




RESEARCH ARTICLE

Beyond skill scores: exploring sub-seasonal forecast value through a case-study of French month-ahead energy prediction

Joshua Dorrington¹  | Isla Finney² | Tim Palmer¹  | Antje Weisheimer^{1,3} ¹Atmospheric, Oceanic and Planetary Physics, University of Oxford, Oxford, UK²Lake Street Consulting Ltd, London, UK³European Centre for Medium-Range Weather Forecasts, Reading, UK**Correspondence**

J. Dorrington, Atmospheric, Oceanic and Planetary Physics, Clarendon Laboratory, University of Oxford, Sherrington Road, Oxford OX1 3PU, UK.

Email:

joshua.dorrington@physics.ox.ac.uk

Funding information

UK NERC, Grant/Award Number: NE/L002612/1

Abstract

We quantify the value of sub-seasonal forecasts for a real-world prediction problem: the forecasting of French month-ahead energy demand. Using surface temperature as a predictor, we construct a trading strategy and assess the financial value of using meteorological forecasts, based on actual energy demand and price data. We show that forecasts with lead times greater than two weeks can have value for this application, both on their own and in conjunction with shorter-range forecasts, especially during boreal winter. We consider a cost/loss framework based on this example, and show that, while it captures the performance of the short-range forecasts well, it misses the marginal value present in medium-range forecasts. We also contrast our assessment of forecast value to that given by traditional skill scores, which we show could be misleading if used in isolation. We emphasise the importance of basing assessment of forecast skill on variables actually used by end-users.

KEYWORDS

applications, energy, forecasting, sub-seasonal prediction

1 | INTRODUCTION

Over the last 15 years, operational forecasting centres have increasingly extended forecasts into the 3–6 week time-scale, often called the monthly, extended or sub-seasonal range. Driven by the notion of seamless prediction (Hoskins 2013), the aim is to fill the gap between conventional two-week weather forecasts and longer-term seasonal projections. This can be done by harnessing the variability of slow drivers such as sea ice (Chevallier *et al.*, 2018), the land surface (Dirmeyer *et al.*,

2019), and atmospheric–oceanic processes such as the Madden–Julian Oscillation (Vitart 2017).

Despite sub-seasonal forecasting becoming well established, the applications and interpretation associated with these extended forecasts can still be unclear, both because robust skill is often only achievable for long time-averages of variables, and because the nature of forecast skill becomes inherently more probabilistic at longer lead times. End-users will get more value from forecasts if we can clearly understand the variables and time-scales that sub-seasonal forecasts can predict well, and where they

have more limited applicability (White *et al.*, 2017 gives a detailed discussion). However, much of the literature focused on verifying sub-seasonal forecast skill uses either mid-tropospheric, large-scale fields (Buizza *et al.*, 2005), or spatially localised station data (Monhart *et al.*, 2018), both of which are somewhat removed from end-user application which in many cases is interested in national averages at high temporal resolution.

Further to this, it is not always clear how directly forecast skill, as measured by the usual metrics, maps onto actual, realisable value to an end-user. This distinction underlines a gap that often exists between those who develop and produce forecasts and those who ultimately use them to make decisions. The importance of linking forecast development to user needs and of addressing the incommensurability of perspectives between these two groups, where one group speaks of correlation skill, and another annual expected return, was emphasised in Palmer (2002). His figure 1 captures schematically the idea of linking forecast verification metrics directly to end-user benefits. More recently, an extensive report on the future development of sub-seasonal forecasts (NASEM 2016) made the same point: their key findings 3.8 and 3.10 highlight the importance of developing dialogue and explanations of reasoning between model developers and users, while their recommendation B explicitly calls for stakeholder involvement in forming verification metrics. In short, we must recognise our models are only as good as end-users find them to be.

In this work we focus on this question of skill scores, taking a specific case-study of French energy markets as an example, and we compare progressively more idealised metrics of value to highlight the different conclusions they might imply if used in isolation. Precisely, we examine the problem of French month-ahead energy forecasting over a 9-year period. The need for improved analysis of sub-seasonal forecasts for the energy sector was emphasised in Orlov *et al.*, (2020) who remark that “the economic value of S2S forecasts has not been sufficiently investigated”, making our choice of case-study particularly relevant.

There has recently developed a valuable body of literature addressing the use of sub-seasonal and seasonal forecasts, many specifically targeted for energy sector applications. Soret *et al.*, (2019) demonstrated that for some extreme event case-studies, use of ensemble forecasts for surface temperature and wind speed in the midlatitudes had potential to improve decision-making even at week 4, while Van Der Wiel *et al.*, (2019) looked at sub-seasonal skill through the lens of weather regimes, showing they are associated with different distributions of energy-relevant variables such as demand and renewable generation. Explicit considerations of economic forecast value are found in Dutton *et al.*, (2013) and Emanuel

et al., (2012) for idealised hypothetical businesses based on trading on seasonal temperature terciles and hurricane activity, respectively. Lynch *et al.*, (2014) go further and presents a detailed analysis of wind energy predictability which shows skill in average wind speed over days 14–21, and possible increased returns of 60% from the utilisation of forecasts in decision-making for a realistic power-generation business case.

We aim to build on this basis, with the distinction that we look at the use of sub-seasonal forecasts at daily resolution, which is more desirable than weekly means for some users, and has not been often looked at beyond two-week lead times in this context. A main difference from prior work in this area is that, while we offer insights into the potential value of long-range forecasts for the French energy market, our primary focus is to emphasise the value of assessing model skill in a user-centred context.

Section 2 discusses the details of the forecast systems and data pre-processing used in this study. In Section 3 we explain the use case of energy stakeholders trading power on the French energy market. We directly compute the value of current sub-seasonal forecasts in this case, using real-world power price and demand data. In Section 4 we extract the users' cost-loss ratio from the energy example, and so can compute the commonly used metric of potential economic value (PEV), which we can then directly compare to our more detailed value framework, in order to assess the validity of the underlying assumptions. Moving to purely meteorological scores typically used by the forecast verification community in Section 5, we examine several common skill metrics for daily temperature forecasts. We can then see how our perspective on forecast value is altered depending on the verification method used. We finish in Section 6 by commenting on the implications of these results, both for using sub-seasonal forecasts for energy, and for the way in which academic meteorology assesses forecast skill.

2 | DATA

We analyse hindcast data from three different operational forecasting systems, covering the common period 1999–2018. In Section 3, we use only forecasts for which corresponding energy data are also available, restricting us to 2010–2018. Figures S1–S4 assess whether any significant difference in skill can be seen between 1999–2010 and 2010–2018, and we observe only a slight increase in week 1 probabilistic skill. Bearing this minor difference in mind, we consider results calculated from the full and restricted datasets comparable to each other. The extended-range forecasting system (hereafter EC45) is initialised twice weekly as a seamless continuation of the

ECMWF 15-day forecast, using the same model cycle and coupled to the NEMO ocean model, running for 46 days. We use the model cycle IFS-45R1, one cycle behind the current operational cycle 46R1 at time of writing, in order to have a full hindcast dataset available. The first 15 days of the EC45 system are run at a spectral resolution of TCo639, and upscaled after that to TCo319, corresponding to ~18 and 36 km respectively. Data from the EC45 system for the 2018 boreal Summer was not available, representing the only departure from the common data period.

The ECMWF seasonal forecasting system, SEAS5, is initialised once per month using an older model cycle (CY43R1), but runs out to 7 months, at the same TCo319 resolution as the extended-range EC45 forecast (Stockdale *et al.*, 2018; Johnson *et al.*, 2019). We make use of the first 45 days of the SEAS5 data as a baseline against which to evaluate the utility of the EC45 system's increased initialisation frequency and model improvements.

To capture differences in sub-seasonal forecasts between modelling centres, we also include the Environmental Modelling Centre's sub-seasonal ensemble forecast system (hereafter SubX), which is being run weekly in real time out to 35 days on an experimental basis as part of the SubX project (SubX 2019), with a comparable spectral resolution of T574 (~33 km). This project collates a number of different sub-seasonal hindcasts from different centres, however we chose not to use other model contributions due to limited hindcast periods and/or small ensemble sizes.

Analysis is performed on both annual and boreal winter (December to February, or DJF) datasets. The number of initialisation dates for each season, and ensemble size are summarised in Table 1. To validate the forecasts we make use of the ERA5 reanalysis (C3S 2017).

In all cases, we analyse daily 1200 UTC 2 m temperature data, on a 1° grid averaged over a box centred on mainland France. A gridpointwise, monthly and lead-time-dependent bias correction was applied, by mapping quantiles of each forecast's temperature distribution to those of ERA5 during the same time period. A 'drop one out' approach was used: for each hindcast year, data from only the other 19 years were used to bias correct. This prevents over-fitting, as long as we accept the mild assumption that the year-to-year correlation of surface temperatures is negligible. A gridpointwise sinusoidal fit to the annual cycle was computed for ERA5 using the first four harmonics and then removed from all datasets, as was a weak linear trend, to leave stationary anomaly fields. The area averaging was then performed by applying a land-sea mask and taking a cosine-latitude weighted mean over the region [5°W–8°E; 42°–51°N], to produce the final scalar anomaly field. This region includes a small number of gridpoints from outside the borders of France, but comparisons

using the ERA5 data show negligible differences in the T2m series obtained with a more detailed mask (Figures S5–S7).

