research questions and their downstream results and framings are more appealing to the researcher than others. This is the case here, and I absolutely concede that the findings were welcome to me and my organization. Nevertheless, to the best of my knowledge, the findings are presented accurately and forthrightly. All the underlying datasets are publicly available, and I invite others to investigate the same questions from the same or different angles.

References

1. Wolfers J, Zitzewitz E. Prediction markets. J Econ Perspect. 2004;18(2) 107–126. https://doi.org/10.1257/0895330041371321
2. Wolfers J, Zitzewitz E. Interpreting prediction market prices as probabilities. NBER Working Paper 12200. 2006. https://doi.org/10.3386/w12200
3. Manski CF. Interpreting the predictions of prediction markets. Econ Lett. 2006;91(3) 425–429. https://doi.org/10.1016/j.econlet.2006.01.004
4. Galton F. Vox populi. Nature. 1907;75(1949) 450–451. https://doi.org/10.1038/075450a0
5. Surowiecki J. The wisdom of crowds. New York. Doubleday; 2004.
6. Sunstein CR. Infotopia. How many minds produce knowledge. New York. Oxford University Press; 2006.
7. Sabin P. The bet. Paul Ehrlich, Julian Simon, and our gamble over Earth's future. New Haven. Yale University Press; 2013.
8. Brown PT. When are scientific claims and expert assertions untrustworthy? And how prediction markets can help. IBKR Campus, Traders' Insight; 10 November 2025. Available from https://www.interactivebrokers.com/campus/traders-insight/ibkr-climate-energy/when-are-scientific-claims-and-expert-assertions-untrustworthy-and-how-prediction-markets-can-help/

9. Brown PT. The value of climate prediction markets. IBKR Campus, Traders' Insight; 28 March 2025. Available from https://www.interactivebrokers.com/campus/traders-insight/ibkr-climate-energy/the-value-of-climate-prediction-markets/

10. Hayek FA. The use of knowledge in society. Am Econ Rev. 1945;35(4) 519–530.

11. Grossman SJ, Stiglitz JE. On the impossibility of informationally efficient markets. Am Econ Rev. 1980;70(3) 393–408.

12. Gneiting T, Raftery AE. Strictly proper scoring rules, prediction, and estimation. J Am Stat Assoc. 2007;102(477) 359–378. https://doi.org/10.1198/016214506000001437

13. Arrow KJ, Forsythe R, Gorham M, Hahn R, Hanson R, Ledyard JO, et al. The promise of prediction markets. Science. 2008;320(5878) 877–878. https://doi.org/10.1126/science.1157679

14. Kagan N, Baiocchi R. Calibration in prediction markets. Theory and evidence. Kalshi Research working paper; August 2026. Available from https://kalshi.com/research/publications/calibration

15. Berg JE, Nelson FD, Rietz TA. Prediction market accuracy in the long run. Int J Forecast. 2008;24(2) 285–300. https://doi.org/10.1016/j.ijforecast.2008.03.007

16. Forsythe R, Nelson F, Neumann GR, Wright J. Anatomy of an experimental political stock market. Am Econ Rev. 1992;82(5) 1142–1161.

17. Rhode PW, Strumpf KS. Historical presidential betting markets. J Econ Perspect. 2004;18(2) 127–142. https://doi.org/10.1257/0895330041371253

18. Snowberg E, Wolfers J, Zitzewitz E. Prediction markets for economic forecasting. NBER Working Paper 18222. 2012. https://doi.org/10.3386/w18222

19. Erikson RS, Wlezien C. Are political markets really superior to polls as election predictors? Public Opin Q. 2008;72(2) 190–215. https://doi.org/10.1093/poq/nfn010

20. Rothschild D. Forecasting elections. Comparing prediction markets, polls, and their biases. Public Opin Q. 2009;73(5) 895–916. https://doi.org/10.1093/poq/nfp082

21. Goel S, Reeves DM, Watts DJ, Pennock DM. Prediction without markets. In: Proceedings of the 11th ACM Conference on Electronic Commerce. New York. Association for Computing Machinery; 2010. https://doi.org/10.1145/1807342.1807400

22. Atanasov P, Rescober P, Stone E, Swift SA, Servan-Schreiber E, Tetlock P, et al. Distilling the wisdom of crowds. Prediction markets vs. prediction polls. Manage Sci. 2017;63(3) 691–706. https://doi.org/10.1287/mnsc.2015.2374

23. Mellers B, Ungar L, Baron J, Ramos J, Gurcay B, Fincher K, et al. Psychological strategies for winning a geopolitical forecasting tournament. Psychol Sci. 2014;25(5) 1106–1115. https://doi.org/10.1177/0956797614524255

24. Cowgill B, Zitzewitz E. Corporate prediction markets. Evidence from Google, Ford, and Firm X. Rev Econ Stud. 2015;82(4) 1309–1341. https://doi.org/10.1093/restud/rdv014

25. Chen KY, Plott CR. Information aggregation mechanisms. Concept, design and implementation for a sales forecasting problem. Pasadena. California Institute of Technology Social Science Working Paper 1131; 2002.

