Forecast performance for a house in Cumberland Center, Maine: results as of 27 October 2026

Weather project, forecasting subproject. Dated analysis of the Forecasts and Model Checks page as built on lamp at 14:20 Eastern on 2026-10-27 (snapshot `forecasting/analyses/2026-10-27`, forecasting code 0.8.0). First edition, revision 2: adds 95% confidence intervals. The snapshot was retaken for this revision; no score quoted in the first version changed by more than 0.02.

Abstract

Purpose. This report records and explains what the Forecasts and Model Checks page showed on 27 October 2026. The page turns free weather forecasts from the National Oceanic and Atmospheric Administration (NOAA) into forecasts of solar production, heating cost and electricity use at one house, and checks each against what happened.

Methods. Each forecast is scored by how far ahead it was made (*lead time*), using the mean absolute error (MAE). Two simple guesses serve as baselines: *persistence* (the latest observation repeated) and the *normal* (the 1991-2020 average for the date). Past forecasts are replayed with models fitted only on earlier data. A block bootstrap, which resamples whole weeks, gives 95% confidence intervals (CI).

Results. Temperature forecasts for the next day miss by about 2.4 degrees Fahrenheit (degF) and are clearly better than the normal for 5 days. Solar production forecasts for tomorrow miss by 7.0 kilowatt-hours (kWh; 95% CI 5.7 to 8.2) on a typical 47 kWh day, against 17 kWh for the normal, and are clearly useful for 5 days. Heating degree day forecasts beat the normal by 62% (57 to 66%) for the coming week and are clearly useful for about two weeks. A gas bill forecast a month ahead misses by 24% (17 to 32%), hardly worse than the 22% with the weather known. Household electricity follows habits more than weather: tomorrow's forecast misses by 7.6 kWh (7.2 to 8.1) of a typical 42 kWh day, mostly gained by tracking recent use. The house weather station now reads about 1.4 degF colder than NOAA's analysis, and the gap is widening.

Conclusion. The weather forecasts are good enough to plan the next five days of solar production and heating. Gas and electricity forecast errors come mostly from household use, not weather. The house station needs attention.

Keywords: forecast verification, lead time, skill score, ensemble forecast, solar production, heating degree days, gas usage, household electricity, station quality control.

How to read this report

In plain terms. This is a dated snapshot. The web page changes every day; this report freezes the numbers and charts of one day and explains them, so later editions can be compared with it. Each figure is a chart from the page as it looked on 27 October 2026; the caption gives its number on the page. How the data is stored and how each check works is in the companion design report (also as a print PDF), referred to below as "the design report".

Contents

Abstract . . . 1
How to read this report . . . 1
1. Introduction . . . 3
2. Data and methods . . . 3
3. Results . . . 4
3.1 Temperature forecasts . . . 4
3.2 Heating degree days and the two-week cost outlook . . . 6
3.3 Gas bills forecast ahead . . . 7
3.4 Solar production . . . 8
3.5 Household electricity and net grid use . . . 10
3.6 Station checks . . . 11
4. Discussion . . . 12
5. Limitations . . . 13
6. Conclusions . . . 13
Data availability and reproducibility . . . 14
References . . . 14
Appendix. Abbreviations . . . 14

List of Figures

1 Hourly temperature forecast error by lead time, last 90 days (page Figure 2). . . . 5
2 Daily high and low forecasts scored against PRISM (page Figure 3). . . . 5
3 Daily heating degree day forecast error and range coverage by lead time, heating seasons since October 2020 (page Figure 6). . . . 6
4 Gas and heating electricity cost for the next two weeks, forecast against normal weather (page Figure 5). . . . 7
5 Average error per gas bill by how many days before the meter read the forecast was made (page Figure 8). . . . 8
6 Daily production against the PVGIS model driven by PRISM's measured sunlight (page Figure 11). . . . 9
7 Daily production forecast error by lead time, 2026-07-29 to 2026-10-24 (page Figure 13). . . . 10
8 Daily household electricity use against the daily mean temperature, last 365 days, with the change-point model (page Figure 15). . . . 11
9 Electricity forecast error by lead time: (A) household use, every GEFS run since October 2020; (B) net grid use, last 90 days (page Figure 17). . . . 11
10 Station temperature minus the HRRR analysis (top) and the house station against PRISM (bottom), 30-day means (page Figure 18). . . . 12

List of Tables

1 Analyses, data windows and sample sizes in this edition. . . . 4
2 Weekly HDD forecasts by heating season (October to April); errors in HDD, coverage of the 8-in-10 range. . . . 7
3 How far ahead each forecast beats the normal for the date (last lead day). . . . 13
4 Abbreviations used in this report. . . . 14

1. Introduction

In plain terms. The house in Cumberland Center, Maine, heats with gas, has a solar array (since January 2026), and keeps records of its gas bills, electricity use and solar production. The page uses free NOAA forecasts, republished by the non-profit dynamical.org, to answer five practical questions:

1. How good are temperature forecasts for this house, and how far ahead?
2. How much will the solar panels make in the coming days?
3. What will the next two weeks of heating cost?
4. How close would a gas bill forecast made weeks ahead come to the real bill?
5. How much electricity will the house use, and how much will it draw from the grid?

Every forecast on the page appears next to its own track record, because a forecast is only worth using if it has beaten simple guesses in the past. A sixth check asks whether the weather stations are still reading correctly.

Three NOAA models supply the forecasts:

- **High-Resolution Rapid Refresh (HRRR):** a 3 km model that looks 48 hours ahead.
- **Global Forecast System (GFS):** a global model that looks 16 days ahead.
- **Global Ensemble Forecast System (GEFS):** runs a global model 31 times from slightly different starts, to 35 days, so the spread of the runs shows how certain the forecast is.

Three sources serve as truth:

- **Observed weather:** the Parameter-elevation Regressions on Independent Slopes Model (PRISM) daily maps from Oregon State University, and two stations: the house's own station (ZASTMET) and the Portland International Jetport station (KPWM).
- **Energy:** the utility bills.
- **Solar:** the SolarEdge monitoring system.

2. Data and methods

In plain terms. Every check follows the same recipe.

1. Store every forecast as it was issued.
2. When the forecast day has passed, compare the forecast with what happened.
3. Average the size of the misses (the MAE) for each lead time, and compare it with the two baselines.

The *skill score* summarises the comparison: skill = 1 - MAE of the forecast / MAE of the baseline. A skill of 1 is perfect, 0 is no better than the baseline, and below 0 is worse (Murphy 1993). Where a forecast gives a range, the share of days on which the actual value fell inside it (the *coverage*) shows whether the range is honest. The target for an 8-in-10 range is 80%.

For energy, a *usage model* turns weather into kWh or therms. Feeding it the weather that actually happened gives the *known-weather floor*: the error that remains even with a perfect weather forecast. The gap between a forecast's error and that floor is the part caused by the weather forecast itself.