3 | FORECAST VALUE IN AN ENERGY FRAMEWORK

3.1 | Methodology

In France as well as many European countries, energy to be delivered on a future date can be bought and sold through a liberalised national market. Power providers and traders are of course interested in buying these energy futures for an optimal price, a quantity which has many complex drivers such as macroeconomic factors, the price of fossil fuels which determines which power plants are operational, changes in installed capacity of renewables, and of course meteorological conditions (Hardt *et al.*, 2019 give a discussion of modelling non-meteorological impacts on energy price and demand on the inter-annual scale). We focus on baseline energy contracts, where energy is to be delivered throughout the day of interest, avoiding the issue of high-frequency sub-daily variability in the energy demand.

Energy contracts tend to be traded on discrete time-scales (quarterly, monthly, daily etc.). Taking the example of a month-ahead contract, the price of energy on a given day will be constant for the whole month of interest. For example in late January, a MWh of electricity purchased for 01 February is priced the same as a MWh for 28 February. However, electricity within the current calendar month is priced on a daily basis, leading to an increased sensitivity to meteorological conditions. While both these prices will evolve over time and a number of different time-scales exist, we consider only two prices of interest as a first approximation; the final month-ahead price, and the day-ahead price, as depicted schematically in Figure 1. The question then is: should energy be bought at the stable, close-to-climatology month-ahead price, or at the more volatile day-ahead price?

We would like to develop a plausible real-world assessment of the value of sub-seasonal surface temperature forecasts, while keeping the meteorological element front and centre. To do this we make use of real-world French energy price data covering the period 2010–2018, both month-ahead and day-ahead spot prices (calculated at around 1100 UTC the day before the target date) (EEX 2020). We also use daily averaged French demand data covering the same period (RTE 2019). The economic value we obtain here will be by necessity a lower-bound, as we are ignoring the economic and societal impacts extreme events can have on the energy sector beyond simple price

TABLE 1 A summary of the operational forecasting systems used in this work. Ensemble sizes are for the hindcast datasets used in this analysis; operational forecast ensembles are considerably larger. For more details see Section 2 and references therein

Name	Originating centre	Forecast period used	Initialisation frequency	No. of Annual/Dec-Feb initialisations	Time range of initialisation dates	Ensemble size
EC45	ECMWF	46 days	2 per week	2146/734	03 Jan 1999–30 May 2018	11
SEAS5	ECMWF	46 days	1 per month	219/ 78	01 Jan 1999–01 Aug 2018	25
GEFS	EMC (SUBX)	35 days	1 per week	1017/336	01 Jun 1999–30 May 2018	11




FIGURE 1 An example using real data of the evolution of energy prices for some given future date, in this case days in February 2018. During early January, the month-ahead February price (which is the same for all days in February) is the only power contract traded (solid blue line) in enough size to give a reliable benchmark of electricity cost. This varies day-to-day for a number of reasons including weather forecast changes, changes in scheduled availability of power plants, and varying coal and gas prices. As we go into February, days are priced and traded individually. A good simplification which sees real-world use is to consider only the day-ahead price, (black line) once the month has begun. This has increased volatility compared to the month-ahead pricing, reflecting the day-to-day variations in both power demand and supply. If the price spikes high, then buying at the end of the previous month is cheaper (shaded green), while if the price drops low, then buying at day-ahead price is cheaper (shaded red)

fluctuations. As one example, extended cold spells in December, January, February 2016/2017 led to a series of French nuclear plants having to come offline for emergency maintenance (ENTSOG 2017).

Modelling the evolution of energy prices is a discipline in and of itself, and we are not focused on producing a high-fidelity price model in this work. Instead we build a “good enough” trading strategy that will allow us to provide at least a lower bound on the usefulness of forecast data in this sector. Extensive prior work shows that surface temperature is the dominant driver of energy demand, with second-order factors including cloud cover, rainfall, wind speed, and (especially in warmer climates) relative humidity and wet bulb temperature (Lam *et al.*, 2008; Psiloglou *et al.*, 2009; Cho *et al.*, 2013; Thornton *et al.*, 2016). We focus only on temperature, and use anomalies of ERA5 daily surface temperature with respect to monthly means to derive a weakly quadratic relationship

between temperature and demand anomalies for each season. We find that, despite a whole host of missing factors, ERA5 T2m makes a reasonable single-variable predictor of demand with an annual correlation of 0.66, especially in winter when the demand for heating is highest, and so consequently the relationship between temperature and demand is strongest. Appendix A shows details of the fitted relationships, and the resulting demand estimates for each season.

We start from the basis that the anomalous demand calculated with respect to month is strongly correlated with the day-ahead price but not with the month-ahead price (not shown). This means that, if we can predict the energy demand from forecasts of surface temperature, then we can also predict whether it is cheaper to buy a unit of energy at the month-ahead price, or to wait until the day before and buy at the cost for immediate delivery, known as the spot price.

Given our estimate of demand, a basic trading strategy is:

- Predict the anomalous demand for a future date.
- If the demand anomaly is greater than a threshold d , buy energy at the month-ahead price.
- Else, buy energy the day before at the spot price.

We make a forecast of demand for every date in the record, making a decision on whether to buy at the month-ahead price at the end of the previous calendar month. This means that forecasts for early in the month will be more skilful than those for the end of the month; the lead time is not fixed. There are two simple situations we might imagine a user to be in: one where they wish to buy a set amount of energy (as for a trader for example), and one in which they wish to purchase a certain fraction of the total energy demand (as might be more relevant for an energy provider trying to prevent a production shortfall). In the first case we can evaluate the average cost per unit of power, C , using this strategy over T days with a predictor of anomalous demand, P_D , as:

$$C_{\text{set amount}} = \frac{1}{T} \sum_t^T (H[P_D(t) - d] \cdot p_{\text{month}} + H[d - P_D(t)] \cdot p_{\text{day}}), \quad (1)$$

where p_{month} is the month-ahead price, p_{day} , the day-ahead price, and H the Heaviside step function. In the second case, we must account for the fact that we are buying more power for likely high-demand days than for low-demand days, so we weight the cost of power by the ratio of the daily demand to the mean demand: $\tilde{D}_t := D_t / \langle D \rangle$. This gives us the average cost of buying a set fraction of power:

$$C_{\text{set fraction}} = \frac{1}{T} \sum_t^T \tilde{D}_t \cdot \{H[P_D(t) - d] \cdot p_{\text{month}} + H[d - P_D(t)] \cdot p_{\text{day}}\}. \quad (2)$$

As references, we consider four predictors that make no use of forecast data. We take two ‘perfect’ predictors, one in which we know the future demand anomaly precisely, and one in which we know only the future surface temperature perfectly, and must estimate the demand using our simple demand curve. This allows us to separate out any deficiency in meteorological skill from the drawbacks of our simplistic demand model. As climatological references, we consider the strategies of always buying at the day-ahead spot price, always buying month-ahead, and a random strategy where the probability of buying day-ahead and month-ahead are both 0.5 for any given

day. Buying for the day-ahead is on average cheaper than buying a month-ahead; a premium is paid for the lower volatility of the long-term contracts.

For each date in the record, we produce a demand forecast from each system using the most recently initialised forecast for which it was still possible to buy power at the month-ahead price. For example, a model may have three forecasts for 09 February, initialised on 02 February, 26 January and 19 January, at lead times of 7, 14 and 21 days respectively. In that case, while the 7-day forecast is most recent, it was initialised after the last opportunity to purchase at the month-ahead price (31 January) and so we would use the day 14 forecast instead. If no forecast had been available, we would instead use the climatological strategy of always buying on the day before. We use the forecasts’ ensemble mean temperature to predict demand. Taking the ensemble mean after predicting demand made no qualitative difference.

In order to better separate out the relative value of the short- and long-range forecasts, we also consider cases where only the first 15 days of the forecast are available, and the case where only long-range forecasts (lead times greater than 15 days) are used. We do not consider these additional cases for the seasonal SEAS5 system, as forecasts are initialised at the beginning of every month and so only lead times >28 days are ever available for decision-making. For the sub-seasonal systems, the median lead times using all forecasts, only the first two weeks, and only days 15+ are 18, 9 and 24 days respectively. Even though this is the same for both EC45 and SubX, the SubX lead times are skewed, with many fewer forecasts for days 1 to 6 contributing due to its lower initialisation frequency.

The obvious action threshold to choose is $d=0$; we should buy on any positive demand anomaly. However, due to the higher average price of month-ahead contracts, and the imperfect skill of forecasts, we also choose to consider a more cautious strategy of only buying in advance if we predict an annual upper tercile demand anomaly (≥ 54.4 GWh).