26. Dreber A, Pfeiffer T, Almenberg J, Isaksson S, Wilson B, Chen Y, et al. Using prediction markets to estimate the reproducibility of scientific research. Proc Natl Acad Sci U S A. 2015;112(50) 15343–15347. https://doi.org/10.1073/pnas.1516179112

27. Page L, Clemen RT. Do prediction markets produce well-calibrated probability forecasts? Econ J. 2013;123(568) 491–513. https://doi.org/10.1111/j.1468-0297.2012.02561.x

28. Snowberg E, Wolfers J. Explaining the favorite–long shot bias. Is it risk-love or misperceptions? J Polit Econ. 2010;118(4) 723–746. https://doi.org/10.1086/655844

29. Bürgi C, Deng W, Whelan K. Makers and takers. The economics of the Kalshi prediction market. University College Dublin working paper; 2026.

30. Diercks AM, Katz JD, Wright JH. Kalshi and the rise of macro markets. Working paper; 2026.

31. Bauer P, Thorpe A, Brunet G. The quiet revolution of numerical weather prediction. Nature. 2015;525(7567) 47–55. https://doi.org/10.1038/nature14956

32. Glahn HR, Lowry DA. The use of model output statistics (MOS) in objective weather forecasting. J Appl Meteorol. 1972;11(8) 1203–1211. https://doi.org/10.1175/1520-0450(1972)011%3C1203:TUOMOS2.0.CO;2>

33. Ghirardelli JE, Glahn B. The Meteorological Development Laboratory's aviation weather prediction system. Weather Forecast. 2010;25(4) 1027–1051. https://doi.org/10.1175/2010WAF2222312.1

34. Vislocky RL, Fritsch JM. Improved model output statistics forecasts through model consensus. Bull Am Meteorol Soc. 1995;76(7) 1157–1164. https://doi.org/10.1175/1520-0477(1995)076%3C1157:IMOSFT2.0.CO;2>

35. Clemen RT. Combining forecasts. A review and annotated bibliography. Int J Forecast. 1989;5(4) 559–583. https://doi.org/10.1016/0169-2070(89)90012-5

36. Craven JP, Rudack DE, Shafer PE. National Blend of Models. A statistically post-processed multi-model ensemble. J Oper Meteorol. 2020;8(1) 1–14. https://doi.org/10.15191/nwajom.2020.0801

37. Glahn HR, Ruth DP. The new digital forecast database of the National Weather Service. Bull Am Meteorol Soc. 2003;84(2) 195–201. https://doi.org/10.1175/BAMS-84-2-195

38. Sanders F. On subjective probability forecasting. J Appl Meteorol. 1963;2(2) 191–201. https://doi.org/10.1175/1520-0450(1963)002%3C0191:OSPF2.0.CO;2>

39. Murphy AH, Winkler RL. Reliability of subjective probability forecasts of precipitation and temperature. J R Stat Soc Ser C Appl Stat. 1977;26(1) 41–47. https://doi.org/10.2307/2346866

40. Baars JA, Mass CF. Performance of National Weather Service forecasts compared to operational, consensus, and weighted model output statistics. Weather Forecast. 2005;20(6) 1034–1047. https://doi.org/10.1175/WAF896.1

41. Novak DR, Bright DR, Brennan MJ. Operational forecaster uncertainty needs and future roles. Weather Forecast. 2008;23(6) 1069–1084. https://doi.org/10.1175/2008WAF2222142.1

42. Leith CE. Theoretical skill of Monte Carlo forecasts. Mon Weather Rev. 1974;102(6) 409–418. https://doi.org/10.1175/1520-0493(1974)102%3C0409:TSOMCF2.0.CO;2>

43. Gneiting T, Raftery AE. Weather forecasting with ensemble methods. Science. 2005;310(5746) 248–249. https://doi.org/10.1126/science.1115255

44. Gneiting T, Raftery AE, Westveld AH, Goldman T. Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. Mon Weather Rev. 2005;133(5) 1098–1118. https://doi.org/10.1175/MWR2904.1

