Title

The cost of day-ahead solar forecasting errors in the United States

Authors
Wang, Yuhan
Millstein, Dev
Mills, Andrew D
et al.

Publication Date
2022

DOI
10.1016/j.solener.2021.12.012

Solar Energy 231 (2022) 846–856

Contents lists available at ScienceDirect

Solar Energy

journal homepage: www.elsevier.com/locate/solener

The cost of day-ahead solar forecasting errors in the United States

Yuhan Wang [a,b], Dev Millstein [a,*], Andrew D. Mills [a], Seongeun Jeong [a], Amos Ancell [a]

[a] *Lawrence Berkeley National Laboratory, United States*
[b] *University of California, Berkeley, United States*

A B S T R A C T

As solar energy contributes an increasing share of total electricity generation, solar forecasting errors become important relative to overall load uncertainty and can add costs to electricity systems. We investigated the costs of day-ahead solar forecast errors across 667 existing solar power plants in the United States (years 2012 through 2019). Our analysis was based on hourly real-time and day-ahead nodal prices. We analyzed two types of solar forecasts: persistence forecasts, a simple approach to forecasting, and a numerical weather prediction forecast, the North American Mesoscale Model (NAM), an improvement over persistence forecasts based on public data and modelling software. We modeled hourly energy forecasts using meteorological forecasts and plant specific characteristics. Hourly plant generation was modeled and debiased with multiple sources of generation records. NAM forecast errors had relatively low costs on average, at no more than \$1/MWh in all years except 2016, when costs rose to \$1.5/MWh. Even after these error costs, the value of solar was marginally higher when simulating solar participation in day-ahead markets versus participation only in real-time markets. On average, the premium for participating in the day-ahead market, based on NAM forecasts, ranged from −0.5 to 5.2 \$/MWh across years. Average error costs were higher in regions with higher solar penetration (i.e., California and New England) compared to regions with low solar penetration. However, California and New England had similar error costs despite higher solar penetration in California, indicating that error costs to date have been only loosely correlated with solar penetration levels.

1. Introduction

The share of electricity generated from solar power is growing in the United States and globally – solar power now accounts for roughly 20% of the electricity generated in California (Bolinger, Seel et al. 2020) and global solar generation more than doubled between 2015 and 2018 (EIA 2021). That solar is a major source of electricity in some regions, and growing quickly in many others, means that solar forecast errors are, or may soon become, an important source of uncertainty to balancing electricity supply and demand. The growth of solar motivates questions about the impact of solar forecasting errors on regional electricity systems. For example, what is the cost of forecast errors to the overall system, or the cost to plant operators? Can improved forecasts reduce costs? Understanding the answers to these questions may help smooth the integration of solar energy into the electricity grid, reduce the cost of electricity, and support decarbonization goals (as solar power is often considered a key component of pathway's for electric sector decarbonization).

Solar forecast errors, or any unanticipated change to load or supply, can cause real costs to energy systems because of inflexibility in demand and supply. For example, prior research suggests that improving short-term load forecasting accuracy by 1% would provide approximately \$1.5 million of value per year to a typical U.S. utility with 5-gigawatts of peak load (Hong 2015). Short-term electricity demand is commonly inelastic, or almost inelastic, to real-time (RT) prices. Suppliers (types of generators) face varying physical limits to the rate at which they can turn on, off, ramp up, or ramp down. Because of this inflexibility in both demand and supply, a large over forecast in solar supply will likely force a different set of generators to react than the set that was originally scheduled. This new set of generators may be more dependent on less efficient, but faster reacting, generators, such as gas turbines. An under forecast of solar can cause transmission congestion, curtailment, and also force a change to the planned set of generators. Past research efforts have explored the impacts of solar forecast errors with a variety of techniques, for example, through econometric analyses of the German market (Gürtler and Paulsen 2018, Kulakov and Ziel 2021) and the California market (Woo, Moore et al. 2016), and through a simulation of solar forecasting impacts in Arizona (Wu, Botterud et al. 2015) and the New England region of the U.S. (Martinez-Anido, Botor et al. 2016).

Other studies have investigated the costs of solar forecast errors, or the value of improving solar forecasts, at one, or a small group of, solar power plants (Kraas, Schroedter-Homscheidt et al. 2013, Luoma, Mathiesen et al. 2014, De Giorgi, Congedo et al. 2015, Ruhnau, Hennig et al. 2015, Kaur, Nonnenmacher et al. 2016, Antonanzas, Pozo-Vázquez et al. 2017, Cirés, Marcos et al. 2019). These studies universally found some benefit to improving forecasts. Antonanzas et al. (2017), for

* Corresponding author.
E-mail address: dmillstein@lbl.gov (D. Millstein).

https://doi.org/10.1016/j.solener.2021.12.012
Received 3 March 2021; Received in revised form 7 December 2021; Accepted 10 December 2021
Available online 20 December 2021

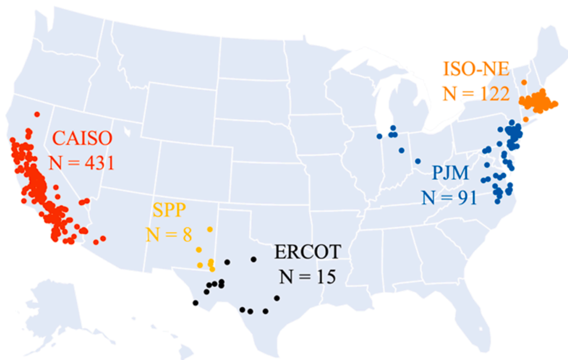

Fig. 1. The geographic distribution of plants included in this study: 667 plants span 5 U.S. ISOs.

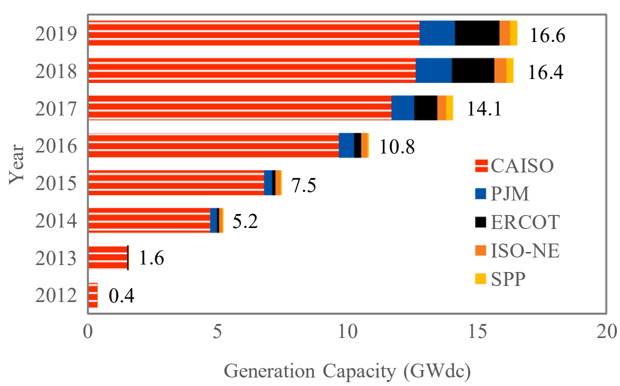

Fig. 2. UPV generation capacity included in this study by year.

example, showed that the best forecasting approach they tested provided a value of 1.2 €/MWh over a simple 'persistence' forecast, for a single plant in Spain in 2010. Kraas et al. (2013) examined the forecasts for a concentrating solar plant in Spain, and found that forecast improvements beyond persistence forecasts offer a roughly 50% reduction to penalties required by the system operator due to forecast errors. Luoma et al. (2014) examined the value of solar forecasts at 63 locations in California in 2010 and 2011. They found that forecasts from a numerical weather prediction model provided 96% of the benefits of having perfect foresight (on average, based on hourly nodal prices). Overall though, it is challenging to meaningfully compare the results across the studies because the studies cover different time periods, employ different methodology, and explore regions with varying compositions of power plant type and varying regulations relevant to the treatment of solar forecast errors.