3.2 | Results

In this Section we present the difference in average energy price between the trading strategies discussed in Section 3.1 and the strategy of always buying energy the day before, which we note again as the best climatological strategy. In Figure 2 we look at the set energy amount scenario of Equation (1) while Figure 3 shows the set energy fraction case of Equation (2). Errors on these differences are estimated using 1,000 bootstrap resamples over forecast target dates, and those strategies whose fifth percentile

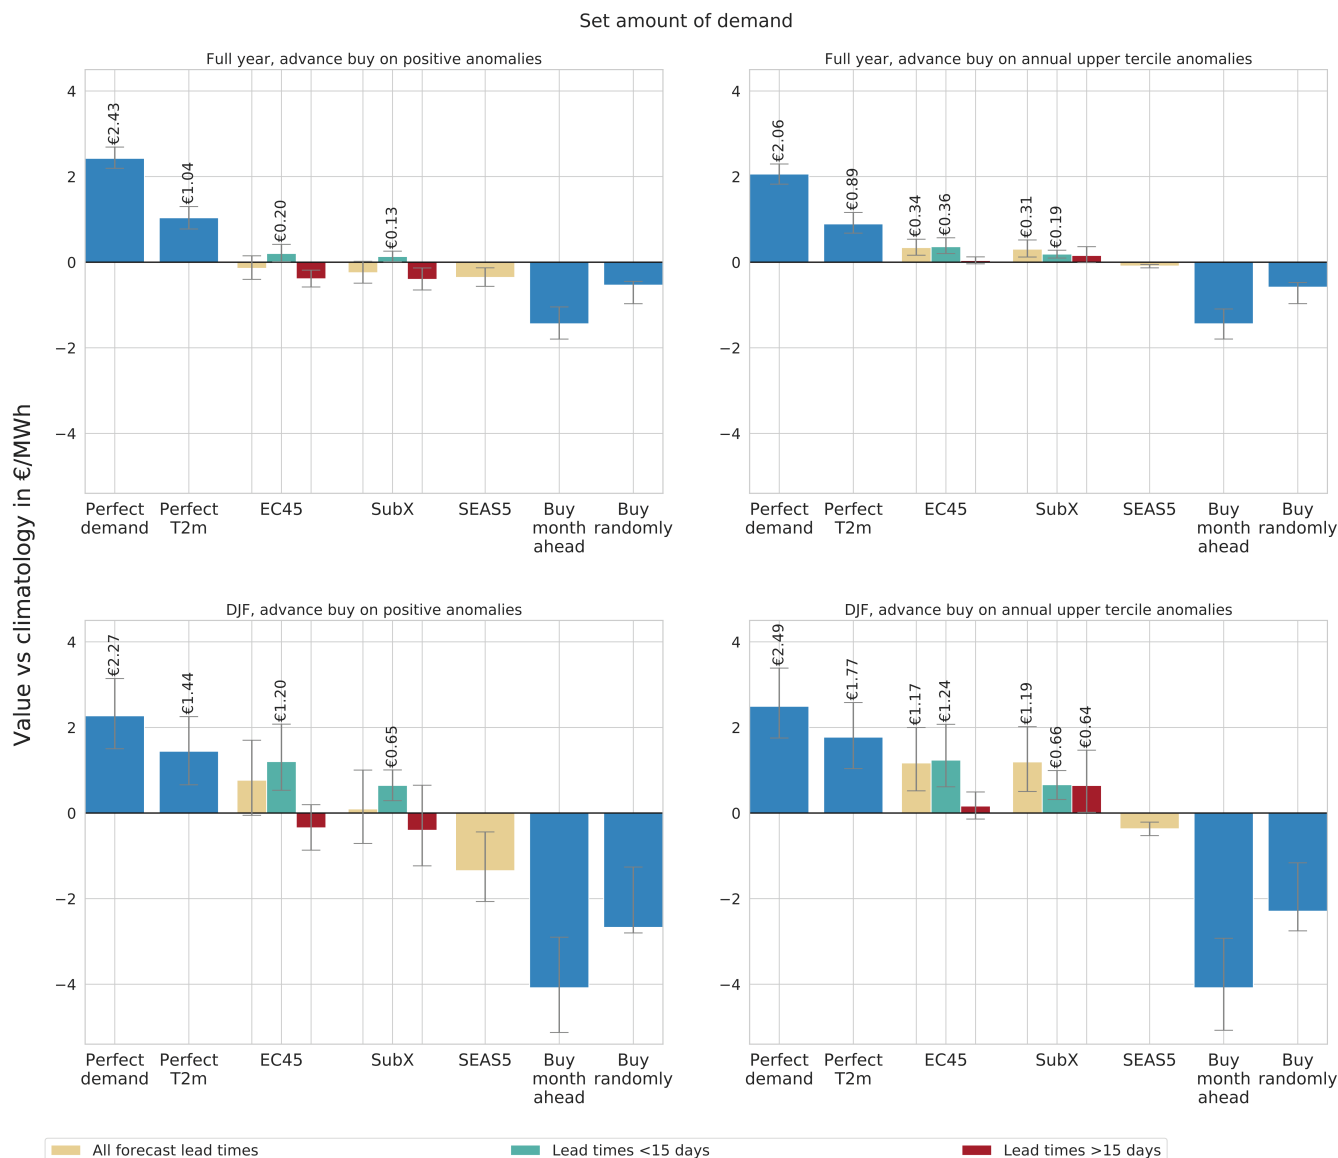


FIGURE 2 The average saving on a MWh of French baseline energy over the period 2010–2018 under different purchasing strategies compared to buying at the day-ahead price. (a, b) are the average cost over the full year, while (c, d) are for December, January, February only. (a, c) are based on buying at month-ahead price when the demand anomaly relative to the month is predicted to be positive, and (b, d) are based on buying at month-ahead rate when demand anomaly is in the upper tercile. Error bars show the 5th to 95th percentiles estimated by bootstrapping over target dates. Where the 5th percentile exceeds 0, the mean saving provided from using that strategy is marked above the bar

is positive are annotated with their associated potential saving per unit of energy. When references to statistically significant differences between forecast systems are noted, it is a result of analysis of this kind, but we do not show differences between all possible combinations of trading strategies for conciseness.

We start by looking at the value associated with a hypothetical perfect forecast of demand. Under our trading paradigm proposed here, there is the potential to save €2.50 per MWh for a user buying a set amount of baseline energy, while for a user buying a constant fraction of daily energy generation this rises to €2.80 per MWh on

an annual basis and up to €3.86 per MWh during winter. With the average cost of power on the order of €45 per MWh, these are hypothetical price savings of roughly 5%, which would certainly be of interest to industry stakeholders. For the December, January, February increased fraction case, the average cost of power, €62 per MWh, is much higher due to the increased effect of high-demand outliers, which also results in a maximum saving of roughly 5%.

Of course our demand model cannot meet these theoretical targets, even when using observational ERA5 temperatures as input. Perfect temperature forecasts for