Uncertainty. Every score is an average over a limited sample, so the report gives a 95% *confidence interval* (CI) for the headline scores, written "7.0 kWh (5.7 to 8.2)". The intervals come from a moving-block bootstrap (Künsch 1989):

1. Group the scored cases by date, so all runs and stations for one day stay together.
2. Rebuild the sample from runs of 7 consecutive dates drawn at random (the circular variant, so the last dates can join the first), until it is as long as the original.
3. Recompute the MAE of the forecast and of each baseline on that sample, and the skill from the same draw.

4. Repeat 2,000 times; the 2.5th and 97.5th percentiles of the results form the interval.

The blocks matter because weather errors persist: a misforecast cold spell spoils several days in a row, so neighbouring days carry less independent information than their count suggests. Seven days is about the life of a weather pattern. Gas bills are resampled the same way, in runs of 3 consecutive bills (about one season). A forecast is called *clearly* better than a baseline when the lower end of its skill interval is above zero. The intervals are not sensitive to the block length: from 1-day to 14-day blocks, the interval for tomorrow's solar error barely changes (2.2 to 2.4 kWh wide), and those for tomorrow's electricity error (0.6 to 1.1 kWh wide) and the weekly HDD error (0.9 to 2.0 HDD wide) widen but lead to the same conclusions. The computation uses NumPy 2.5.4 and pandas 3.1.0 with a fixed random seed, so it gives the same intervals on every run.

The data windows differ by analysis (Table 1). The design report gives the full methods in section 4.3 and the station checks in section 4.2.

Table 1. Analyses, data windows and sample sizes in this edition.

Analysis	Forecast	Truth	Window	Cases
Temperature	HRRR and GFS, every run	ZASTMET and KPWM hourly; PRISM daily high and low	Last 90 days (from 2026-07-29)	363 HRRR and 364 GFS runs
Heating degree days	GEFS, every daily run	KPWM daily mean temperature	Heating seasons since October 2020	1,198 weekly forecasts
Gas bills	GEFS runs replayed 0-35 days before each meter read	Billed therms	Bills 2021-10-05 to 2026-08-05	46 bills (15 in winter)
Solar production	HRRR (days 0-1) and GFS (days 0-15) sunlight	SolarEdge daily production	2026-07-29 to 2026-10-24	88 days, up to 351 run-day pairs per lead
Electricity use	GEFS temperature through a usage model	Daily household use (meter, SolarEdge)	2020-10-01 to 2026-10-26	2,218 runs; about 2,100 days per lead
Net grid use	GEFS use minus GFS solar	Use minus production	2026-07-29 to 2026-10-24	70-84 days per lead

3. Results

3.1 Temperature forecasts

In plain terms. For the next few hours, HRRR misses the hourly temperature at the house and the airport by 1.9 degF on average, and 2.7 degF on day 2 (Figure 1). GFS misses by 2.9 degF on day 2. Persistence ("same hour as yesterday") misses by 6.3 degF, so both models add a lot in the first days.

By one week, GFS misses the hourly temperature by 5.1 degF (4.4 to 6.0), no better than the usual temperature for that week and hour (4.9 degF, 4.2 to 5.6). For daily highs and lows the picture is similar (Figure 2):

- **Day 1:** GFS highs miss by 2.4 degF (1.9 to 2.9), against 4.5 degF (3.9 to 5.2) for the normal.
- **Horizon:** on average, GFS beats the normal out to day 6 for highs and day 8 for lows. Only through day 5 is the gain clear for both: from day 6 the skill intervals include zero.

Today's forecast (page Figure 1) is a two-sided week:

- **Wednesday 28 October:** a cold, clear night. HRRR takes the house to about 26 degF, below freezing; GFS stops near 34 degF.
- **Thursday 29 and Friday 30 October:** a warm spell, with highs near 66 degF, about 13 degF above normal.
- **Into November:** near-normal days follow.

The two models disagree by 8 degF on Wednesday's low. That is typical of calm, clear nights, when the coldest air settles in hollows: a 3 km model cell can be colder or warmer than one backyard.

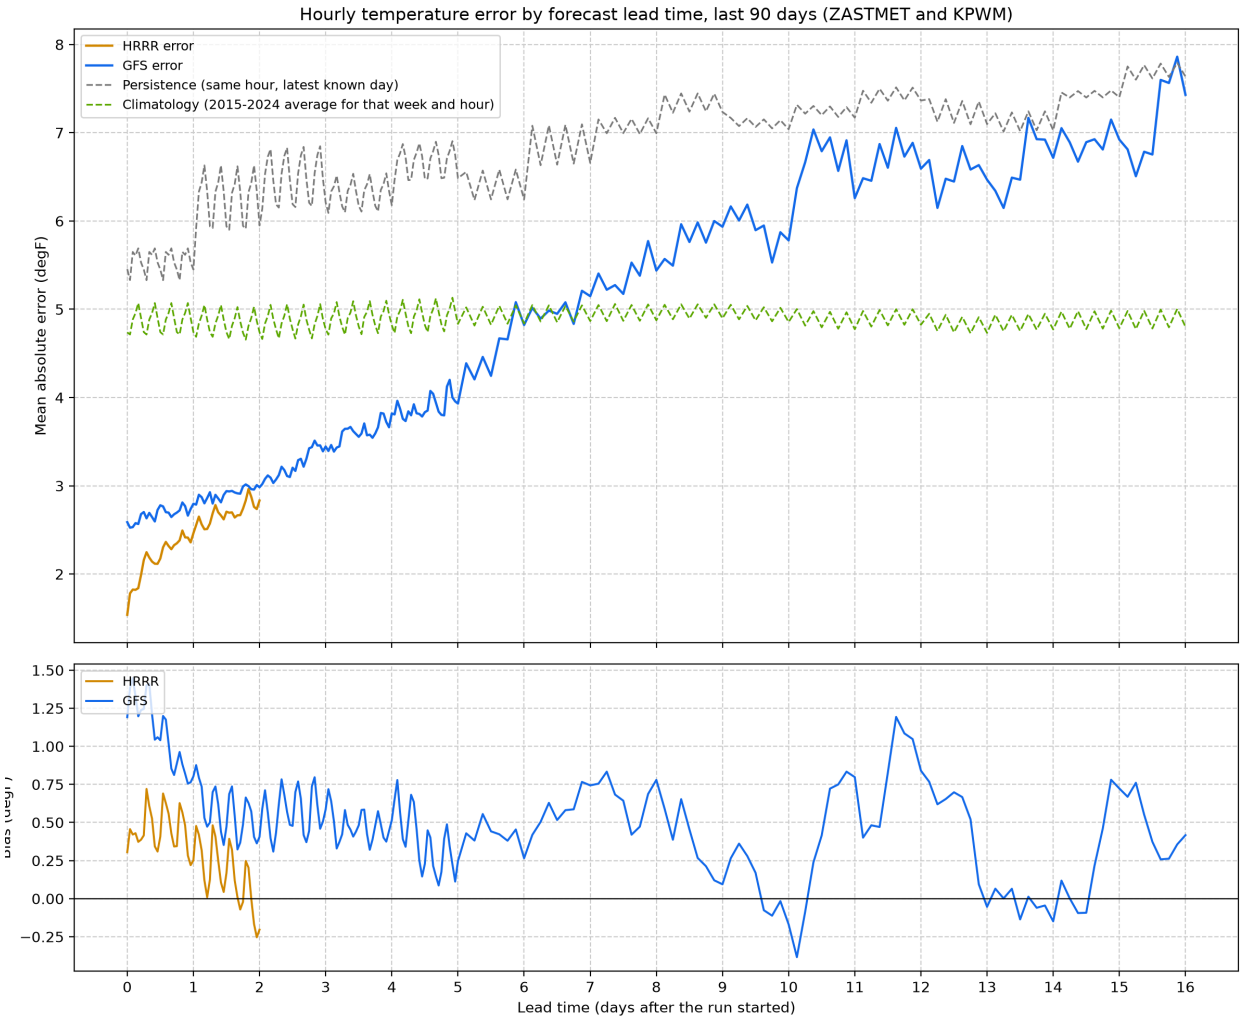


Figure 1. Hourly temperature forecast error by lead time, last 90 days (page Figure 2).