Existing research leaves important questions unanswered. For example, how has the cost of solar forecasting errors evolved over time, or how has the cost of solar forecast errors varied by region? Existing empirical studies have focused on a limited time period (usually a year) and a single region. Thus, they provide limited insight into how the cost of forecasting errors has varied between years, evolved with additional solar penetration, or varied across regions. Additionally, most of the empirical studies cited explored the cost of forecast errors at one to ten total plants, perhaps obscuring variation within regions.

This study aims to address these knowledge gaps. We investigate the cost of solar forecasting errors at 667 existing utility-scale photovoltaic (UPV) plants across five Independent System Operator (ISOs) regions in the United States. Our study period runs from 2012 through 2019.

We estimate the cost of solar forecast errors using empirical prices from the RT and day-ahead (DA) energy markets in each ISO. Hourly price time-series for each solar plant were derived from locational marginal price (LMP) nodes closest to each plant (there are over 50,000 LMP nodes in the United States from which to match to solar plants). With this scope we examine how the cost of forecasts errors has changed overtime and across regions, and examine the results in the context of the rapid solar deployment observed over the past decade.

We compare two solar forecasting strategies, simple persistence (e.g., the solar forecast for 1 pm today = solar production at 1 pm yesterday), and a numerical weather prediction forecast (NAM, the North American Mesoscale model). The NAM forecast allows us to evaluate improvements to persistence forecasts without requiring proprietary data. RT generation is modeled based on individual plant characteristics and finely resolved meteorological data. Importantly, RT generation has been debiased with regional hourly generation records and monthly plant-level generation records, in order to provide an accurate representation of actual solar generation.

In the Methodology section (2) we provide details on the models and empirical data used in the study, as well as additional contextual discussion of RT and DA market prices and various forecasting techniques. In the Results and Discussion section (3) we first present national average results and then focus on results at the regional-level and at individual plants. We discuss how error costs are related to overall solar penetration in different regions and examine the value to solar of participating in the DA and RT markets versus participating only in RT markets. Finally, section (4) concludes the paper with a summary of the most important findings.

2. Methodology

2.1. Sample of utility-scale solar plants

The cost of forecast errors was evaluated at plants across U.S. regions with regulated wholesale electricity markets. We excluded the regions of the Midwest Independent System Operator and of the New York Independent System Operator, because at the time of writing, we were unable to debias generation records from those regions due to data limitations. For similar reasons, we have excluded plants from other regions, including states such as Nevada and Arizona. Plants were excluded if they began operating after 2018, if they were less than 1 MW in capacity, or if we were missing data needed for debiasing their historical generation records. After these exclusions, 667 utility-scale photovoltaic (UPV) plants remained and were included in our analysis. We derived plant characteristics from EIA Form 860 (Energy Information Administration 2020a) and on the data set associated with Bolinger et al. (2020). The plants are located in five Independent System Operators (ISOs), or in some cases Regional Transmission Organizations, including the California ISO (CAISO), the New England ISO (ISO-NE), the Electric Reliability Council of Texas (ERCOT), the Southwest Power Pool (SPP), and PJM. Fig. 1 plots the locations of these 667 sample plants, colored by corresponding ISOs.

Among the 667 UPV plants, 431 plants are in CAISO, together with 122 in ISO-NE, 91 in PJM, 15 in ERCOT, and 8 in SPP. Our study period covered years 2012–2019. For plants completed after 2012, we began analysis when they commenced operation, as reported by the U.S. Energy Information Administration (EIA). Some plants were excluded from year 2019 due to missing data, so 617 out of 667 plants were included in 2019 analysis. Fig. 2 shows the generation capacity included in our analysis each year.

In 2018, our sample plants accounted for 16.4 GW_{dc} generation capacity in total, 44% of cumulative U.S. UPV capacity in that year (Bolinger, Seel et al. 2020). The majority of our sample resides in the CAISO market (76.9%, by capacity), followed by ERCOT (10.0%), PJM (8.9%), ISO-NE (2.6%) and SPP (1.6%). While the national-level results from our analysis is therefore largely driven by market conditions in CAISO, ISO-specific analysis and discussion will also be presented.

2.2. Measuring the cost of forecast errors

We define the cost of a solar forecast error by the reduction in revenue caused by the error, assuming that the plant uses forecasts to participate in the DA market. For example, if a plant forecasts it will generate 100 MWh in a particular hour the next day, we assume it will sell that amount in the DA market at the DA price. If the plant then generates only 90 MWh in that hour, it will be forced to purchase 10 MWh (the amount that was under forecast) in the RT market at the RT price. If the DA and RT price are the same, the solar plant will be no worse off than had it correctly forecast 90 MWh in that hour. But if the RT price is higher than the DA price, the solar plant will have to buy the 10 MWh in RT at a higher price than it sold its generation for in the DA market, leading to an incremental cost. The incremental cost is, in this case, the cost of the forecast error. Note that we only investigate DA forecast errors, and do not evaluate the costs of forecast errors on intraday time periods. This is similar to the method applied by Luoma et al (2014) and Hong (2015).

More generally, the revenue for any plant participating in a typical two-settlement DA and RT market (R_{dart}) is equal to:

$$R_{dart} = \frac{\sum_{i=1}^{N}\left[P_{da}\bullet E_f + P_{rt}\bullet\left(E_a - E_f\right)\right]}{\sum_{i=1}^{N} E_a} \tag{1}$$

where P_{rt} is the RT market price, P_{da} is the DA market price, E_f is the energy generation forecast, E_a is the actual generation, and N is the number of hours in the year of interest.

Eq. (1) can be reformulated as:

$$R_{dart} = \frac{\sum_{i=1}^{N} P_{da}\bullet E_a}{\sum_{i=1}^{N} E_a} - \frac{\sum_{i=1}^{N}(P_{rt} - P_{da})\bullet\epsilon}{\sum_{i=1}^{N} E_a} \tag{2}$$

where $\epsilon = E_f - E_a$. In Eq. (2), the second term equals the cost of forecast errors, and this term goes to zero when the error is zero, leaving the 1st term as the full R_{dart} assuming there is no forecast error (i.e., $E_a = E_f$). Note that the denominator in Eq. (1) and (2) is equal to the sum of recorded generation, so that the units of R_{dart} are \$/MWh.

An over forecast of solar (ε greater than 0) accompanied by a higher RT price, relative to DA ($P_{rt} > P_{da}$), will see a reduction in R_{dart} relative to an accurate forecast. Similarly, an under forecast accompanied by an RT price below the DA price will reduce R_{dart}. Note there are two situations in which forecast errors can increase R_{dart}, when an under forecast of solar (ϵ less than 0) is accompanied by a RT price larger than DA price, or when those two conditions are reversed. In this situation, the error in the forecast leads to additional revenue, rather than a cost. We would expect forecast errors to result in additional revenue (rather than cost) for a portion of hours each year due to random chance, but we would expect the frequency of this outcome to decline in situations where deviations of RT prices from DA prices are more correlated with solar forecast errors. The majority of the paper will focus on the cost of solar forecast errors (defined as the second term of Eq. (2)).

2.3. Benefit of participating in DA markets

In many studies, the value of solar is estimated based on the RT price of energy at the time at which the solar generation is realized. Focusing on RT energy value ignores the potential value of participating in the DA market.