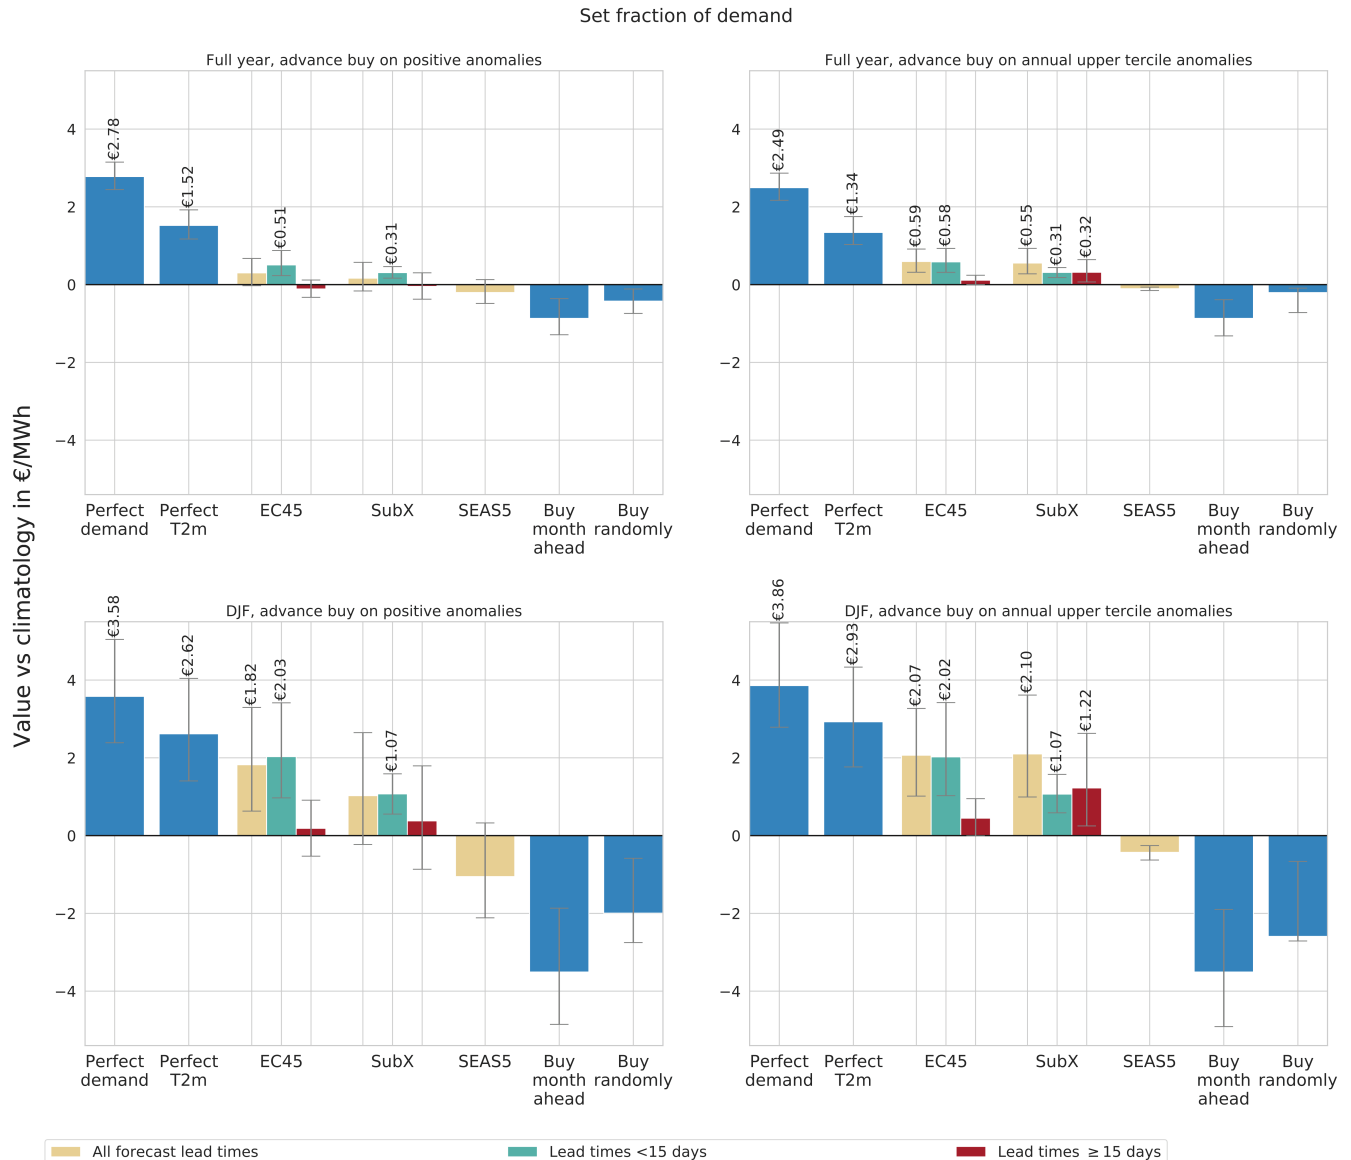



FIGURE 3 As Figure 2, but showing the average saving on a MWh of French baseline energy weighted by the amount of demand on that day. This represents the scenario where a user is interested in buying a set fraction of the daily demand. Prices are explicitly valid for a user whose average energy obligation is 1 MWh

December, January, February come closer to the theoretical maximum saving than perfect annual forecasts, realising 63–75% of the maximum saving depending on price scenario, and 43–55% respectively. This is due simply to the much better temperature–demand fit obtained for December, January, February than other seasons, while the improved value obtainable for the set fraction case is due to the increased influence of high-demand days under this scenario; as our model captures only some of the variability, we are more able to predict the sign of strong outlier events than weak ones.

Looking now at the performance of our forecast systems, we immediately see that forecasts for the full year hold much less value than during December, January,

February, which again follows quite naturally from the better temperature–demand fit for December, January, February than other seasons. However we do see statistically significant, albeit low, value on an annual basis for both sub-seasonal systems, at least in the first two weeks, for both the set amount and set fraction cases. When looking at trades on positive anomalies, longer-range forecasts are indistinguishable from or robustly worse than always buying energy the day ahead, and their presence removes value from the forecast system. This is because a forecast with no skill has negative value compared to climatological action, as shown by the negative value of the random trading strategy. However, When only the first 15 days of forecasts are used, both EC45 and SubX

FIGURE 4 The cumulative distribution of forecast lead times used to make trading decisions in Figures 2 and 3 for each forecasting system. The increased initialisation frequency of the EC45 system compared to SubX means that the average lead time of forecasts used for decision-making is lower



have statistically significant savings associated with them, which while no larger than €0.50 per MWh on an annual basis, are between €0.65 per MWh and €2.03 per MWh during December, January, February. For the set fraction case, December, January, February EC45 forecasts are valuable even when taking the system as a whole, although the longer-range forecasts individually are indistinguishable from climatology.

The first 15 days of EC45 are significantly better than for SubX for the set fraction case during December, January, February as a result of more regular forecast initialisation. We also see quite clearly that for the paradigm described here, there is no value in the SEAS5 system, despite (as we will show in later sections) it having T2m skill comparable to the sub-seasonal systems. This is purely due to the fact it is always initialised on the 1st of the month – the worst possible time from an energy perspective! This is made explicit in Figure 4 which shows the cumulative fraction of forecasts used in this section for increasing lead time, and where the shortest lead time SEAS5 forecast used is a 28 day forecast.

When basing our strategies on the more conservative upper tercile threshold, annual forecasts are more valuable, worth ~5% more with respect to a perfect demand forecast than when trading on positive anomalies. In December, January, February the picture is more complex with short-range forecasts becoming 1–2% less valuable, but with some longer-range forecasts now becoming useful. The forecasts beyond 15 days for the SubX system now have significant value on their own, in all scenarios except for the annual set amount case, and while the EC45 system does not show the same long-range skill, it is not significantly worse than the SubX system. The greater value of the first 15 days of EC45 compared to SubX is statistically significant for the same reason as before: its increased initialisation frequency. Crucially

however, in this case the presence of the longer-range forecasts substantially improves the quality of the SubX system as a whole, making it equally skilful to EC45, and during December, January, February almost doubling the associated cost saving under both user cases.

In summary, the simple demand model and trading strategy used here certainly limits the possible value temperature forecasts can provide in this framework, however it provides enough skill to show that under some conditions forecasts can have value for daily prediction at lead times exceeding two weeks. Results were optimal when trading on upper tercile anomalies, and tuning the action threshold more precisely could conceivably increase value yet further. We found forecasts gave better savings in general for December, January, February than for the full year, as a result of both higher surface temperature skill (Figure 6 below) and a closer temperature–demand relationship. We also saw that while increased initialisation frequency of course makes a forecast system more valuable, as shown by the comparison between the first two weeks of EC45 and SubX, the addition of forecasts ≥ 15 days can make up for that lower initialisation frequency at a lower computational cost. We now proceed to examine how closely these observations match up to the implications provided by more traditional, meteorological skill frameworks.

4 | POTENTIAL ECONOMIC VALUE

We have seen above that, under realistic decision-making conditions, longer-range temperature forecasts can provide value in some cases for end-users in the energy sector. We would also now like to consider how closely the value of forecasts shown above can be emulated within the more

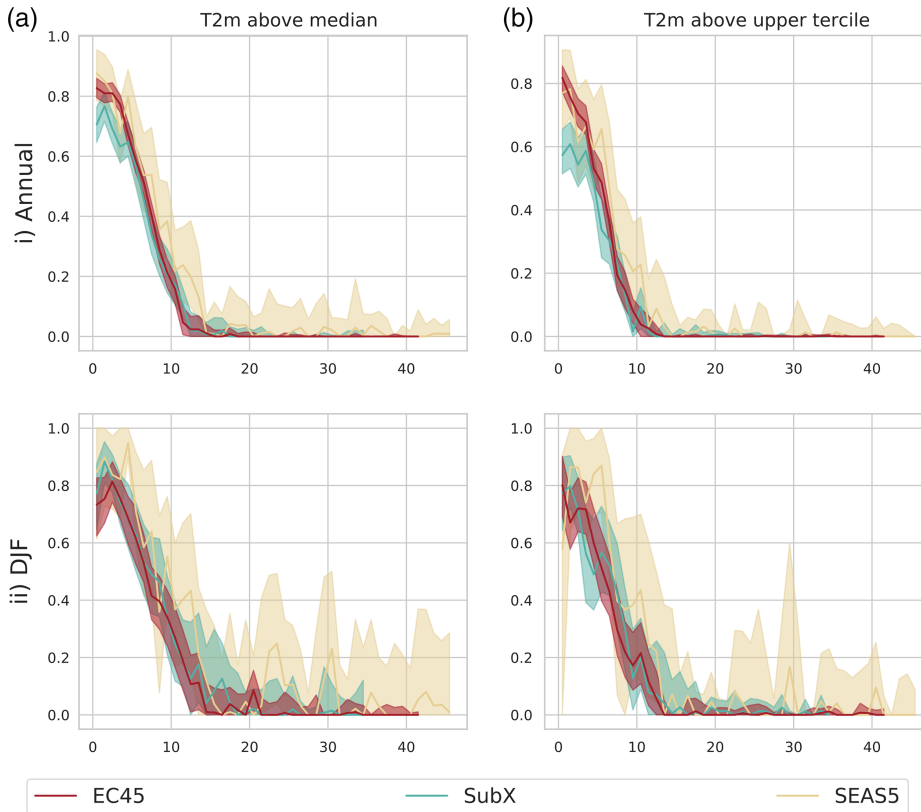


FIGURE 5 The potential economic value of forecasts for forecasts of (a, c) T2m (as a proxy for demand) exceeding the median, and (b, d) exceeding the upper tercile, annually in (a, b) and for December, January, February only in (c, d). A cost–loss ratio of 0.65 is used in all cases. Shading shows the first standard deviation estimated by bootstrap resampling of forecast days