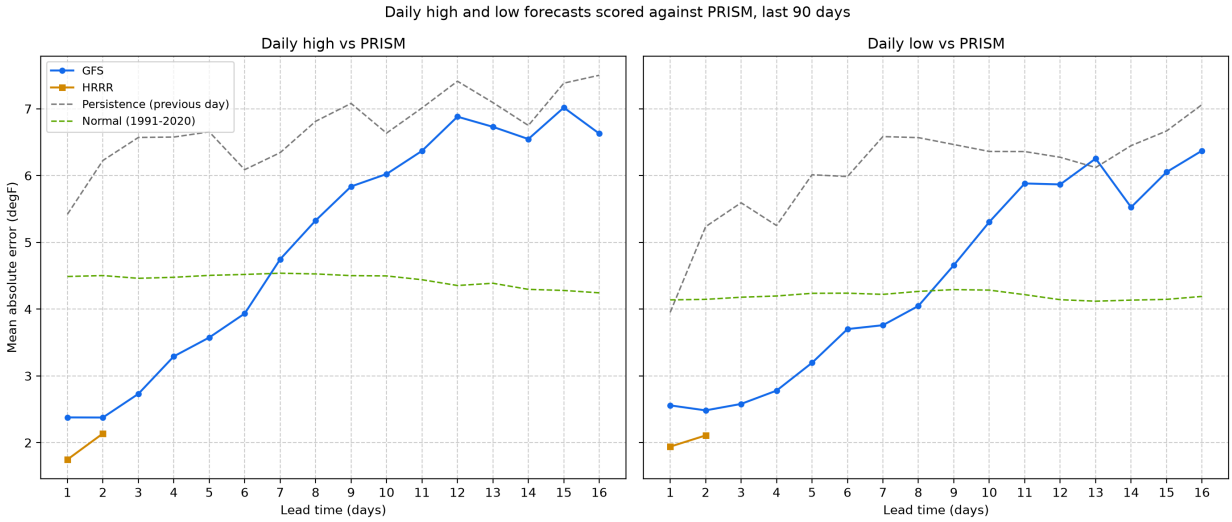


Figure 2. Daily high and low forecasts scored against PRISM (page Figure 3).

3.2 Heating degree days and the two-week cost outlook

In plain terms. *Heating degree days* (HDD) measure the cold: how far each day's mean temperature falls below 65 degF, summed over the days. The house's gas use follows them closely. Over six heating seasons, the GEFS forecast of the coming week's total missed by 10.1 HDD on average (9.3 to 11.0), against 26 for the normal and 41 for persistence (Figure 3). At today's gas price that is about $4 of heating a week. Day by day, the forecast is clearly better than the normal through day 13. From day 14 on, its skill stays below 0.05 and its intervals include or touch zero. The page, which counts any positive average, reports about day 26; that figure overstates how far ahead the forecast is useful.

The raw ensemble range is too narrow: it held the actual week only 55% of the time. A calibration fitted on past errors (shift the centre, widen the spread) brings this to 81% for the coming week (77 to 84%) and 77% for the week after (74 to 81%). Each season is scored with a calibration fitted on the other seasons only. Table 2 shows that this holds season by season, not just on average.