To determine if solar plants would find it valuable to use forecasting to participate in DA markets, we calculate the Day-Ahead Premium (DAP). The DAP is the difference between revenues from participation in both the DA and RT market, R_{dart}, and revenues from selling solar energy only in RT market, R_{rt}. Note that the DAP depends not just on forecast accuracy but also any systematic differences in the P_{da} and P_{rt}.

We define revenue on a per unit basis such that R_{dart}, R_{rt} and DAP are all in the \$/MWh-generated. Eq. (3) and (4) describe the calculation of DAP, which also build on Eq. (1) for R_{dart}.

$$R_{rt} = \frac{\sum_{i=1}^{N} P_{rt}\bullet E_a}{\sum_{i=1}^{N} E_a} \tag{3}$$

$$DAP = R_{dart} - R_{rt} \tag{4}$$

Below we note a few details related to the calculation of DAP. Although prices are positive in most situations, RT and DA wholesale market prices drop below zero on occasion, which is a signal of oversupply in energy markets. To account for that, it is assumed that negative RT market prices will lead to economic curtailment, i.e. $E_a = 0$ if $P_{rt} < 0$. However, the total energy generation in the denominators (not numerators) in Equations (3) and (4) above is based on total potential generation without curtailment. This is to avoid an increase in calculated value because of curtailment. In other words, R_{rt} represents total revenue after curtailment at prices below zero divided by the total generation possible (ignoring curtailment) given weather conditions and plant characteristics. A similar logic stands for R_{dart}.

Additionally, the pricing in ERCOT markets include a specific mechanism named “Operating Reserve Demand Curve” (ORDC). The ORDC mechanism began in 2014. In addition to the standard pricing system (Locational Marginal Pricing, LMP), ERCOT’s ORDC is an administrative process to increase the raw RT prices during times of low reserves and help incentivize sufficient capacity to meet high net-load hours. The ORDC adder was added to the RT market prices when we calculated P_{rt} for ERCOT plants.

All nodal prices are reported by each ISO and gathered through ABB’s Velocity Suite data product. Nodes were matched to solar plants either within the Velocity Suite product, or based on proximity. ISOs generally have RT pricing with sub-hour time resolution. However, we used hourly average prices for our analysis for tractability.

Finally, we note that the RT and DA energy markets do not cover all system costs, as discussed in prior literature (Luoma, Mathiesen et al. 2014, Martinez-Anido, Botor et al. 2016, Antonanzas, Pozo-Vázquez et al. 2017). For example, we do not account for any regional specific regulations that might impact solar plants, such as the possibility of additional fees tied to missed forecasts, or other market details (e.g. settlement times, sub-hourly RT pricing, etc.). We also do not explore any impact of forecasting on capacity markets. Capacity markets exist in all regions we study (except Texas) and are separate from the energy markets. Capacity markets are designed to ensure that enough capacity exists to cover a small portion of the highest load hours (energy markets alone might not provide sufficient financial incentives for developers to build enough capacity to cover these particular hours). Our focus on energy markets will leave out any impact that uncertainty in solar output has on prices in capacity markets. That said, in most regions the capacity markets tend to provide much less revenue to solar plants than energy markets, and solar capacity revenue tends to decline with high solar penetration (Bolinger, Seel et al. 2020, Millstein, Wiser et al. 2021).

2.4. Error metrics

To quantify forecast errors, we used the metric of normalized Root Mean Square Error (nRMSE). nRMSE is defined as the Root Mean Square Error (RMSE) normalized by the maximum actual generation (see Eq. 5–6).

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(E_f - E_a\right)^2} \tag{5}$$

$$nRMSE = \frac{RMSE}{\max\limits_{1\le i\le N} E_a} \tag{6}$$

where E_f is the energy generation forecast, E_a is the actual generation, and N is the number of hours.

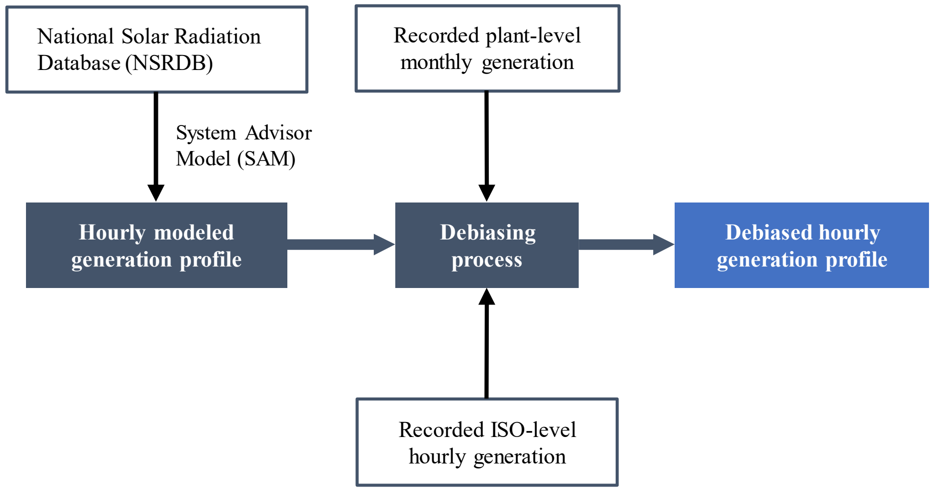

Fig. 3. The approach used to estimate plant-level hourly RT generation data.

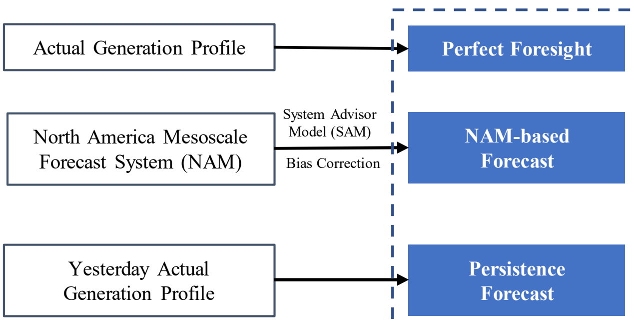

Fig. 4. Preparation of NAM-based and persistence forecasts. Perfect foresight assumes forecast generation equals actual generation.

2.5. Hourly generation time-series

Forecast errors are calculated as the difference between forecast generation and actual hourly generation. However, generation records are not publicly available at the plant and hourly level. To overcome this challenge, we developed a 'debiased hourly generation profile' for each plant. Fig. 3 summarizes the process we used to create the debiased generation profiles. The debiased hourly generation profiles were utilized as actual generation at sample plants in this study.