traditional cost/loss framework of the potential economic value (PEV; Murphy, 1985; Sultan *et al.*, 2010). The main simplifications involved in moving to the PEV framework is that the cost and loss no longer depend on the magnitude of the demand continuously, only on whether the action threshold is exceeded or not, and that we switch from looking at the predictability of demand directly, and use T2m as a proxy.

The PEV can be defined in terms of the confusion matrix, $M^{(\text{conf})}$, which describes how likely we are to predict or miss an impactful event e , and the cost matrix, $M^{(\text{cost})}$ which tells us the associated cost of acting on our forecast:

$$M^{(\text{conf})} = \begin{matrix} & \begin{matrix} \text{Predicted} & \text{Not Predicted} \end{matrix} \\ \begin{matrix} \text{Event} \\ \text{No Event} \end{matrix} & \begin{pmatrix} P_{\text{True Positive}} & P_{\text{False Negative}} \\ P_{\text{False Positive}} & P_{\text{True Negative}} \end{pmatrix} \end{matrix}$$

$$M^{(\text{cost})} = \begin{matrix} & \begin{matrix} \text{Action} & \text{No Action} \end{matrix} \\ \begin{matrix} \text{Event} \\ \text{No Event} \end{matrix} & \begin{pmatrix} C & L \\ C & 0 \end{pmatrix} \end{matrix}$$

where P is probability, C is the cost of taking action (i.e., buying month-ahead energy), and L is the loss associated with inaction, in our case paying a hefty day-ahead price for power during a period of high demand. The four probabilities in $M^{(\text{conf})}$ are estimated using our demand

forecasts and observational data from the previous section. If the elements of these matrices can be defined and calculated, then the PEV is given by

$$PEV = \frac{C_{\text{clim}} - \sum_{i,j} M_{ij}^{(\text{conf})} M_{ij}^{(\text{cost})}}{C_{\text{clim}} - C \cdot P(e)}, \quad (3)$$

where

$$C_{\text{clim}} = \begin{cases} C & \text{if } C/L < P(e), \\ L \cdot P(e) & \text{otherwise.} \end{cases} \quad (4)$$

This is simply the cost saving provided by acting on the forecast, expressed as a fraction of the difference between the climatological cost, C_{clim} , and the cost unavoidable with even perfect knowledge, $C \cdot P(e)$, where $P(e)$ is the probability of the event (i.e., it ranges from 1 for a perfect forecast to 0 for a skill-less one). We are in a position to directly extract the ratio from our applied scenario above, and do this by taking the cost of action as the average amount by which month-ahead price exceeds day-ahead price, and by taking the loss as the average amount by which day-ahead cost during a period of high demand exceeds day-ahead cost during low demand. Depending on season or trading scenario chosen, we find the C/L ratio ranges from 0.6 to 0.7. We find our results are not

qualitatively sensitive to the small differences within this range, and so use an intermediate value of 0.65 for our analysis here.

Our construction of PEV in this way is similar to the approach laid out in table 12.1 of Dutton *et al.*, (2018), although we have found no explicit use of it, or estimation of a cost–loss ratio in this manner, in the energy or meteorology literature. Indeed, while there are a range of studies that have looked at the potential economic value of sub-seasonal forecasts (Zhu *et al.*, 2002; Emanuel *et al.*, 2012; Materia *et al.*, 2020), or with a more complex trinary decision model (Sultan *et al.*, 2010), results are normally presented for a range of cost–loss ratios, with no comment made on what value is most realistic for a given application. The optimal PEV will always be found for a user with a cost–loss ratio equal to the probability of the event being predicted (i.e., in our case 0.5 or 0.33 depending on action threshold). As most work looks at action on upper tercile, quintile or decile events, this skews the applicability of a forecast heavily towards low cost–loss ratios, and it is of course not *a priori* given that the optimal user of a forecast actually exists. In our case, our cost–loss ratio is quite high compared to the optimal value and we will see this reflected in relatively modest PEV values.

We find in Figure 5 that there are only small differences between PEV for December, January, February and annually, and between action thresholds, with 1-day forecasts tending to realise 70–80% of the potential value. The PEV falls in all cases below 40% between days 7–10, and becomes statistically indistinguishable from 0 between days 10 and 14. We also see immediately that the large differences between models vanish, with the exception of a slightly lower annual PEV for SubX.

Based on these results in isolation, we would expect to see no real-world value in forecasts greater than 15 days, and an average PEV of approximately 0.4 over the first two weeks of a forecast, for both action thresholds and across seasons. This matches broadly speaking with the annual values calculated in the previous section, especially for the ultimately optimal approach of trading on upper tercile events, which also showed 40% of the potential saving provided by a perfect temperature forecast being realised. However, the savings seen for extended-range forecasts are in direct contrast to the zero-value PEV and results for December, January, February, where models were able to provide upward of 60% of the value a perfect temperature forecast would provide.

Therefore we see the PEV in this case misses a valuable component of the user end-case and causes an overly pessimistic estimate of forecast skill; more extreme demand anomalies have a larger price impact. This could be partially accounted for by extending the event/non-event dichotomy to include strong events, weak events etc.,

which would give a finer-grained perspective but which also would require additional free parameters representing more detailed costs and losses for different categories of event. Despite this limitation of the PEV, we see a broad qualitative agreement between the two approaches, especially for the annual average, with the clearest value in the first two weeks of the forecast, and PEV over that period falling in the 30–60% range.

5 | NON-ECONOMIC SKILL SCORES

Having evaluated forecast value in economic terms, it is now useful to examine more purely meteorological skill scores, and we focus on some of the most frequently employed skill metrics in the literature. In Figure 6 we show the Pearson correlation and root-mean-square error (RMSE) between the forecast ensemble mean and observations, and the continuous ranked probability score (CRPS; a measure of probabilistic skill; Hersbach, 2000 gives details) for daily French surface temperatures. As for the PEV, we see relatively minor differences in skill between the different forecast systems, and most differences are confined to week 1 of the forecast.

Correlation skill remains significantly above zero out to days 22 and 27 for SubX and EC45 respectively, both during December, January, February and annually (Figure 6a, b), while it is not possible to distinguish robust skill in SEAS5 beyond day 15 in December, January, February or day 18 annually. These scores are more optimistic than any other metric, and from this we might well conclude that the daily forecast retains some marginal value into the extended range as we have realised in Section 3, although of course there is no way to quantify the term 'marginal' with a non-economic score. It is worth noting that, despite correlations being low after week 2 (<0.3), they are still significantly different from zero, and as there clearly is skill at these longer lead times shown in Section 3, they are also meaningful.

In contradiction, annual RMS Error saturates between days 11 and 15 (Figure 6c, d), which would then imply essentially no useful information existed in the extended range. The CRPS (Figure 6e, f) provides an equivalent message, saturating between days 15 and 20, and implying no added value could be gained even from a full probabilistic perspective. This can be best understood in light of the fact that our trading strategy in Section 3 is based on threshold exceedance, reducing sensitivity to the differences between model and verification that would degrade RMS error and CRPS. Within the energy framework, as long as sign and approximate magnitude of anomalies are well predicted, value can be extracted.