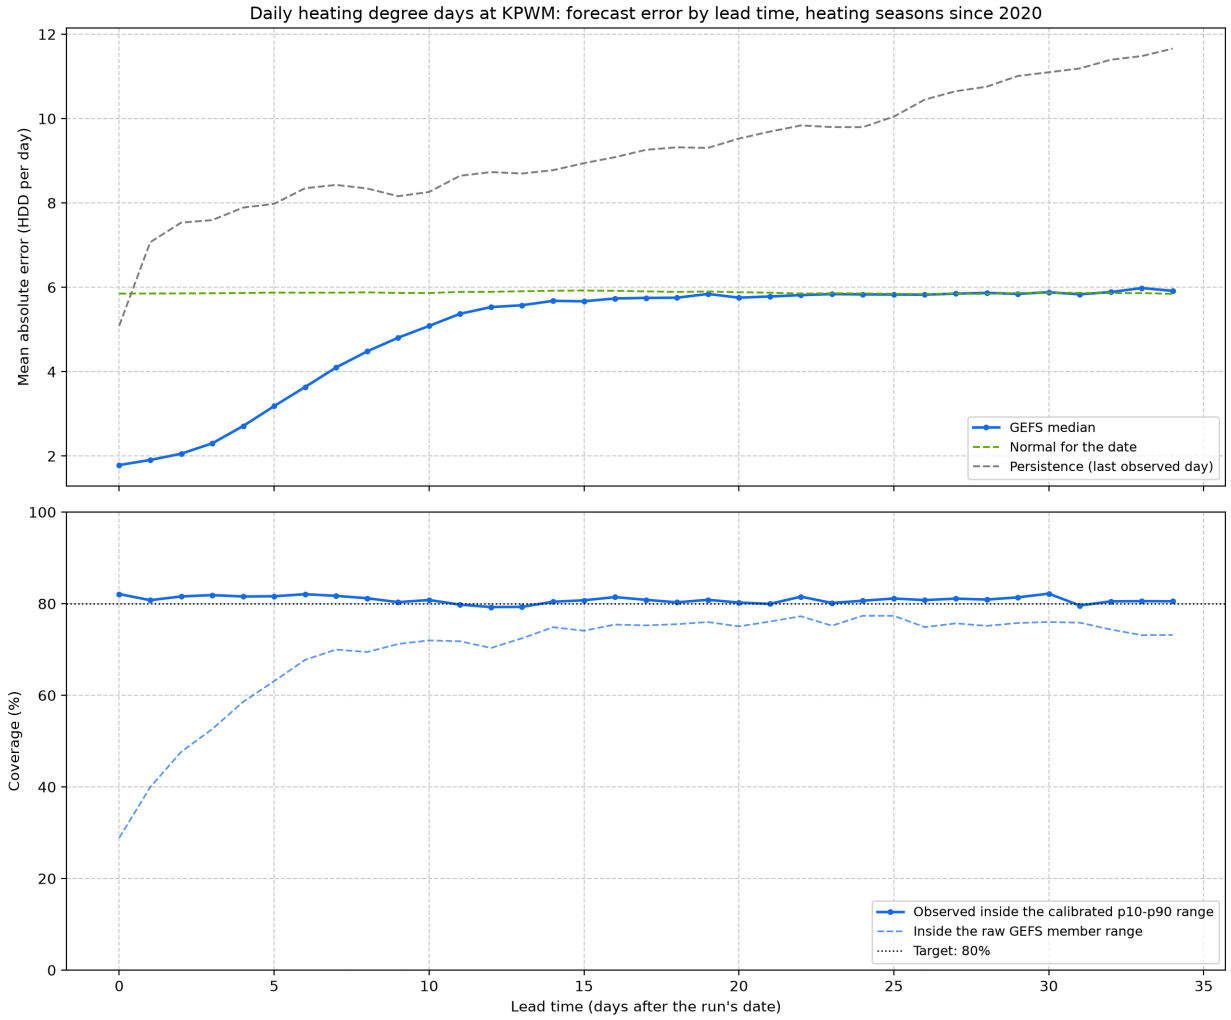


Figure 3. Daily heating degree day forecast error and range coverage by lead time, heating seasons since October 2020 (page Figure 6).

Table 2. Weekly HDD forecasts by heating season (October to April); errors in HDD, coverage of the 8-in-10 range.

Season	Weekly forecasts	GEFS error	Normal error	Persistence error	Inside raw range	Inside calibrated range	Inside normal range
2020-21	204	11	25	43	53%	75%	81%
2021-22	212	9	24	43	67%	86%	87%
2022-23	212	11	34	42	49%	79%	74%
2023-24	146	9	27	37	57%	86%	75%
2024-25	212	11	24	42	50%	78%	79%
2025-26	192	10	27	39	51%	82%	81%
2026-27 (so far)	20	7	8	30	60%	95%	100%

Today's outlook (Figure 4) is for a milder-than-normal fortnight:

- **This week:** $59 of gas and heating electricity (8 in 10 chance between $53 and $66), against $68 in normal weather. That is about 22 therms of gas.
- **Next week:** $67 ($55 to $79), against $72.

The dollars use the September bill's rate of $1.90 per therm. Winter supply prices have been much higher (over $3.80 per therm in March 2026), so as the season turns, the therms are the firmer number.

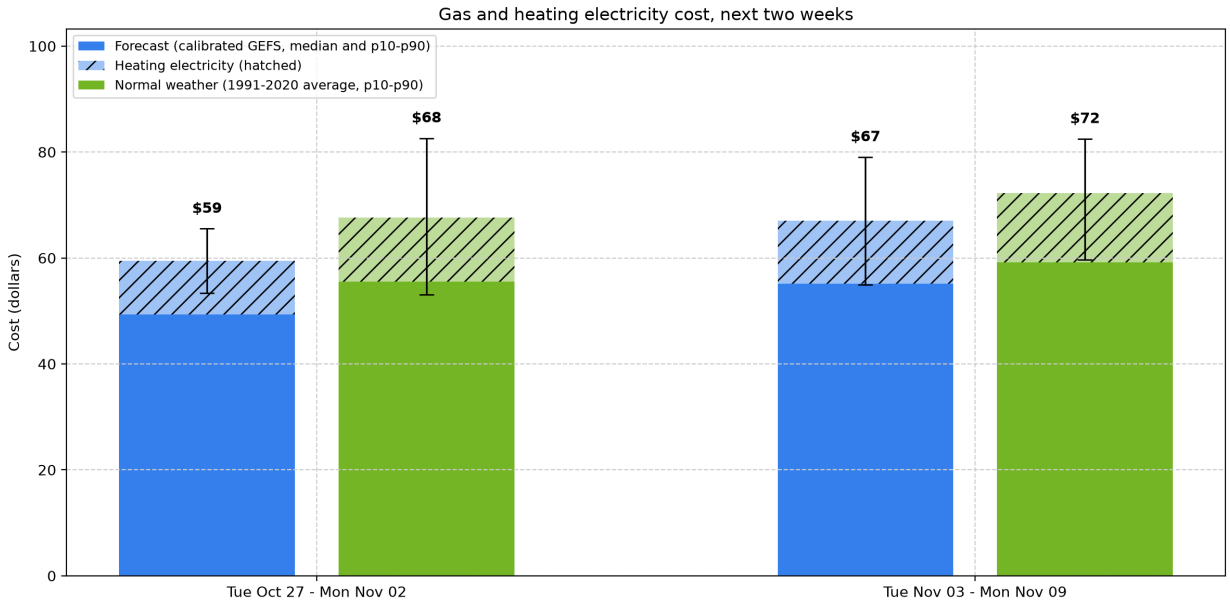


Figure 4. Gas and heating electricity cost for the next two weeks, forecast against normal weather (page Figure 5).

3.3 Gas bills forecast ahead

In plain terms. Gas is metered once a month, so the cost outlook can only be checked one bill at a time. Each of 46 bills since October 2021 was forecast again, as it would have been 0 to 35 days before its meter read (Figure 5).

- **Weather known (0 days):** the usage model misses by 22% per bill on average. Most of that comes from small summer bills; the 2026 summer bills were about double the model's base use, so something besides heating changed this year.
- **Uncertainty:** with 46 bills, the error itself is uncertain: 22% has an interval of 15 to 29%.

- **A month ahead:** forecast 29 to 35 days before the meter read, when almost the whole bill is still in the future, the error rises only to 24% (17 to 32%). Normal weather gives 29% and persistence 39%. Scored on the same bills, the forecast's advantage over the normal is clear: skill 0.17 (0.07 to 0.25).
- **Winter bills** (starting November to February, 15 bills): the usage model alone misses by 6% (4 to 8%), and the month-ahead forecast also by 6% (4 to 8%). Normal weather gives 11% and persistence 21%; the skill against the normal is 0.47 (0.20 to 0.65).
- **In dollars:** a month-ahead forecast is off by about $28 per bill at each bill's own rate, against $37 for the normal.

The weather forecast is not what limits the gas bill forecast; the usage model is. The usage model (0.62 therms a day plus 0.149 therms per HDD, fitted on 33 bills, r = 0.97) will improve most from understanding the higher non-heating use in 2026.