Meteorological data was derived from the National Solar Radiation Database (NSRDB), which is widely used and provides three solar radiation measurements: global horizontal irradiance (GHI), direct normal irradiance (DNI), and diffuse horizontal irradiance (DHI) (Sengupta, Xie et al. 2018). A raw generation profile was modeled at each plant based on the NSRDB data using the System Advisor Model (SAM) (Blair, DiOrio et al. 2018). These profiles were then adjusted for curtailment where and when such data was available. An iterative debiasing process was performed, to ensure that the generation from each plant matched both ISO-level hourly generation and plant-level monthly generation. The iterative process contained two main steps. First, each plant's generation was scaled so its hourly profile summed to match to its recorded monthly total. These totals were available from the U.S. Energy Information Administration (EIA) form 923 (Energy Information Administration 2020b). In the second step, hourly scaling factors were developed for all plants in each region. Hourly total solar generation is reported by each ISO, and a single scaling factor for each hour is applied across all plants in a region so that the modeled hourly total matches the recorded total for each hour. These two steps are then repeated until the change in values between each cycle falls below a minimum threshold. Note that output is scaled within the physical limits possible for each plant (i.e., once a plant reached 100% of its output, it could no longer be scaled up, and other plants (or other hours, depending on the step of the iteration) would be scaled up instead. The debiasing process has been described in previous work (Wiser, Millstein et al. 2020, Wiser, Bolinger et al. 2020, Millstein, Wiser et al. 2021).

2.6. Solar forecast techniques

Forecasting techniques can be generally classified into three major categories: physical, statistical, and ensemble methods (Chaturvedi and Isha 2016, Sobri, Koohi-Kamali et al. 2018). Physical methods are based on knowledge of the underlying processes that link solar irradiance and physical conditions in the atmosphere (Lorenz, Heinemann et al. 2007). Meteorological models—either based on local measurements, satellite imaging, or both—are a fundamental part of the physical solar forecasting methods. In contrast, statistical approaches are purely data-driven (Martín, Zarzalejo et al. 2010). These approaches seek to establish a relationship between parameters and power output based on historical observations, and then use it for prediction. Machine learning methods, a large contributor to recent advancements in forecasting techniques, falls into the statistical category. Ensemble methods are combinations of both physical and statistical methods (Gala, Fernández et al. 2016).

In this paper we compare the cost of forecast errors of a physical forecast method (NAM) to the costs of errors from persistence based forecasts (Fig. 4). Though there are many ways one could improve the physical forecast we use (it is simply based on publicly available meteorological forecast data), optimizing forecasts to reduce error is not the focus of this paper. Instead, we are more interested in the costs of overall forecast errors, and how these change across regions and times. By bounding the maximum forecasts error costs with persistence forecasts and observing the cost reduction derived from using NAM forecasts, we are also able to gain insight into the benefits from improving forecasts.

The physical forecast approach is based on the North American Mesoscale Forecast system (NAM) (NCEP 2020). NAM is a numerical weather prediction model that generates 12 km-resolution forecasts 4 times every day on 00/06/12/18 UTC. It predicts weather and irradiance up to 84 h ahead, with hourly forecasts in the first 36 h. In this study, 24–29 h ahead NAM forecasts from each run were compiled as a continuous DA forecast time series at each sample plant. Then the System Advisor Model (SAM) was used to translate NAM irradiance forecasts into DA energy generation forecasts. We assume that plant owners will at least perform a basic bias correction before using NAM-based forecast for market bidding, therefore, we scale hourly NAM forecasts

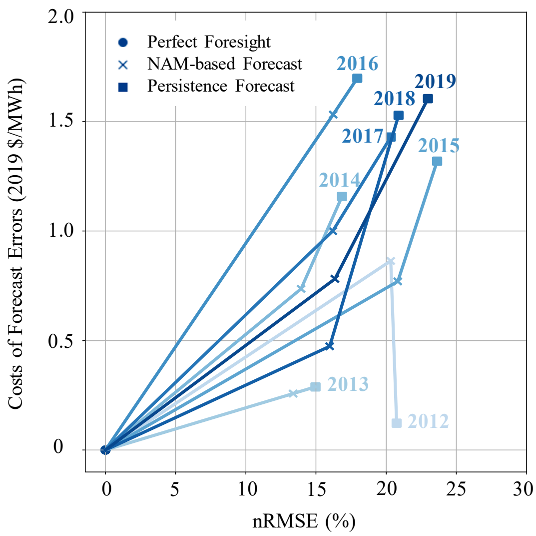

Fig. 5. In most years, the cost of forecast errors increases at an accelerating rate with the magnitude of the forecast errors (nRMSE). The annual, fleet-wide weighted average is shown for years 2012–2019, with weights based on recorded generation by plant.

such that the total output forecast for each month matches the recorded monthly total at each plant (while maintaining physical limits to plant output in individual hours). This final time series is referred as a “NAM-based forecast”.

Finally, the naïve one-day persistence model, assumes tomorrow’s solar energy generation equals today’s generation (i.e., generation in hour T + 24 equals generation in hour T). This is referred to as a “persistence forecast”.

3. Results and discussion

3.1. 2012–2019: Overall cost of forecast errors

We define the cost of solar forecast errors as the cost of compensating for DA forecasts errors through purchasing or selling energy in the RT markets (this is the second term of Eq. (2)). Though it is possible for these transactions to have a positive value (for example, an over forecast in the case where DA price is greater than RT price), on average, these compensating transactions were costly.

Fig. 5 shows the average costs of forecast errors across all plants in the study for NAM and persistence forecasts and for years 2012 through 2019. NAM nRMSE varied between 13.4% and 20.8%. The cost of these NAM forecast errors varied by year, ranging from 0.3 to 1.5 $/MWh, though staying at 1 $/MWh or less in all years except 2016. The persistence forecast nRMSE ranged between 15.0% and 23.6% during the study period, and was larger than NAM nRMSE in all years. The average cost of these persistence errors was ~ 50% greater than the average cost of NAM errors, though this margin also varied significantly by year. The difference between the NAM and persistence forecasts nRMSE was relatively small, ~3% on average (e.g., in 2015 NAM nRMSE was 21% and persistence nRMSE was 24%).

In 5 of 8 years, cost per % of nRMSE was larger for persistence forecasts than NAM forecasts. In those years, the results suggested diminishing returns to forecast improvements. That is, improvements in nRMSE produce smaller cost benefits as forecasts move closer and closer to being perfect. Data from three years (2012, 2013, and 2016) did not follow this pattern, and these exceptions are discussed at the end of section 3.2.1. The pattern of diminishing returns to forecast improvements is consistent with a study of forecast values in ISO-NE (Martinez-Anido, Botor et al. 2016). However, the results presented in Fig. 5 are in contrast to results from a previous study in Spanish markets, which suggested the relationship between market-defined forecast value and forecast error follows an almost linear fit (Antonanzas, Pozo-Vázquez et al. 2017).

3.2. Cost of forecast errors by ISO in the context of growing solar deployments

In 2019, CAISO and ISO-NE were the only two regions where solar accounted for more than 2% of total electricity generation. These two regions had, on average, higher costs to forecast errors than the other regions. For example, if we examine the recent years 2017 – 2019, the costs of NAM forecast errors in CAISO and ISO-NE ranged from 0.3 to 1.7 $/MWh, and the cost of persistence forecast errors ranged from roughly 1.5 to 2.9 $/MWh (see Fig. 6). For context, many recent U.S. utility-scale solar power purchase agreements (PPAs) have contracted at below $40/MWh (Bolinger, Seel et al. 2020). In most other regions, the costs of forecast errors were of lower magnitude in most years. In 2019, the cost of NAM forecast errors was negative in SPP and ERCOT, and the cost of persistence errors was negative in PJM and SPP. Negative cost indicates that on balance, the errors would have increased solar plant revenue rather than imposed costs. We discuss ‘negative costs’ later in this section.