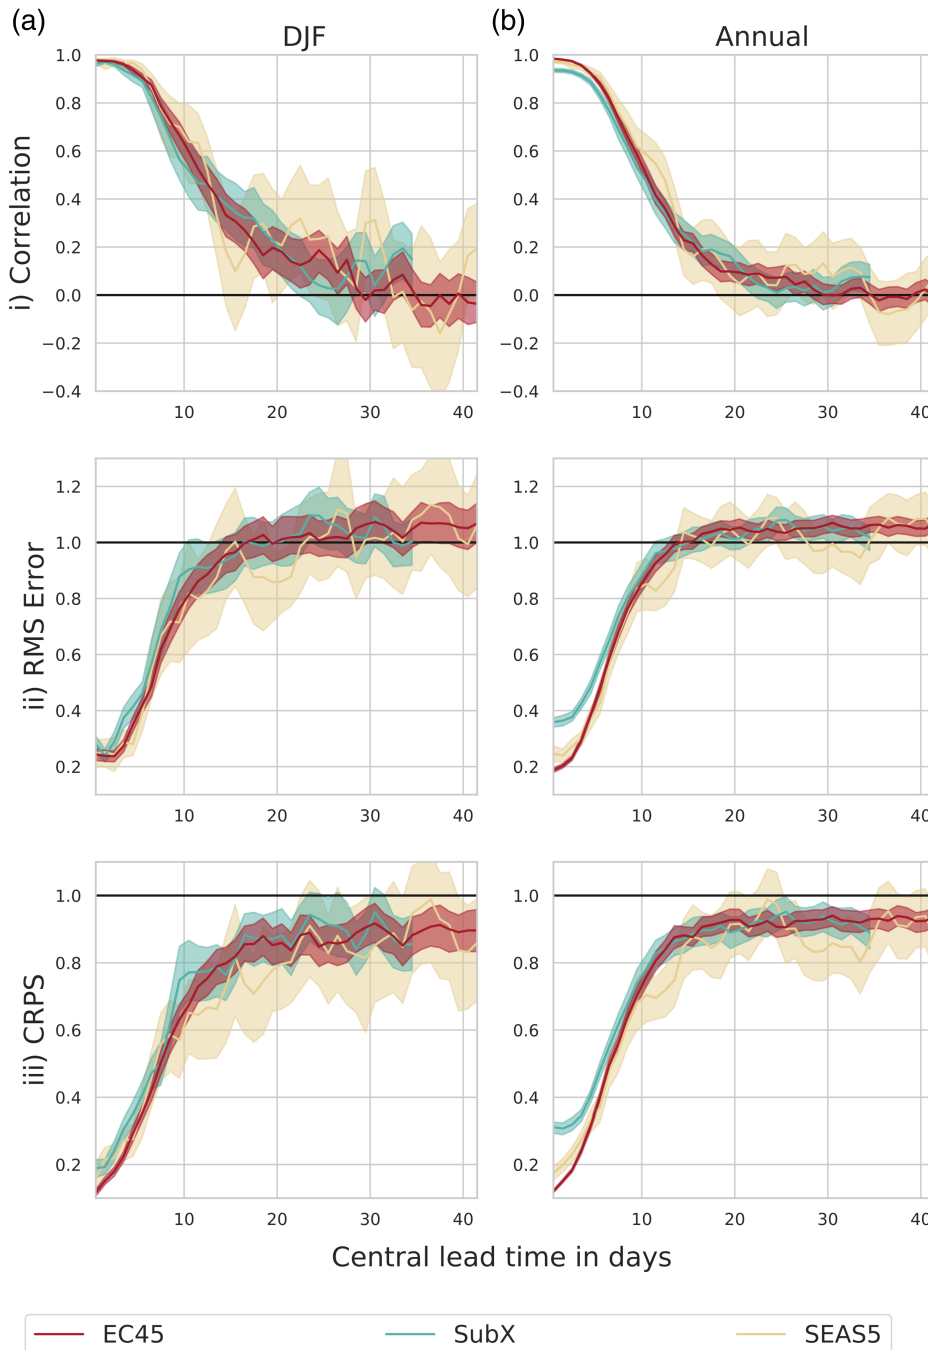


FIGURE 6 Measures of ensemble forecast skill for daily French 2 m temperature anomalies (a, c, e) during December, January, February and (b, d, f) annually averaged. (a, b) Correlations and (c, d) RMS error are with respect to the ensemble mean, whereas (e, f) the CRPS is a fully probabilistic score. RMSE and CRPS have been normalised by the seasonal climatological variance estimated from ERA5. Shaded regions represent the first standard deviation of bootstrap resampling over forecast initialisation dates

6 | DISCUSSION

In Section 3 we have quantified the real-world value of sub-seasonal temperature forecasts for energy trading on a national market. Using a simple trading strategy, we show that, compared to a climatological strategy, sub-seasonal forecast systems can provide savings of €0.30–2.10 per MWh, which represents 40% of the value a perfect temperature forecast could provide annually, and up to ~70% during December, January, February. Due to the simplicity of our model, the largest added value is for December, January, February when demand is most sensitive to

temperature, but we also find some significant value on an annual basis. Not all this value comes from the first two weeks of the forecast systems; we show that SubX forecasts beyond two weeks have value on their own, 40% as much as a perfect forecast during December, January, February, and that, without these extended-range forecasts, the total value of the system is reduced. We also saw clearly in this framework how differences in forecast initialisation frequency had a substantial impact on the usefulness of the system, with SEAS5 showing no value at all, despite being a very similar model to the EC45 system.

Having found these results, in Sections 4 and 5 we explore how well they are mirrored by more idealised/abstracted skill scores. In Section 4 we have managed to extract a cost–loss ratio from our data, and compare the true forecast value to the estimate of a simplified PEV model. We find broad qualitative agreement with Section 3, especially for shorted lead times, but the binary cost–loss assumption underlying the PEV causes an underestimation of forecast value in cases where cost correlates strongly with the magnitude of the weather event (i.e., the binary assumption of event/no event is simplistic). In our case this means PEV is indistinguishable after week 2, and the value in December, January, February is underestimated. Additionally if we had averaged over different cost–loss ratios, or assumed an optimal value as is commonly done, we may have arrived at quite different results, as the cost–loss ratio we calculated, 0.6–0.7, is much higher than for the optimal value. Simplifying our skill scores yet further, a comparison to conventional, non-economic scores highlights that, even when RMS error and CRPS are at climatological levels, end-users might still be able to extract value from forecasts, even with very low positive correlations.

As we have emphasised, our estimations of value using our trading framework form only a lower bound. A more careful consideration of action thresholds, use of population-weighted area means, and detailed analysis of very extreme events which affect energy supply as well as demand (as mentioned in the Introduction) could increase the forecast value further, making the underestimation of value from traditional skill scores all the more severe.

We acknowledge that we have considered in this work only a single end-user sector in a single country, and that, even within this domain, other frameworks for making forecast-based decisions exist. We present these results then not in order to draw far-reaching conclusions about current operational forecast value, but to provide an argument and conceptual model for assessing sub-seasonal forecasts in a way that directly connects with end-user requirements.

Extending the methodology of this study to additional national domains, and in general trying to use specific end-user examples to choose variables for forecast skill verification, could help shape model development and increase enthusiasm amongst forecast users. As discussed in the Introduction, there is a growing body of literature presenting user-driven assessments of model value, with greater or lesser degrees of complexity, whose methodologies could potentially be used more routinely to inform model development. Of course there are many benefits to simple skill scores: they are easy to understand and calculate, and have general applicability, but much as the importance of a hierarchy of models is recognised as


vital to atmospheric science, adopting a hierarchy of skill scores, including intermediate complexity options as we present here, may prove to be of equal use in converting science into societal benefit.

ACKNOWLEDGEMENTS


We acknowledge the National Environmental Research Council for funding of an industrial internship on which the origins of this work are based, ECMWF for making IFS hindcast data available by commercial license, and the agencies that support the open-source SubX system: NOAA/MAPP, ONR, NASA, and NOAA/NWS.

We thank Dave Macleod and Daniel Befort for valuable discussions that helped shape the direction and enhanced the overall understanding of this work. and the reviewers for their constructive comments which have been of great help in improving this manuscript.

ORCID

Joshua Dorrington  <https://orcid.org/0000-0003-3802-457X>

Tim Palmer  <https://orcid.org/0000-0002-7121-2196>

Antje Weisheimer  <https://orcid.org/0000-0002-7231-6974>

REFERENCES

- Buizza, R., Houtekamer, P.L., Pellerin, G., Toth, Z., Zhu, Y. and Wei, M. (2005) A comparison of the ECMWF, MSC, and NCEP global ensemble prediction systems. *Monthly Weather Review*, 133, 1076–1097. <https://doi.org/10.1175/MWR2905.1>
- Chevallier, M., Goessling, H., Guémas, V., Jung, T. and Massonnet, F. (2018). The role of sea ice in subseasonal predictability, Chapter 10 in Sub-Seasonal to Seasonal Prediction. Robertson, A., Vitart, F. (eds), Elsevier, Amsterdam, Netherlands.
- Cho, H., Goude, Y., Brossat, X. and Yao, Q. (2013) Modeling and forecasting daily electricity load curves: A hybrid approach. *Journal of the American Statistical Association*, 108, 7–21. <https://doi.org/10.1080/01621459.2012.722900>
- C3S (2017). ERA5: Fifth generation of ECMWF atmospheric reanalyses of the global climate. <https://cds.climate.copernicus.eu/cdsapp#!/home>; accessed 10 July 2020.
- Dirmeyer, P.A., Gentile, P., Ek, M.B. and Balsamo, G. (2019). Land surface processes relevant to sub-seasonal to seasonal (S2S) prediction, Chapter 8 in Sub-Seasonal to Seasonal Prediction. Robertson, A., Vitart, F. (eds), Elsevier, Amsterdam, Netherlands.
- Dutton, J.A., James, R.P. and Ross, J.D. (2013) Calibration and combination of dynamical seasonal forecasts to enhance the value of predicted probabilities for managing risk. *Climate Dynamics*, 40, 3089–3105. <http://www.worldclimate>
- Dutton, J.A., James, R.P. and Ross, J.D. (2018). Probabilistic forecasts for energy: weeks to a century or more, Chapter 12 in Weather & Climate Services for the Energy Industry. Troccoli, A. (ed.), Springer, Cham, Switzerland.
- EEX (2020) *Power Futures Overview*. <http://www.eex.com/en/products/power-derivatives-market/power-futures/power-futures-products>; accessed 10 July 2020.