45. Vannitsem S, Bremnes JB, Demaeyer J, Evans GR, Flowerdew J, Hemri S, et al. Statistical post-processing for weather forecasts. Review, challenges, and avenues in a big data world. Bull Am Meteorol Soc. 2021;102(3) E681–E699. https://doi.org/10.1175/BAMS-D-19-0308.1

46. Bi K, Xie L, Zhang H, Chen X, Gu X, Tian Q. Accurate medium-range global weather forecasting with 3D neural networks. Nature. 2023;619(7970) 533–538. https://doi.org/10.1038/s41586-023-06185-3

47. Lam R, Sanchez-Gonzalez A, Willson M, Wirnsberger P, Fortunato M, Alet F, et al. Learning skillful medium-range global weather forecasting. Science. 2023;382(6677) 1416–1421. https://doi.org/10.1126/science.adi2336

48. Price I, Sanchez-Gonzalez A, Alet F, Andersson TR, El-Kadi A, Masters D, et al. Probabilistic weather forecasting with machine learning. Nature. 2025;637(8044) 84–90. https://doi.org/10.1038/s41586-024-08252-9

49. Lang S, Alexe M, Chantry M, Dramsch J, Pinault F, Raoult B, et al. AIFS, ECMWF's data-driven forecasting system [Preprint]. arXiv 2406.01465. 2024.

50. Rasp S, Hoyer S, Merose A, Langmore I, Battaglia P, Russell T, et al. WeatherBench 2. A benchmark for the next generation of data-driven global weather models. J Adv Model Earth Syst. 2024;16(6) e2023MS004019. https://doi.org/10.1029/2023MS004019

51. Roll R. Orange juice and weather. Am Econ Rev. 1984;74(5) 861–880.

52. Boudoukh J, Richardson M, Shen Y, Whitelaw RF. Do asset prices reflect fundamentals? Freshly squeezed evidence from the OJ market. J Financ Econ. 2007;83(2) 397–412. https://doi.org/10.1016/j.jfineco.2004.06.017

53. Zeng L. Weather derivatives and weather insurance. Concept, application, and analysis. Bull Am Meteorol Soc. 2000;81(9) 2075–2082. https://doi.org/10.1175/1520-0477(2000)081%3C2075:WDAWIC2.3.CO;2>

54. Dutton JA. Opportunities and priorities in a new era for weather and climate services. Bull Am Meteorol Soc. 2002;83(9) 1303–1312. https://doi.org/10.1175/1520-0477-83.9.1303

55. Jewson S, Brix A. Weather derivative valuation. The meteorological, statistical, financial and mathematical foundations. Cambridge. Cambridge University Press; 2005.

56. Roulston M, Kaivanto K. Can expert prediction markets forecast climate-related risks? Bull Am Meteorol Soc. 2024;105(10) E1898–E1914. https://doi.org/10.1175/BAMS-D-24-0135.1

57. Hsu SL. A prediction market for climate outcomes. Univ Colo Law Rev. 2011;83 179. https://doi.org/10.2139/ssrn.1770882

58. Vandenbergh MP, Raimi KT, Gilligan JM. Energy and climate change. A climate prediction market. UCLA Law Rev. 2013;61 1962.

59. Lucas GM Jr, Mormann F. Betting on climate policy. Using prediction markets to address global warming. UC Davis Law Rev. 2018;52 1429–1486.

60. Aliakbari E, McKitrick R. Information aggregation in a prediction market for climate outcomes. Energy Econ. 2018;74 97–106. https://doi.org/10.1016/j.eneco.2018.06.002

61. Roulston M, Kaplan T, Day B, Kaivanto K. Prediction-market innovations can improve climate-risk forecasts. Nat Clim Change. 2022;12(10) 879–880. https://doi.org/10.1038/s41558-022-01467-6

62. Boslough M. Using prediction markets to generate probability density functions for climate change risk assessment. American Geophysical Union Fall Meeting; 2011. Abstract GC11A-0885.