In addition to interannual variation, the costs of forecast errors varied by plant. Fig. 7 shows the cost of NAM forecast errors at each plant in 2018 and 2019. The costs of NAM forecast errors varied across plants in CAISO and ERCOT more than across plants in other regions, and the costs of errors at individual plants in both regions ranged across positive and negative values. The variation across regions and plants indicates that the costs of forecast errors are somewhat stochastic in nature.

We note the caveat that sample sizes for ERCOT and SPP are smaller than other regions, and thus we are cautious about the conclusions we make for these two regions. All plants in the ERCOT sample started operation after 2012, and we do not begin covering ERCOT results until 2016–2019, when there were at least five plants in operation. PJM and SPP results start from year 2014 for the same reason. Additionally, CAISO and ISO-NE results for years 2012–2013 should be treated with appropriate caution as sample sizes in those years are limited.

In one sense, the results shown in Figs. 6 and 7 are surprising: we expect there to be a cost to forecast errors but these figures show some examples of forecast errors increasing solar plant revenues. One explanation for the incidences of beneficial forecasting errors is that at low solar penetration, LMP prices are mostly independent from solar forecast errors. If prices were completely independent of solar forecasts, the impact to the cost of forecast errors would be random, occasionally positive and occasionally negative. All instances in Fig. 6 in which forecast errors increased revenue occurred in regions and years where solar penetration was equal to or less than 2% (implying that RT prices were relatively independent of solar). Specifically, solar accounted for 1% or less of generation in SPP, PJM, and ERCOT through 2019, and solar penetration was less than 2% in ISO-NE through 2016, and in CAISO through 2012. Additionally, we note that sometimes, in low solar penetration regions, the higher nRMSE persistence forecasts can lead to lower error costs than NAM forecasts (e.g., ISONE in 2016 and ERCOT in 2018). The explanation for this phenomenon is the same as for why some low penetration regions sometimes show benefits to forecasting errors, that is, prices in these cases are not particularly correlated with solar forecasting errors.

A broader question, related to the more general concept of ‘solar integration costs,’ is, does the cost of forecast errors continue to grow with solar penetration? The growth of error costs would point to increasing costs associated with integrating higher levels of solar generation into an electricity system. Generally, given a well-functioning market, the costs of forecast errors to solar plants should roughly equal the cost to the grid of balancing those forecast errors. Previous work indeed suggests error costs may increase with solar penetration.

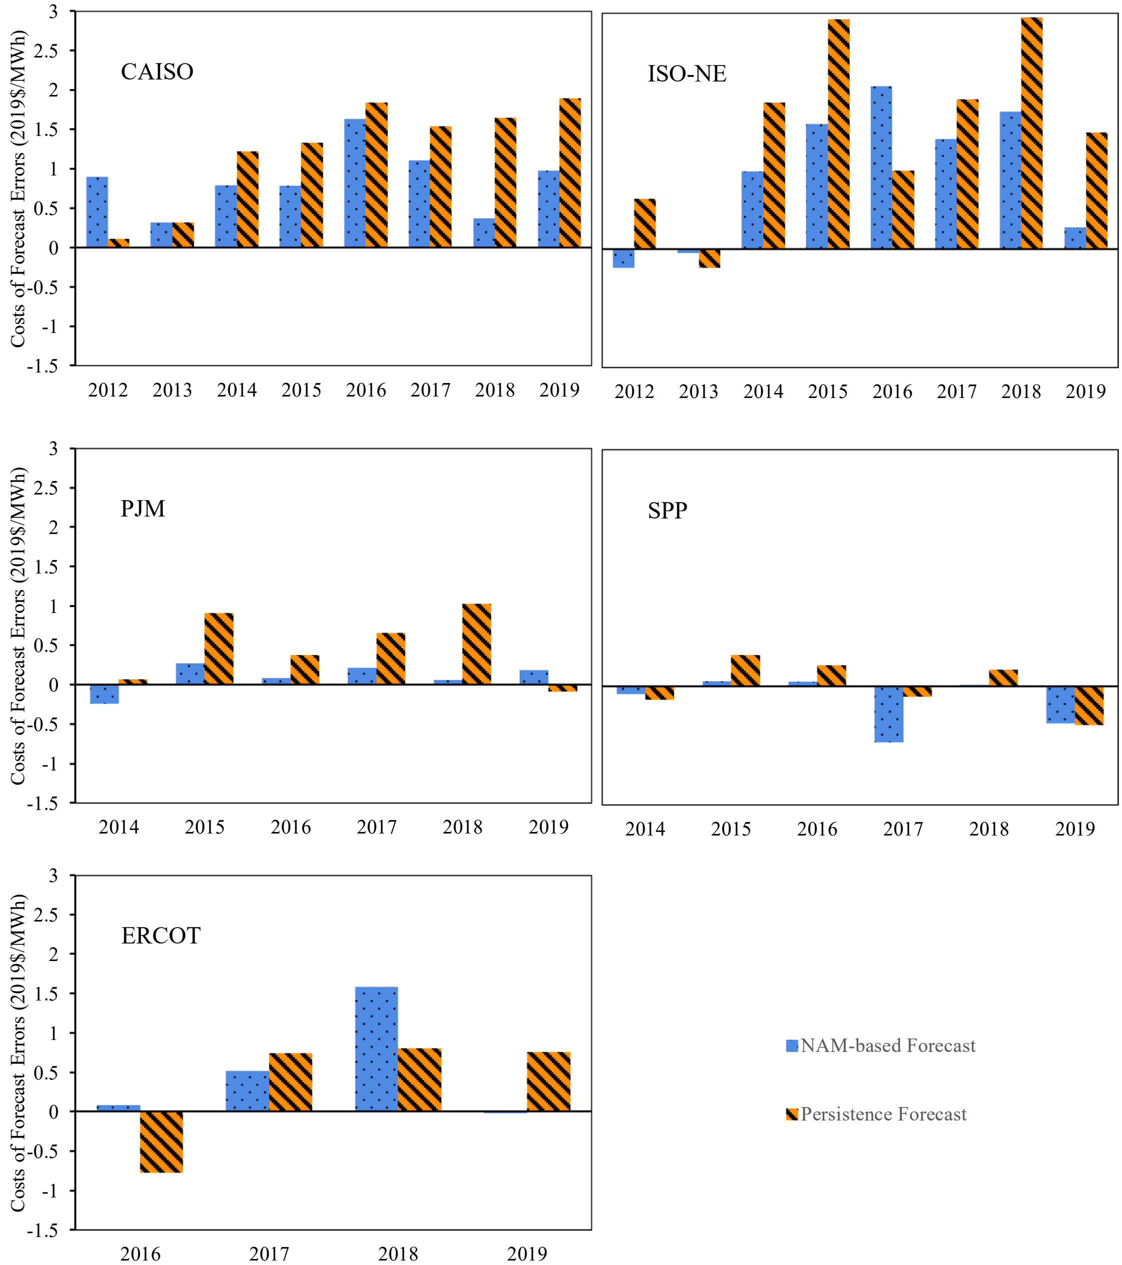

Fig. 6. Forecast error costs by ISO, year, and forecast type, with average costs weighted by the recorded generation by plant.