- Emanuel, K.A., Fondriest, F. and Kossin, J. (2012) Potential economic value of seasonal hurricane forecasts. *Weather Climate and Society*, 4, 110–117. <https://doi.org/10.1175/WCAS-D-11-00017.1>
- ENTSOG (2017). Winter review 2016/2017. <https://www.entsog.eu/outlooks-reviews#winter-outlooks-and-reviews>; accessed 10 July 2020.
- Hardt, L., Brockway, P., Taylor, P., Barrett, J., Gross, R. and Heptonstall, P. (2019). Modelling demand-side energy policies for climate change mitigation in the UK: a rapid evidence assessment. <http://www.ukerc.ac.uk/publications/modelling-demand-in-energy-system-models.html>; accessed 10 July 2020.
- Hersbach, H. (2000) Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather and Forecasting*, 15, 559–570
- Hoskins, B.J. (2013) The potential for skill across the range of the seamless weather–climate prediction problem: a stimulus for our science. *Quarterly Journal of the Royal Meteorological Society*, 139, 573–584
- Johnson, S.J., Stockdale, T.N., Ferranti, L., Balmaseda, M.A., Molteni, F., Magnusson, L., Tietsche, S., Decremet, D., Weisheimer, A., Balsamo, G., Keeley, S.P.E., Mogensen, K., Zuo, H. and Monge-Sanz, B.M. (2019) SEAS5: the new ECMWF seasonal forecast system. *Geoscientific Model Development*, 12, 1087–1117. <https://doi.org/10.5194/gmd-12-1087-2019>
- Lam, J.C., Wan, K.K., Cheung, K.L. and Yang, L. (2008) Principal component analysis of electricity use in office buildings. *Energy and Buildings*, 40, 828–836
- Lynch, K.J., Brayshaw, D.J. and Charlton-Perez, A. (2014) Verification of European subseasonal wind speed forecasts. *Monthly Weather Review*, 142, 2978–2990. <https://doi.org/10.1175/MWR-D-13-00341.1>
- Materia, S., Muñoz, Á.G., Álvarez-Castro, M.C., Mason, S.J., Vitart, F. and Gualdi, S. (2020) Multimodel subseasonal forecasts of spring cold spells: potential value for the hazelnut agribusiness. *Weather and Forecasting*, 35, 237–254. <https://doi.org/10.1175/WAF-D-19-0086.1>
- Monhart, S., Spirig, C., Bhend, J., Bogner, K., Schär, C. and Liniger, M.A. (2018) Skill of subseasonal forecasts in Europe: effect of bias correction and downscaling using surface observations. *Journal of Geophysical Research: Atmospheres*, 123, 7999–8016. <https://doi.org/10.1029/2017JD027923>
- Murphy, A.H. (1985) Decision making and the value of forecasts in a generalized model of the cost-loss ratio situation. *Monthly Weather Review*, 113, 362–369
- NASEM (2016). Enhancing the value and benefits of S2S forecasts, Chapter 3 in Next Generation Earth System Prediction: Strategies for Subseasonal to Seasonal Forecasts. <http://www.nap.edu/21873>, accessed 10 July 2020.
- Orlov, A., Sillmann, J. and Vigo, I. (2020) Better seasonal forecasts for the renewable energy industry. *Nature Energy*, 5, 108–110. <https://doi.org/10.1038/s41560-020-0561-5>
- Palmer, T.N. (2002) The economic value of ensemble forecasts as a tool for risk assessment: from days to decades. *Quarterly Journal of the Royal Meteorological Society*, 128, 747–774. <https://doi.org/10.1256/0035900021643593>
- Psiloglou, B.E., Giannakopoulos, C., Majithia, S. and Petrakis, M. (2009) Factors affecting electricity demand in Athens, Greece and London, UK: a comparative assessment. *Energy*, 34, 1855–1863
- RTE (2019) *Record of power consumption*. http://clients.rte-france.com/lang/an/visiteurs/vie/vie_stats_conso_inst.jsp; accessed 10 July 2020.
- Soret, A., Torralba, V., Cortesi, N., Christel, I., Palma, L., Manrique-Suñén, A., Lledó, L., González-Reviriego, N. and Doblas-Reyes, F.J. (2019) Sub-seasonal to seasonal climate predictions for wind energy forecasting. *Journal of Physics: Conference Series*, 1222. <https://doi.org/10.1088/1742-6596/1222/1/012009>
- Stockdale, T.N., Balmaseda, M.A., Johnson, S.J., Ferranti, L., Molteni, F., Magnusson, L., Tietsche, S., Vitart, F., Decremet, D., Weisheimer, A., Roberts, C.D., Balsamo, G., Keeley, S.P.E., Mogensen, K., Zuo, H., Mayer, M. and Monge-Sanz, B.M. (2018) SEAS5 and the future evolution of the long-range forecast system. <https://www.ecmwf.int/node/18750>; accessed 10 July 2020.
- SubX (2019). The SubX model dataset. <http://iridl.ldeo.columbia.edu/SOURCES/.Models/.SubX/>; accessed 10 July 2020.
- Sultan, B., Barbier, B., Fortilus, J., Mbaye, S.M. and Leclerc, G. (2010) Estimating the potential economic value of seasonal forecasts in West Africa: a long-term ex-ante assessment in Senegal. *Weather Climate and Society*, 2, 69–87. <https://doi.org/10.1175/2009WCAS1022.1>
- Thornton, H.E., Scaife, A.A. and Hoskins, B.J. (2016) The role of temperature in the variability and extremes of electricity and gas demand in Great Britain. *Environmental Research Letters*, 11. <https://doi.org/10.1088/1748-9326/11/11/114015>
- Van Der Wiel, K., Bloomfield, H.C., Lee, R.W., Stoop, L.P., Blackport, R., Screen, J.A. and Selten, F.M. (2019) The influence of weather regimes on European renewable energy production and demand. *Environmental Research Letters*, 14. <https://doi.org/10.1088/1748-9326/ab38d3>
- Vitart, F. (2017) Madden–Julian Oscillation prediction and teleconnections in the S2S database. *Quarterly Journal of the Royal Meteorological Society*, 143, 2210–2220. <https://doi.org/10.1002/qj.3079>
- White, C.J., Carlsen, H., Robertson, A.W., Klein, R.J., Lazo, J.K., Kumar, A., Vitart, F., Coughlan de Perez, E., Ray, A.J., Murray, V., Bharwani, S., MacLeod, D., James, R., Fleming, L., Morse, A.P., Eggen, B., Graham, R., Kjellström, E., Becker, E., Pegion, K.V., Holbrook, N.J., McEvoy, D., Depledge, M., Perkins-Kirkpatrick, S., Brown, T.J., Street, R., Jones, L., Remenyi, T.A., Hodgson-Johnston, I., Buontempo, C., Lamb, R., Meinke, H., Arheimer, B. and Zebiak, S.E. (2017) Potential applications of subseasonal-to-seasonal (S2S) predictions. *Meteorological Applications*, 24, 315–325. <https://doi.org/10.1002/met.1654>
- Zhu, Y., Toth, Z., Wobus, R., Richardson, D.S. and Mylne, K. (2002) The economic value of ensemble-based weather forecasts. *Bulletin of the American Meteorological Society*, 83, 73–83

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Dorrington J, Finney I, Palmer T, Weisheimer A. Beyond skill scores: exploring sub-seasonal forecast value through a case-study of French month-ahead energy prediction. *Q J R Meteorol Soc.* 2020;146:3623–3637. <https://doi.org/10.1002/qj.3863>

APPENDIX

In this appendix, the details of the fitted relationships between ERA5 surface temperature and demand, and between demand and energy price for the period 2010–2018 are shown in Figures A1 and A2 respectively, and briefly discussed.