63. Hanson R. Combinatorial information market design. Inf Syst Front. 2003;5(1) 107–119. https://doi.org/10.1023/A:1022058209073

64. Roulston M, Kaivanto K, Toumi R. Tracking evolving views of the 2024 Atlantic hurricane season with an expert prediction market. Weather. 2025;80(8) 277–280. https://doi.org/10.1002/wea.7730

65. Roulston MS, Kaivanto K. Joint-outcome prediction markets for climate risks. PLoS One. 2024;19(8) e0309164. https://doi.org/10.1371/journal.pone.0309164

66. Crosier AW. Prediction markets beat the weather forecast on tomorrow's high temperature [Preprint]. arXiv 2609.23969. 2026. Available from https://arxiv.org/abs/2609.23969

67. ForecastEx LLC. Daily temperature forecast contract terms and conditions [Internet]. Chicago. ForecastEx; 2026 [cited 2026 Sep 16]. Available from https://data.forecastex.com/regulatory/DailyTemperatureTermsandConditions.pdf

68. National Oceanic and Atmospheric Administration, Department of Defense, Federal Aviation Administration, United States Navy. Automated Surface Observing System (ASOS) user's guide. 1998. Available from https://www.weather.gov/media/asos/aum-toc.pdf

69. Zippenfenig P. Open-Meteo.com weather API [Software]. Zenodo; 2023. https://doi.org/10.5281/zenodo.7970649

70. Janjić ZI, Gerrity JP Jr, Nickovic S. An alternative approach to nonhydrostatic modeling. Mon Weather Rev. 2001;129(5) 1164–1178.

71. NOAA National Weather Service Meteorological Development Laboratory. NAM MOS [Internet]. Silver Spring (MD). [cited 2026 Oct 2]. Available from https://vlab.noaa.gov/web/mdl/nam-mos

72. Deutscher Wetterdienst. Model Output Statistics-MIX (MOSMIX) [Internet]. Offenbach am Main. [cited 2026 Oct 2]. Available from https://www.dwd.de/EN/ourservices/met_application_mosmix/met_application_mosmix.html

73. National Weather Service. National Digital Forecast Database and local database description and specifications. NWS Instruction 10-201. Silver Spring (MD). National Oceanic and Atmospheric Administration; 2 January 2019.

74. Lang S, Alexe M, Clare MCA, Roberts C, Adewoyin R, Ben Bouallègue Z, et al. AIFS-CRPS. Ensemble forecasting using a model trained with a loss function based on the continuous ranked probability score [Preprint]. arXiv 2412.15832. 2024.

75. Molteni F, Buizza R, Palmer TN, Petroliagis T. The ECMWF ensemble prediction system. Methodology and validation. Q J R Meteorol Soc. 1996;122(529) 73–119. https://doi.org/10.1002/qj.49712252905

76. European Centre for Medium-Range Weather Forecasts. IFS documentation [Internet]. Reading. [cited 2026 Oct 2]. Available from https://www.ecmwf.int/en/publications/ifs-documentation

77. Han J, Pan HL. Revision of convection and vertical diffusion schemes in the NCEP Global Forecast System. Weather Forecast. 2011;26(4) 520–533. https://doi.org/10.1175/WAF-D-10-05038.1

78. NOAA National Centers for Environmental Prediction. Global Forecast System [Internet]. College Park (MD). [cited 2026 Oct 2]. Available from https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php

79. Zhou X, Zhu Y, Hou D, Fu B, Li W, Guan H, et al. The development of the NCEP Global Ensemble Forecast System version 12. Weather Forecast. 2022;37(6) 1069–1084. https://doi.org/10.1175/WAF-D-21-0112.1

80. Zängl G, Reinert D, Rípodas P, Baldauf M. The ICON (ICOsahedral Non-hydrostatic) modelling framework of DWD and MPI-M. Description of the non-hydrostatic dynamical core. Q J R Meteorol Soc. 2015;141(687) 563–579. https://doi.org/10.1002/qj.2378

81. Deutscher Wetterdienst. Numerical weather prediction models, ICON and ICON-EPS [Internet]. Offenbach am Main. [cited 2026 Oct 2]. Available from https://www.dwd.de/EN/ourservices/nwp_forecast_data/nwp_forecast_data.html

82. Côté J, Gravel S, Méthot A, Patoine A, Roch M, Staniforth A. The operational CMC-MRB Global Environmental Multiscale (GEM) model. Part I, design considerations and formulation. Mon Weather Rev. 1998;126(6) 1373–1395.

83. Houtekamer PL, Deng X, Mitchell HL, Baek SJ, Gagnon N. Higher resolution in an operational ensemble Kalman filter. Mon Weather Rev. 2014;142(3) 1143–1162. https://doi.org/10.1175/MWR-D-13-00138.1

84. Environment and Climate Change Canada. Global Ensemble Prediction System [Internet]. [cited 2026 Oct 2]. Available from https://eccc-msc.github.io/open-data/msc-data/nwp_geps/readme_geps_en/

85. Walters D, Baran AJ, Boutle I, Brooks M, Earnshaw P, Edwards J, et al. The Met Office Unified Model Global Atmosphere 7.0/7.1 and JULES Global Land 7.0 configurations. Geosci Model Dev. 2019;12(5) 1909–1963. https://doi.org/10.5194/gmd-12-1909-2019