For example, Martinez-Anido et al. (2016) simulate increasing levels of solar deployment within ISO-NE and find that the system-level marginal (per MWh) cost savings of improved solar forecasting rises with higher solar penetration. However, the amount of increased costs is sensitive to composition of a region's electric grid. In regions with little flexibility, increased solar penetration may lead to high levels of solar curtailment, and in this case, solar forecast errors are often close to costless as they simply impact the amount of curtailment in RT. Wu et al. (2015) find a situation like this is possible in their simulation of a region with limited trade and relatively high nuclear capacity.

A generalized argument can be made that the costs of solar forecast errors will likely grow with solar penetration: (1) When solar is over forecasted for a certain hour RT prices will rise, relative to DA prices, as more expensive generators must make up for the absence of forecasted solar. (2) In this case, most solar plants participating in the DA market will need to purchase RT energy to make up for their actual generation shortfall (the cost of the forecast error depending on their shortfall and the difference in DA and RT prices). Note that under forecasts can also lead to costs, as in this case, solar plants will sell additional solar generation for possibly low RT prices. (3) As solar penetration grows, the impact of correlated forecasting errors across a region will have a larger impact on prices, and thus lead to larger costs to forecast errors.

We can look to our results for evidence that penetration is impacting the cost of forecast errors. Specifically, we can compare results in CAISO and ISO-NE to the other regions, as CAISO and ISO-NE have notably greater shares of solar than the other regions. In CAISO, the share of electricity produced by solar grew from 15% in 2017 to 19% in 2019. In ISO-NE, the solar share grew from 3% to 4% over the same period. From 2017 through 2019, we see NAM forecast errors are always less costly than persistence forecast errors in CAISO and ISO-NE. This is not the case in PJM, SPP, or ERCOT, in which there are instances during this time period when NAM errors were more costly than persistence errors. This suggests that the cost of forecast errors is more closely linked to the size of the error in the higher penetration regions of CAISO and ISO-NE

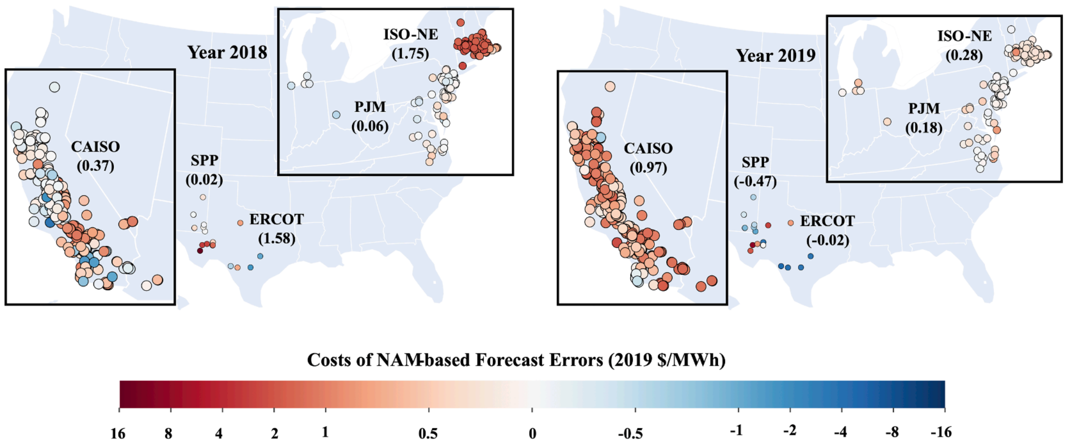

Fig. 7. Costs of NAM Forecast Errors in year 2018 and 2019. Regional averages are included in paratheses under the region names. Note that negative costs indicate an increase in revenue associated with solar forecast errors.

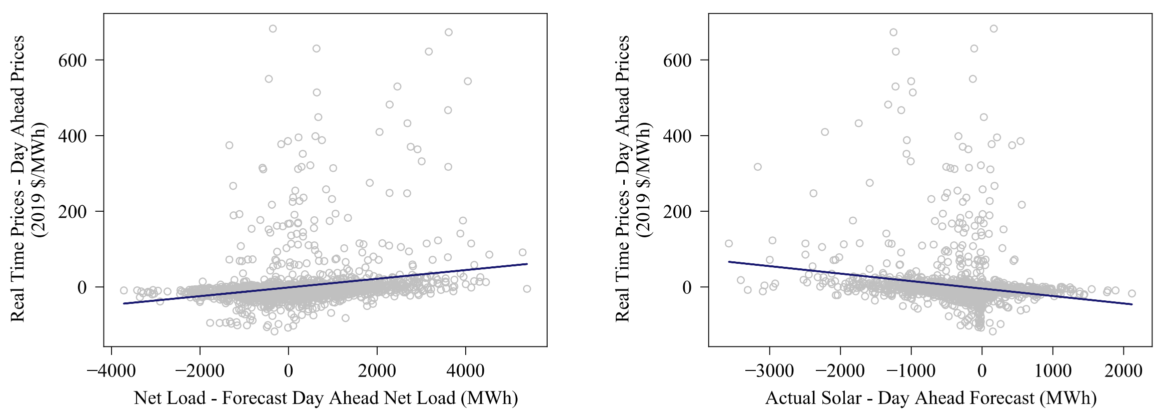

Fig. 8. Correlation between RT-DA price difference and net load forecasting errors (left) and solar forecasting error (right). Each point represents one hour in the local afternoon (between 21:00 and 2:00 UTC) in 2019. Prices are averaged across three hubs (SP, NP, and ZP), and solar and load forecasts are as reported by CAISO, and summed across regions to provide a CAISO-wide total. Data is derived from all hours.

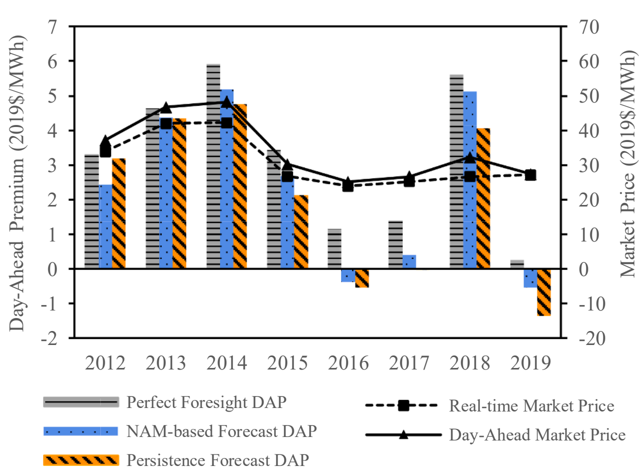

Fig. 9. The weighted average value, across all plants, of participating in the DA market (DAP) and the weighted average wholesale market prices in 2012–2019. Weights were based on recorded generation by plant.

compared with the remaining low penetration regions. However, there was no indication that the higher penetration region CAISO sees more expensive errors than ISO-NE.