A.1 Fitting temperature to energy demand

In order to fit a relation between surface temperature and demand, we first need to remove a seasonal cycle from both time series. Different deseasonalising methods were trialled, and removing the calendar monthly mean from each daily data series was found to give the best fit. The upper row of Figure A1 shows the relationship of these temperature anomalies against demand anomalies for each season, with a quadratic least-squares fit shown by the blue line. Polynomial fits up to fifth order were trialled but gave no significant improvements. These quadratic fits have been used to produce ERA5 estimates of demand anomalies, as shown on the lower row of Figure A1. As we see, demand anomalies are most sensitive to temperature anomalies during December, January, February, with JJA showing a particularly weak relationship. The bimodality in JJA demand represents a week-day/weekend split, presumably due to increased commercial air-conditioning during the summer months. Remembering that the proportion of demand variance that the temperature anomaly explains is the correlation squared,

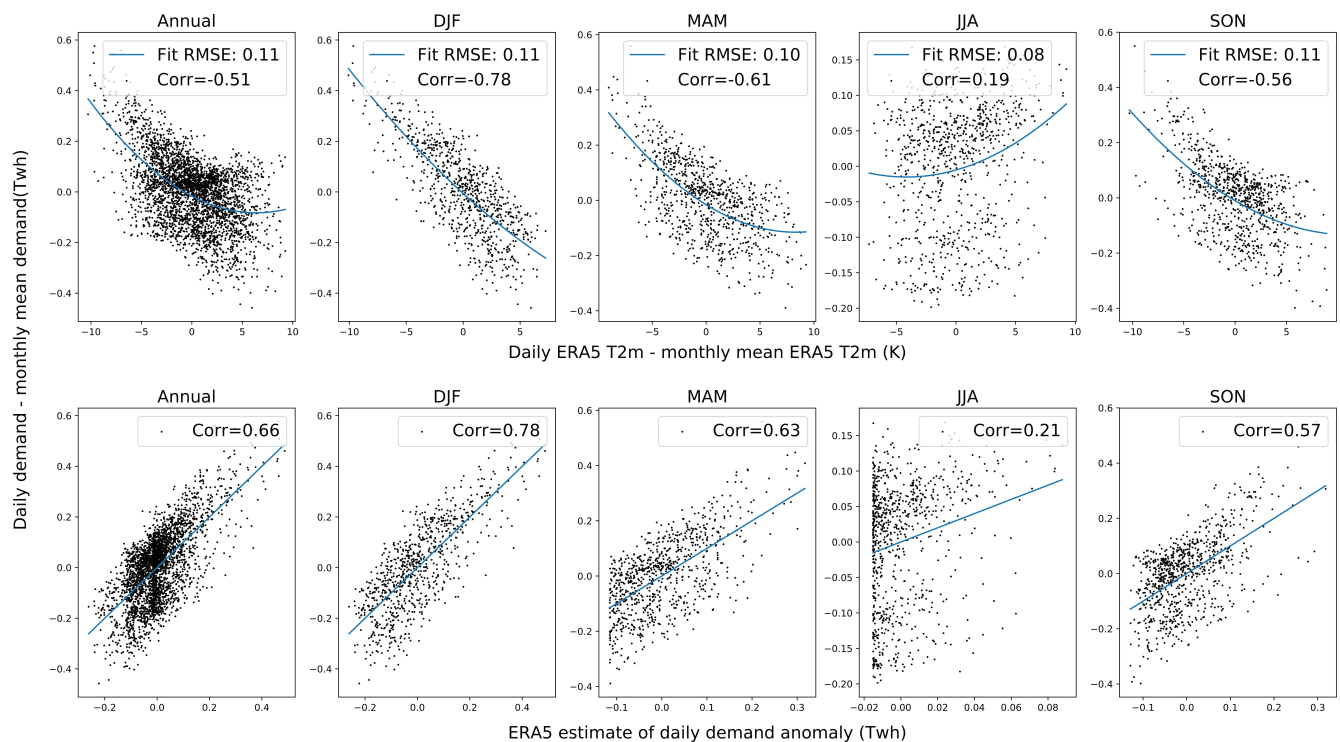


FIGURE A1 Upper row: anomalies of daily ERA5 surface temperature against anomalies of daily energy demand, with anomalies calculated with respect to the mean value for each calendar month. The solid line shows the best quadratic fit to the data. Lower row: ERA5 estimates of daily energy demand anomaly, calculated using the quadratic fit above, plotted against the same energy demand anomaly series

we see why we were only able to obtain forecast value during December, January, February and on an annual basis; in spring to autumn, temperature explains less than 40% of the demand variability.

A.2 Relationship between energy demand and potential price saving

The upper row of Figure A2 shows the correlation between daily French energy demand and the month-ahead price saving, which we define as the difference between the day-ahead price, and the final

month-ahead price. The lower row shows the same but using daily demand anomalies, calculated with respect to the mean demand for each calendar month. We see that deseasonalising most improves the fit during the transition months (i.e., MAM and SON) when the weather and consumers' behavioural habits are changing most rapidly. We see that after deseasonalising, the relationship is relatively stable across seasons, with the correlation of the anomaly time series varying between 0.61 and 0.72 (or equivalently, explaining between 37% and 52% of the variance in price differences).

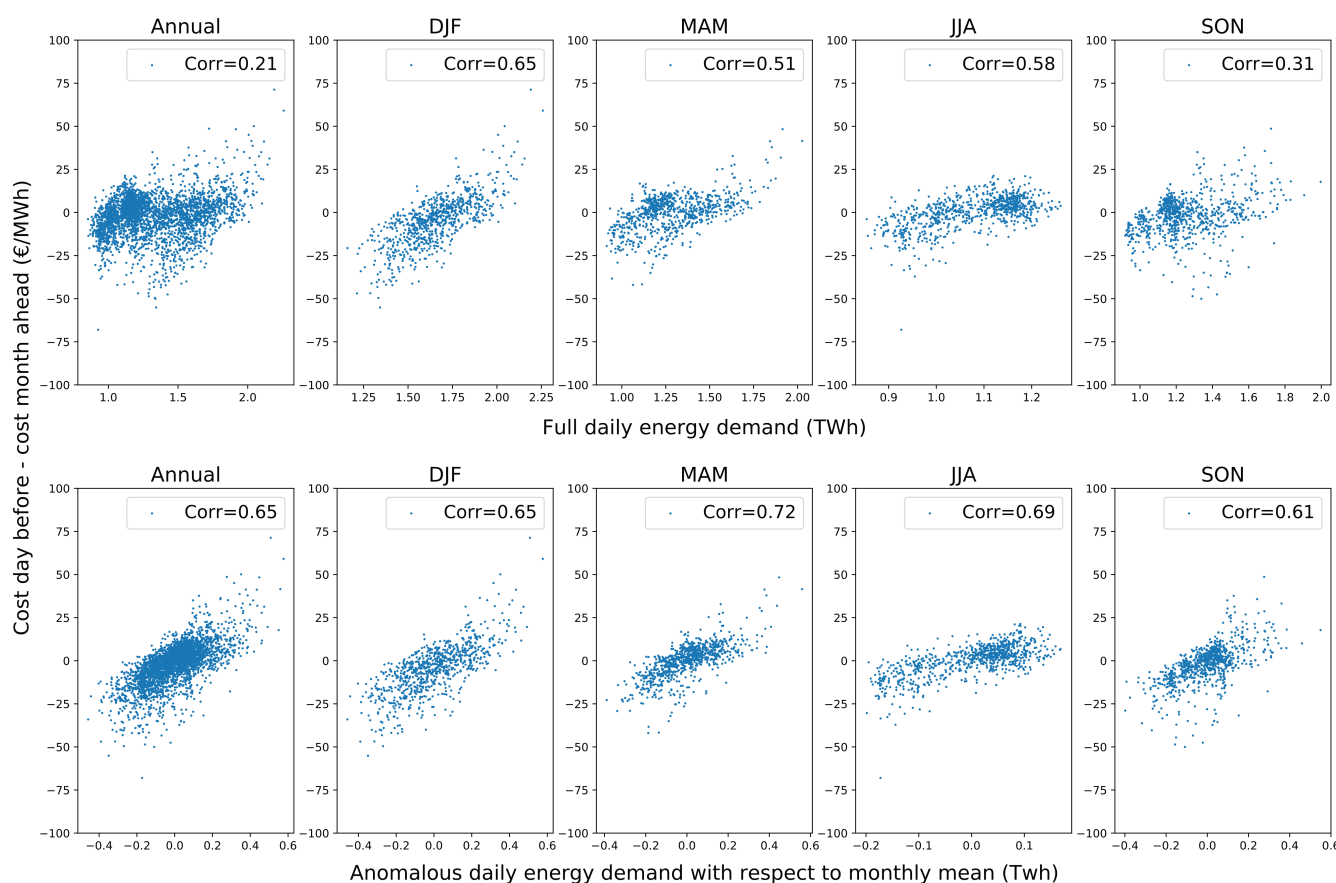


FIGURE A2 Upper row: daily French energy demand against the difference in price between buying energy the day before delivery and buying power at the end of the previous month, termed the month-ahead cost saving. Lower row: the daily demand anomaly with respect to the mean value for each calendar month against the same month-ahead cost saving as above