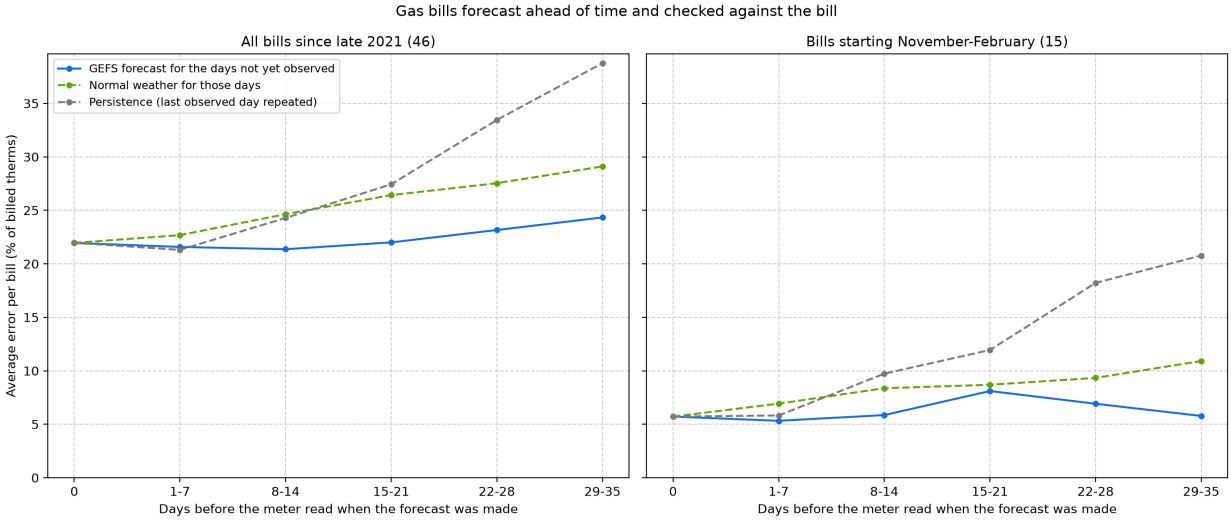


Figure 5. Average error per gas bill by how many days before the meter read the forecast was made (page Figure 8).

3.4 Solar production

In plain terms. With the actual sunlight known, the production model explains the array well (Figure 6):

- **How the model works:** PRISM's daily sunlight is converted with the European Commission's Photovoltaic Geographical Information System (PVGIS) model of this exact array.
- **How well it fits:** the array delivers 101% of the unshaded PVGIS model, with a correlation of r = 0.970 and a typical daily error of 8% over 264 ordinary days since January 2026.
- **Snow:** seven days made less than half the expected amount despite real sun. These are snow on the panels, which the model does not include.

Forecasts start from predicted rather than measured sunlight (Figure 7). Over the last 88 days, the array averaged 47 kWh a day:

- **Tomorrow:** GFS missed by 7.0 kWh (15%; 95% CI 5.7 to 8.2) and HRRR by 7.9 kWh (6.7 to 9.2). The normal missed by 16.8 kWh (14.5 to 19.1) and repeating the last day by 21.6 kWh.
- **Lead time:** the error grows steadily. GFS is clearly better than the normal through day 5. On days 6 and 7 it is better on average, but not clearly (skill intervals include zero), and it is worse from day 8.
- **Floor:** even with the sunlight known, the model misses by 4.3 kWh (3.3 to 5.3) using PRISM, or 5.7 kWh (4.9 to 6.6) using NOAA's HRRR analysis. So about 3 of tomorrow's 7 kWh of error comes from forecasting the weather.

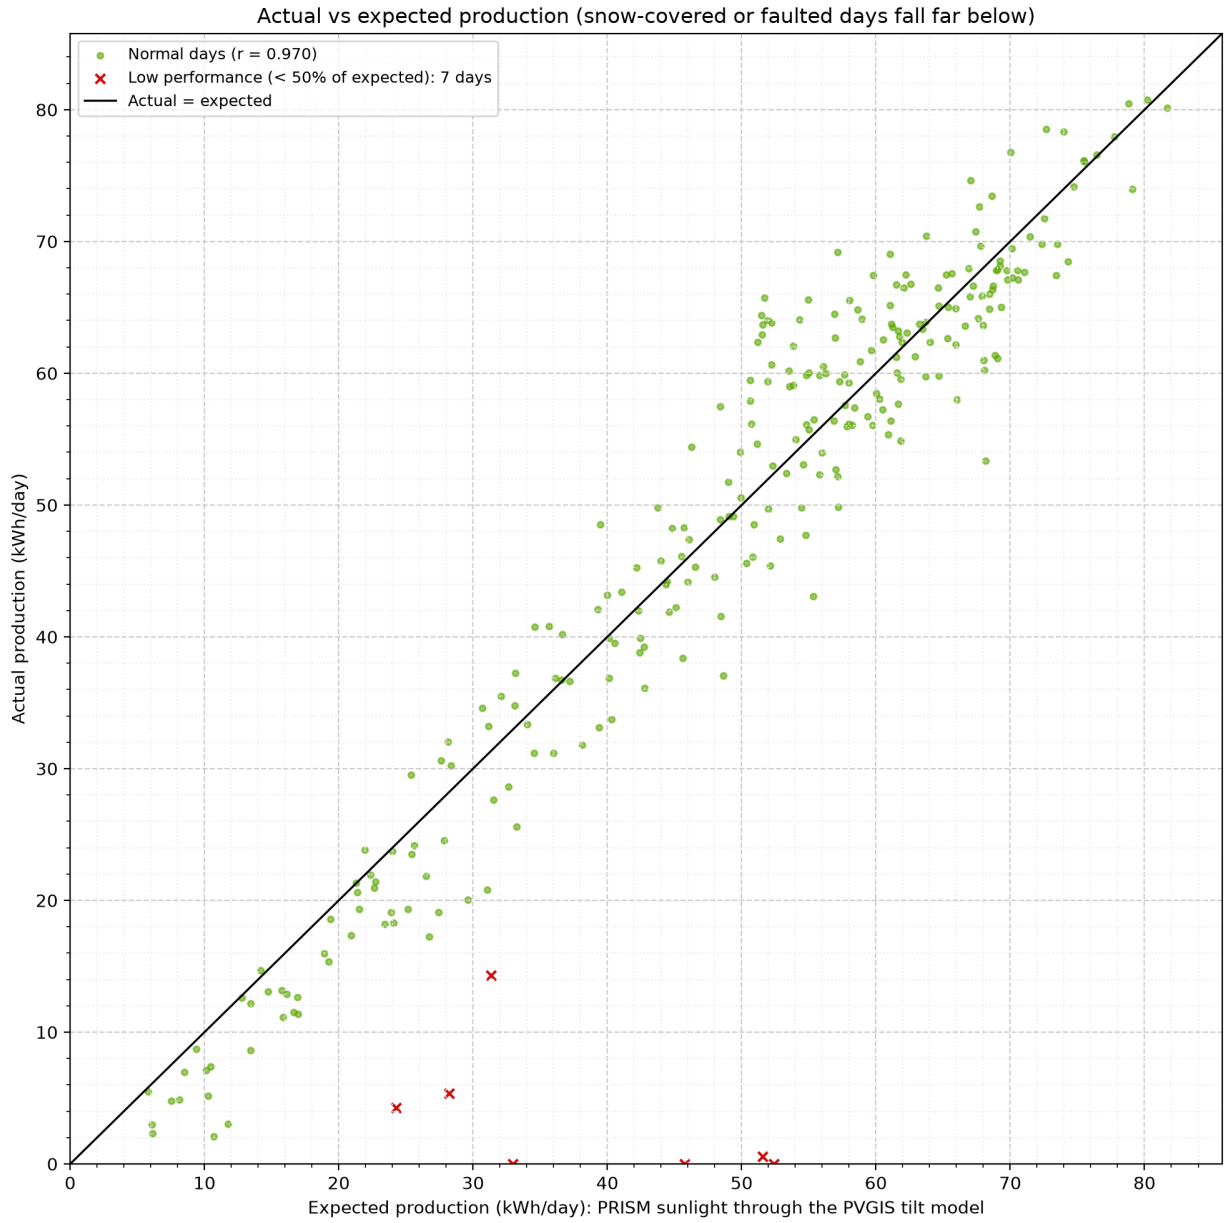


Figure 6. Daily production against the PVGIS model driven by PRISM's measured sunlight (page Figure 11).