3.2.1. Empirical relationship between solar forecast errors and RT-DA price differences

One possible cause of the lack of differentiation between CAISO and ISO-NE forecast error costs is the minimal overall correlation between solar forecast errors and RT-DA price differences. This was true even in CAISO in 2019, despite the high penetration of solar. For example, Fig. 8 shows the correlation between RT-DA price differences (RT minus DA prices) and CAISO reported net load forecast errors, and also RT-DA price differences and CAISO reported solar forecast errors. Note that the y-axis is the "(P_{rt}-P_{da})" part of the cost of forecast error term in Eq. (2) and that the x-axis is the ε part of Eq. (2), but aggregated for all plants in CAISO. The more positive the slope, the bigger the cost of forecast errors. It is clear from visual inspection that there is minimal correlation of RT-DA price differences to both netload forecast errors and solar forecast errors. A simple linear regression between the prices and the errors shows a coefficient of determination of 0.05 and 0.03 for the netload errors and solar forecast errors, respectively. Clearly, variables other than net load or solar forecast errors more strongly influenced the RT-DA price spread (for example, unexpected transmission interruptions or unexpected plant maintenance issues).

However, despite the low coefficients of determination, there was clearly some impact of net load errors on price differences. For example, we see no major RT price spikes when solar power was under forecast by more than ~ 500 MWh, but we do see price spikes when there was an over forecast (or a small under forecast). Also, the linear regression for solar forecasts has a negative slope of 0.02 $/MWh^2, implying that when

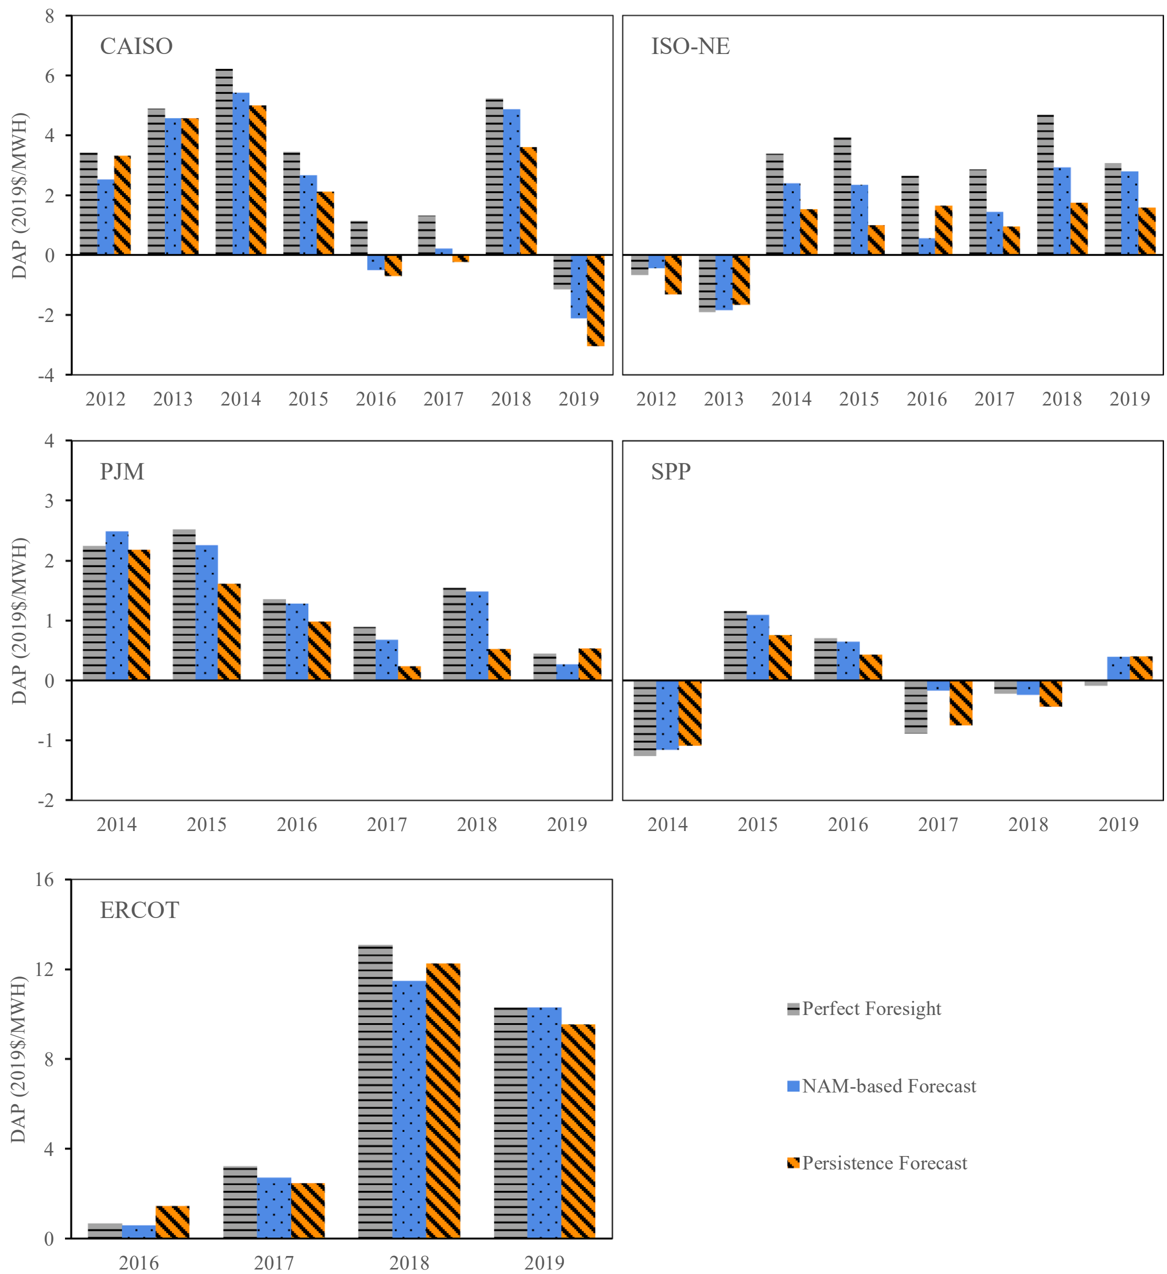

Fig. 10. The weighted average value, across all plants in each region, of participating in the DA market (DAP) in 2012–2019. Weights were based on recorded generation by plant.

actual solar generation is 1000 MWh below its DA forecast, RT prices rise by $20/MWh above the DA price. Similarly, the linear regression for netload forecasts has a positive slope of 0.01 $\$/MWh^2$, implying that when actual netload is 1000 MWh above its DA forecast, RT prices rise by $10/MWh. Solar forecast errors, of course, contribute to net load forecast errors. Currently, net load errors span a wider range than solar forecasting errors. For example, the maximum under forecast for solar is roughly 3000 MWh, while the maximum under forecast for netload is more than 4000 MWh.

We can also examine the relationship between plant-level forecast errors and local nodal RT-DA price differences. In Fig. 1A (in appendix A), we present the average slope, across all plants, of the linear fit of RT-DA price differences to NAM or persistence errors. In all years, and for both types of forecasts, the slope is mildly positive. For example, in 2019, a NAM over forecast equal to half the capacity of a plant would, on average, be associated with a $7/MWh increase in the RT-DA price difference. Another observation is that the slopes were increasing gradually over the study period. This gradual increase in slope values is generally consistent with the expected effects of increasing solar penetration over the study period.