86. Déqué M, Dreveton C, Braun A, Cariolle D. The ARPEGE/IFS atmosphere model. A contribution to the French community climate modelling. Clim Dyn. 1994;10(4–5) 249–266. https://doi.org/10.1007/BF00208992

87. Météo-France. Modèle ARPEGE [Internet]. [cited 2026 Oct 2]. Available from https://donneespubliques.meteofrance.fr/?fond=produit&id_produit=130&id_rubrique=51

88. Nakagawa M. Outline of the high resolution global model at the Japan Meteorological Agency. RSMC Tokyo-Typhoon Center Technical Review. 2009;11 1–13.

89. Japan Meteorological Agency. Outline of the operational numerical weather prediction at the Japan Meteorological Agency [Internet]. Tokyo. [cited 2026 Oct 2]. Available from https://www.jma.go.jp/jma/jma-eng/jma-center/nwp/outline-nwp/index.htm

90. Dowell DC, Alexander CR, James EP, Weygandt SS, Benjamin SG, Manikin GS, et al. The High-Resolution Rapid Refresh (HRRR). An hourly updating convection-allowing forecast model. Part I. Motivation and system description. Weather Forecast. 2022;37(8) 1371–1395. https://doi.org/10.1175/WAF-D-21-0151.1

91. National Weather Service Meteorological Development Laboratory. NDFD verification help [Internet]. Silver Spring (MD). [cited 2026 Sep 22]. Available from https://vlab.noaa.gov/web/mdl/ndfd-verification-help

92. Fréchet M. Sur les tableaux de corrélation dont les marges sont données. Ann Univ Lyon Sect A. 1951;14 53–77.

93. Matheson JE, Winkler RL. Scoring rules for continuous probability distributions. Manage Sci. 1976;22(10) 1087–1096. https://doi.org/10.1287/mnsc.22.10.1087

94. Hersbach H. Decomposition of the continuous ranked probability score for ensemble prediction systems. Weather Forecast. 2000;15(5) 559–570. https://doi.org/10.1175/1520-0434(2000)015%3C0559:DOTCRP2.0.CO;2>

95. Murphy AH, Winkler RL. A general framework for forecast verification. Mon Weather Rev. 1987;115(7) 1330–1338. https://doi.org/10.1175/1520-0493(1987)115%3C1330:AGFFFV2.0.CO;2>

96. Gneiting T, Balabdaoui F, Raftery AE. Probabilistic forecasts, calibration and sharpness. J R Stat Soc Series B Stat Methodol. 2007;69(2) 243–268. https://doi.org/10.1111/j.1467-9868.2007.00587.x

97. Cox DR. Two further applications of a model for binary regression. Biometrika. 1958;45(3–4) 562–565. https://doi.org/10.1093/biomet/45.3-4.562

98. Brown PT, Farnham DJ, Caldeira K. Meteorology and climatology of historical weekly wind and solar power resource droughts over western North America in ERA5. SN Appl Sci. 2021;3 814. https://doi.org/10.1007/s42452-021-04794-z

99. Brown PT, Li W, Cordero EC, Mauget SA. Comparing the model-simulated global warming signal to observations using empirical estimates of unforced noise. Sci Rep. 2015;5 9957. https://doi.org/10.1038/srep09957

100. Efron B, Tibshirani RJ. An introduction to the bootstrap. New York. Chapman and Hall; 1993.

101. Künsch HR. The jackknife and the bootstrap for general stationary observations. Ann Stat. 1989;17(3) 1217–1241. https://doi.org/10.1214/aos/1176347265

102. Hamill TM. Hypothesis tests for evaluating numerical precipitation forecasts. Weather Forecast. 1999;14(2) 155–167. https://doi.org/10.1175/1520-0434(1999)014%3C0155:HTFENP2.0.CO;2>

103. Holm S. A simple sequentially rejective multiple test procedure. Scand J Stat. 1979;6(2) 65–70.

104. Wilson EB. Probable inference, the law of succession, and statistical inference. J Am Stat Assoc. 1927;22(158) 209–212. https://doi.org/10.1080/01621459.1927.10502953

105. Iowa Environmental Mesonet. Model output statistics and LAMP archive [Internet]. Iowa State University. [cited 2026 Sep 15]. Available from https://mesonet.agron.iastate.edu/mos/

106. Brown PT. Weather prediction market accuracy [Internet]. [cited 2026 Oct 1]. Available from https://weather.weatherclimatehumansystems.org/accuracy.html