- **HRRR:** NOAA's high-resolution model is no better than GFS for daily totals.

Today's forecast (page Figure 12) calls for:

- **This week:** five sunny days of about 49 to 56 kWh through Saturday 31 October, then about 2 kWh on Monday 2 November, under a thick overcast.
- **Tomorrow:** GFS forecasts 49 kWh and HRRR (run of 8 AM) 41 kWh.
- **Later in November:** forecasts made that far ahead have no skill, so read those days as "the normal, about 33 kWh".

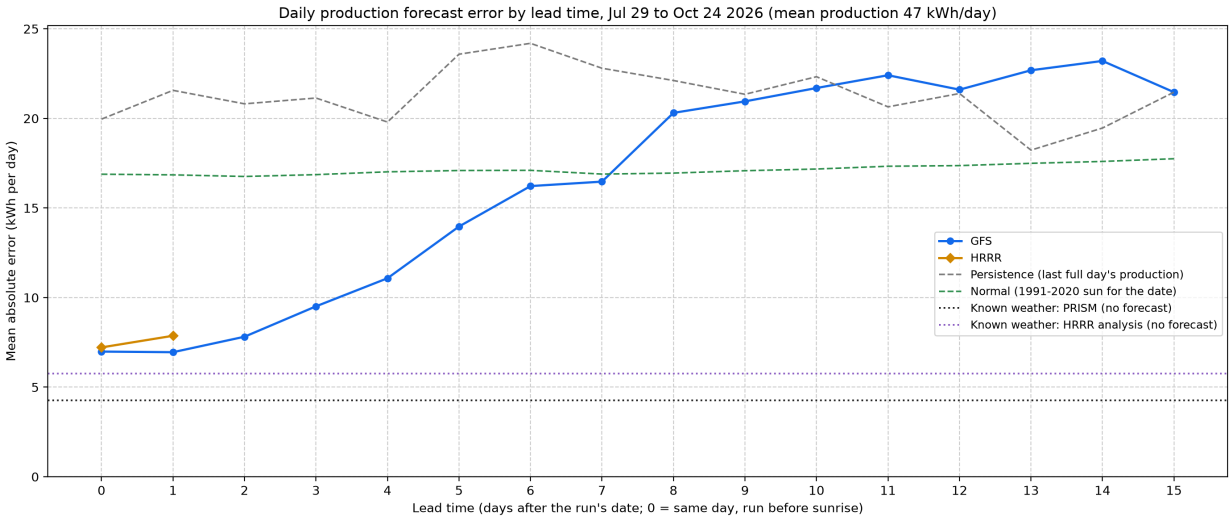


Figure 7. Daily production forecast error by lead time, 2026-07-29 to 2026-10-24 (page Figure 13).

3.5 Household electricity and net grid use

In plain terms. The house used 42 kWh a day on average since October 2020. A *change-point model* (Figure 8), fitted on the past year, describes how use depends on the day's mean temperature at the airport. The method comes from the Princeton Scorekeeping Method (Fels 1986), which is unrelated to the PRISM maps:

- **Base:** 39 kWh a day.
- **Cooling:** 1.5 kWh for each degree above 60 degF, from air conditioning.
- **Cold:** 0.8 kWh for each degree below 38 degF.
- **Weekends:** 2.2 kWh more.

Temperature explains only part of the day-to-day swing ($r = 0.61$; typical scatter 10 kWh a day).

Every GEFS run since October 2020 was replayed through this model, adding the household's *recent level*: how far use ran above or below the model over the previous four weeks (Figure 9, panel A).

- **Tomorrow:** the forecast misses by 7.6 kWh (7.2 to 8.1), against 9.5 kWh for the normal, 9.1 kWh for repeating the last day, and 8.0 kWh for the normal plus the recent level.
- **What the weather forecast adds:** only 0.4 kWh a day in the first week. That gain is small but real: its skill against the normal plus the level is 0.05 (0.02 to 0.07) tomorrow, and stays clearly above zero through day 11. After two weeks it is nil. The rest of the gain comes from tracking habits.
- **Floor:** even with the actual temperature known, the error is 7.5 kWh (7.1 to 8.0). For the current day, "same as yesterday" (7.4 kWh) is as good as any model.
- **Weekly totals:** these are steadier. For the coming week the forecast misses by 37 kWh (33 to 40) of a typical 293 kWh (12%), against 51 kWh for the normal and 52 kWh for persistence.
- **Range:** the 8-in-10 range held the actual day 83% of the time (80 to 85%).

Net grid use is use minus solar production: what the house draws from the grid, or sends to it when negative (Figure 9, panel B). Over the 88 days with solar forecasts:

- **Tomorrow:** the net forecast misses by 9.9 kWh (8.5 to 11.5), against 17.1 kWh (14.5 to 19.8) for normal weather.
- **Horizon:** it is clearly better through day 4 and better on average through day 5, the same horizon as the solar forecast. Beyond that, cloud cover is the larger unknown.

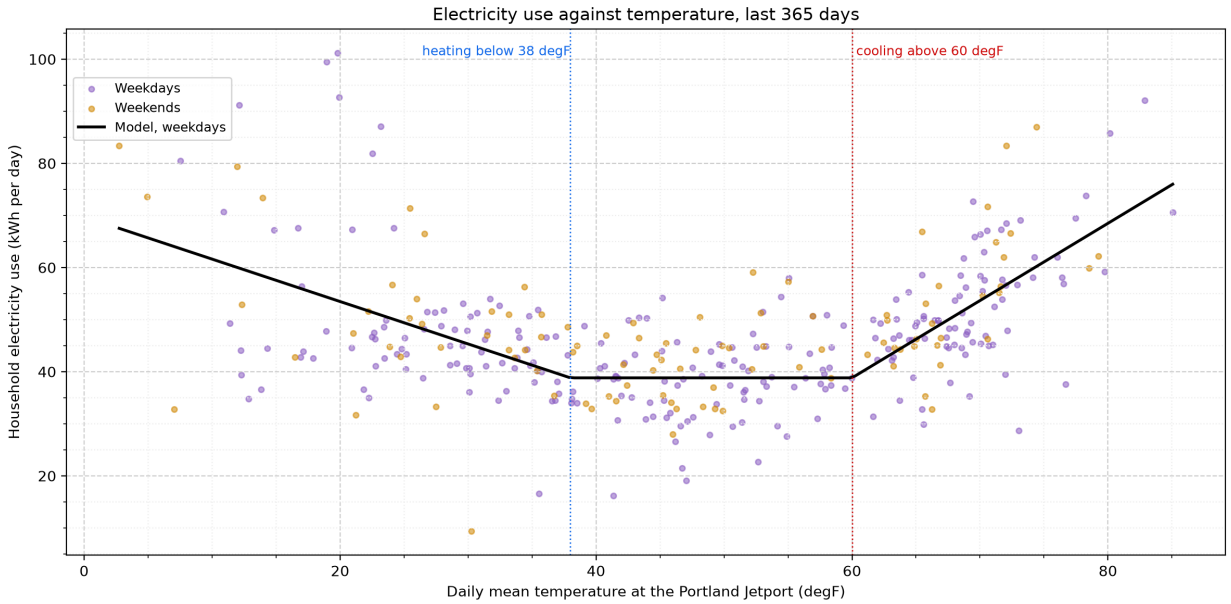


Figure 8. Daily household electricity use against the daily mean temperature, last 365 days, with the change-point model (page Figure 15).

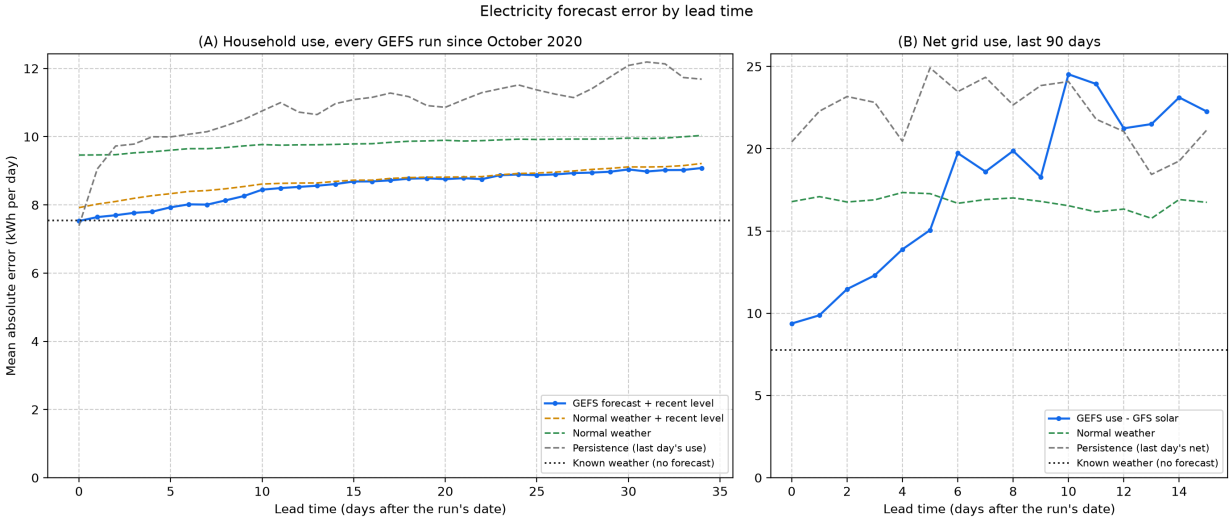


Figure 9. Electricity forecast error by lead time: (A) household use, every GEFS run since October 2020; (B) net grid use, last 90 days (page Figure 17).

Today's outlook (page Figure 16) expects the sunny week to roughly balance household use. Over the next seven days the house should use about 303 kWh and the panels make about 311 kWh, so it should export a little on net. Cloudy days in early November will reverse that.

3.6 Station checks

In plain terms. A weather station can drift slowly as its sensors age. Each station is compared with two references at the same spot (Figure 10):

- **NOAA's HRRR analysis:** the model's hourly best estimate.

- **PRISM:** daily highs and lows. PRISM does not use the house station, so it is a fully independent check.

What the comparison shows:

- **The house station (ZASTMET):** it averaged 0.8 degF colder than HRRR over the past year, while the airport (KPWM) agreed within 0.1 degF. The house station's offset has moved from about +0.5 degF in 2025 to -1.4 degF now, and it has fallen again since September 2026. Over the same period the airport's offset rose slightly, to +0.3 degF, so the two stations now disagree by about 1.7 degF.
- **Against PRISM:** over 499 days the house's highs average 1.0 degF above PRISM and its lows 0.4 degF below.
- **Gap:** the PRISM comparison has a gap from mid-August to late September 2026, while the HRRR comparison does not. The cause has not been checked yet.

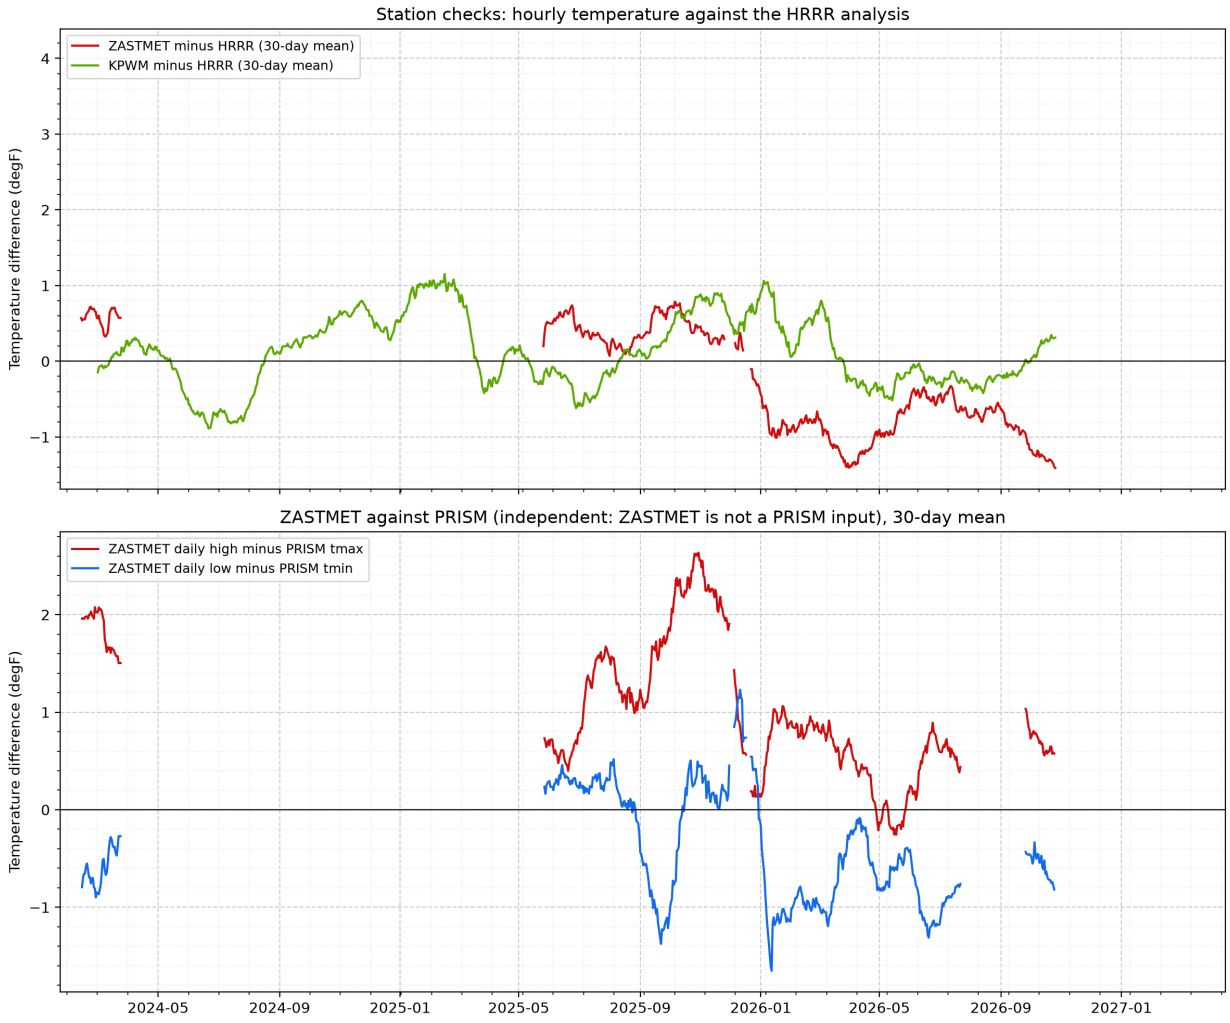


Figure 10. Station temperature minus the HRRR analysis (top) and the house station against PRISM (bottom), 30-day means (page Figure 18).

4. Discussion

In plain terms. Three patterns stand out in this edition.

First, **each forecast has its own useful horizon**, and it is shortest for sunshine (Table 3):

Table 3. How far ahead each forecast beats the normal for the date (last lead day).

Forecast	Better on average	Clearly better (95% CI)
Sunshine (GFS)	day 7	day 5
Net grid use	day 5	day 4
Temperature highs (GFS)	day 6	day 5
Temperature lows (GFS)	day 8	day 5
Degree days (GEFS, daily)	about day 26, with skill below 0.05 after day 13	day 13

Note. "Clearly better" means the lower end of the skill interval is above zero on that lead day and on every earlier one.

Clouds are hard to predict beyond a few days. Temperature averaged over days and weeks keeps more of the large-scale pattern that models predict well. For planning, the GFS sunshine forecast is worth reading for the coming week; beyond that, the normal is the better guess.

Second, **for energy, the household matters more than the weather**:

- **Gas:** a month-ahead bill forecast is almost as good as one made with the weather known, and both miss small summer bills by much more than winter ones.
- **Electricity:** the weather forecast adds only 0.4 kWh a day; following the recent level of use adds more. Better energy forecasts will come from understanding the household's base use (hot water, cooking, appliances, the 2026 change in summer gas use), not from better weather data.

Third, **the house station is drifting**. The offset against HRRR has widened twice, in December 2025 and again since September 2026, while the airport stayed steady. The design report (section 5.1.2) lists possible causes: an ageing sensor, a dirty or blocked radiation shield, or new shade near the mast. The forecast scores use both stations and PRISM, so they are not affected much. But any analysis that relies on the house temperature alone should be read with a 1 to 2 degF caution until the station is checked.

5. Limitations

- **Short solar record:** the solar and net-grid scores cover 88 days of late summer and autumn, with no snow. Winter scores may differ.
- **Few winter gas bills:** only 15 of the 46 bills are winter bills.
- **Model fit:** the production model's share of the PVGIS model is fitted on the days it is checked on, which flatters the known-weather floor slightly.
- **Interval assumptions:** the block bootstrap assumes errors more than about a week apart are nearly independent, and that the recent window represents the future. Slower changes, such as a new season or the drifting house station, are not covered by the intervals. With only 15 winter bills, the gas intervals are approximate.
- **Many comparisons:** the report reads skill intervals at many lead days; a few may exclude zero by chance. The horizons in Table 3 require every earlier lead day to pass as well.
- **Moving numbers:** the page rebuilds daily with new forecasts and with PRISM revisions (PRISM revises each day for six months), so its numbers will drift slightly from this snapshot.

6. Conclusions

1. Temperature and sunshine forecasts are clearly better than the normal for about 5 days, and degree-day forecasts for about two weeks. The two-week heating cost outlook rests on solid ground, and its calibrated ranges are honest (81% coverage for an 80% range).
2. Solar production can be forecast to within about 15% for tomorrow. About 40% of that error is the weather forecast; the rest is the production model.

3. Gas and electricity forecasts are limited by household behaviour more than by weather. Improving the usage models is the next step that would pay off.
4. The house weather station has drifted about 2 degF colder relative to NOAA's analysis since 2025 and should be inspected.
5. This week (from 27 October 2026): a frost risk at the house on Wednesday night, a warm Thursday and Friday, a heating cost below normal (about $59), and solar production roughly matching household use.

Data availability and reproducibility

- **This edition:** charts, numbers and this text are in `forecasting/analyses/2026-10-27/` of the weather repository (`numbers.json` holds every value quoted here).
- **How to produce the next edition:** see `forecasting/analyses/README.md`.
- **Code:** the analysis code is the forecasting package, version 0.8.0. The bootstrap is in `forecasting/verify.py` (`block_means`, `intervals`); every interval quoted here is in `numbers.json` as the `_lo` and `_hi` fields next to its score.
- **Data sources:** NOAA model data comes via dynamical.org under the CC BY 4.0 licence. PRISM data comes from the PRISM Group, Oregon State University. Household records are the project's own.

References

- Fels, M. F. (1986). PRISM: An introduction. *Energy and Buildings*, 9(1-2), 5-18. doi:10.1016/0378-7788(86)90003-4
- Künsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. *The Annals of Statistics*, 17(3), 1217-1241. doi:10.1214/aos/1176347265
- Murphy, A. H. (1993). What is a good forecast? An essay on the nature of goodness in weather forecasting. *Weather and Forecasting*, 8(2), 281-293. doi.org/10.1175/1520-0434(1993)008
- PRISM Group, Oregon State University. prism.oregonstate.edu.
- dynamical.org. NOAA HRRR, GFS and GEFS datasets. dynamical.org.
- Forecasting data design report (companion). forecasting_design.pdf.

Appendix. Abbreviations

Table 4. Abbreviations used in this report.

Abbreviation	Meaning
CI	Confidence interval (95%, from the block bootstrap)
degF, degC	Degrees Fahrenheit and Celsius (1 degF = 0.56 degC)
GEFS	Global Ensemble Forecast System (NOAA, 31 members, 35 days)
GFS	Global Forecast System (NOAA, 16 days)
HDD	Heating degree days: degrees the daily mean falls below 65 degF (18.3 degC), summed
HRRR	High-Resolution Rapid Refresh (NOAA, 3 km, 48 hours)
KPWM	Portland International Jetport weather station
kWh	Kilowatt-hour, unit of electrical energy
MAE	Mean absolute error: average size of the forecast errors, ignoring sign
NOAA	National Oceanic and Atmospheric Administration
PRISM	Parameter-elevation Regressions on Independent Slopes Model (Oregon State University daily maps); not the Princeton Scorekeeping Method of Fels (1986)
PVGIS	Photovoltaic Geographical Information System (European Commission)
Therm	Unit of gas energy on the bill, 100,000 British thermal units, about 29.3 kWh
ZASTMET	The weather station at the house