We note that in most cases NAM error slopes were lower than persistence error slopes with the exception of three years (2012, 2013, and 2016), the same years we highlighted as following a unique pattern in our discussion of Fig. 5. Though 2012, and even 2013, may be sensitive to the relatively low numbers of plants in our sample those years, 2016 shows that even at higher sample sizes, it is possible for the NAM slopes to be steeper than the persistence error slopes. Whether RT-DA price differences were more sensitive to NAM or to persistence errors helps to explain the varying, non-linear, relationships we found between errors costs and error magnitude in Fig. 5. This raises the question, is there some fundamental process that would cause prices to be more sensitive to persistence errors than NAM errors in most years? One could hypothesize that perhaps the relationship between error costs and RT prices is not linear, and so the wider range of errors found with persistence errors leads to generally steeper linear slopes. Or, one could also hypothesize other causes, such as that persistence errors may be more correlated across a region than NAM errors, leading to a higher price sensitivity. However, we will leave investigation of these issues to further research.

Overall, both the CAISO-wide and plant-level observations suggests that solar forecast errors did cause a small increase to the RT-DA price spread, consistent with earlier research (Woo, Moore et al. 2016). As solar deployment continues, we expect the magnitude of region-wide solar forecast errors to become larger than the historical sample of region-wide net load forecast errors. One possibility is that much larger forecast errors could more clearly drive the RT-DA price spread, leading to larger costs to forecast errors. This is another question for future research, and a question that will be made more complicated, and interesting, by the deployment of energy storage.

3.3. Benefit to solar plants of participating in DA markets

In this section we raise the question of whether it is more attractive for a solar plant to participate in the DA market and be exposed to the cost of forecast errors versus just participating in the RT market. In general, we do not expect a large price premium in the DA market versus the RT market. If a predictable price premium existed, it would likely be quickly reduced through market trading. There is perhaps some incentive for a mild price premium in DA markets due to the value of reducing risk for certain parties.

Our findings reflect these conditions, that is, mild price premiums for participation in the DA market. For example, Fig. 9 shows that with perfect foresight, and averaged across all solar plants in our sample, participation in the DA market afforded a small premium over the RT market in all years. This day-ahead premium (DAP) varied between $0.3 to $5.9/MWh, depending on the year. Forecast errors degraded this DAP in all years. The value of the NAM-based DAP ranged from −0.5 to $5.2/MWh, depending on the year. Further degradation of forecasting accuracy, represented by persistence forecast, eroded the annual DAP value to -$1.4 to $4.8/MWh. Thus, the range in DAP across years is larger than the range of the cost of forecast errors, and so, despite the reduction to DAP due to forecast errors, the NAM DAP was positive in all years except 2016 and 2019.

From 2012 through 2015, the average DAP was closely correlated with average wholesale prices (see Fig. 9). However, this correlation diminished after 2015, as the margin between average DA and RT collapsed in 2016, 2017, and 2019. Because our sample of plants is heavily weighted toward plants in California, these trends in DAP may not be perfectly representative of the situation in other regions. But, before we dive into regional results, there are important broad trends which can be illustrated here. While wholesale electricity prices can be driven by many factors, including thermal or other power plant retirements, regulations, and other factors, the cost of natural gas has been particularly important in driving wholesale prices over the last decade. Natural gas price increases drove electricity prices up from 2012 through 2014 (EIA 2014, EIA 2015). Natural gas prices (and average electricity prices) then fell in 2015 and again in 2016 (EIA 2016, EIA 2017), remaining at relatively low levels through the rest of the study period. However, one important exception to the low gas (and electricity) prices occurred in July of 2018 in CAISO, when recorded electricity demand and a brief spike in gas prices lead to the highest monthly electricity price observed in CAISO since 2009 (EIA 2019). This price spike in 2018 also corresponded to a peak in DAP for solar (and to an increase in the average margin between DA and RT prices). Below we delve further into the drivers of DAP variation, but the general context is that annual DAP is loosely correlated with electricity and natural gas prices, and DAP can drop to close to $0/MWh when electricity and gas prices are low and the average margin between the DA and RT markets collapses, and conversely, may spike when the opposite conditions are seen, especially if high electricity prices correspond to sunny times of year in solar heavy regions.

DAP varies on a regional basis (see Fig. 10). Though the cost of forecast errors reduces the premium for participating in the DA market, regional and year-to-year variations in the DA to RT price spread have, to date, been larger than the costs of errors. The factors that drive these DA to RT price differences have little to do with solar and more to do with macro factors like gas prices, etc. With this context in mind, we summarize the regional differences in DAP below.

CAISO DAP values were highest in 2014 and 2018. In 2019, DAPs of all three forecasts were negative, indicating that even with a perfect foresight, market conditions were unfavorable for participating the DA market that year. For other years (2012–2018), perfect foresight DAP values in CAISO were positive in the range of $1.1 - $6.2/MWh, while NAM DAP obtained 57% of this value on average, and persistence DAP obtained 46%.

Compared to CAISO, ISO-NE exhibited less variable DAP values, especially in recent years. Since 2014, DAPs in ISO-NE were positive for both forecasting techniques (and perfect foresight). This implies solar energy could have consistently earned a premium—even with the least accurate persistence forecasts—through participating in the DA market, compared to participating only in the RT market. For ISO-NE, DAPs under perfect foresight in 2014–2019 ranged from $2.6 to $4.7/MWh. Quite similar to CAISO, NAM-based forecast obtained about 59% of perfect foresight DAP in ISO-NE, and persistence forecasts obtained about 43% of perfect foresight DAP in ISO-NE.

In PJM and SPP, DAP values were generally lower than those in CAISO and ISO-NE. The cost of forecast errors in these regions was also lower (in some cases forecast errors increased revenue). Across all years, the NAM DAP ranged from roughly −1 to 2 $/MWh in these regions. This is in contrast to ERCOT, which had relatively large DAP values in recent years. ERCOTs high DAP values coincided with larger-than-average price volatility, which was partially associated with ERCOT's energy-only market design. The increasing price and volatility in 2018–2019, in comparison with those in 2017, elevated the average DAP at ERCOT plants. An example of the pricing volatility occurred in August 2019, when record high energy demand drove RT power prices to reach their $9000/MWh price cap for a few hours (EIA 2020). Thus, in ERCOT, the difference between DA and RT prices can reach thousands of dollars per MWh. With such high differences between RT and DA prices, the daily DAP total can vacillate between large positive and negative values. Overall, the high volatility in ERCOT prices in 2018 and 2019 led to relatively high annual DAP values, and in comparison to DAP, the cost of forecasting errors in ERCOT was small.

4. Conclusion

In this paper we analyzed the cost of solar forecast errors and the benefit of participating in DA markets at 667 utility-scale plants across the U.S. Our study period covered 2012 – 2019 and costs were assessed based on hourly price differences between DA and RT market prices. We calculated the costs of forecast errors from two sources, first a publicly available meteorological forecast model (NAM forecasts), and second, a persistence algorithm. Nationally, the costs of NAM forecasting errors at solar plants averaged to no more than $1/MWh in all years except 2016 when it rose to $1.5/MWh. In a majority, but not all years, the NAM error costs were substantially lower than persistence error costs despite relatively small improvements to nRMSE, suggesting diminishing returns to reducing forecast errors.

We found some evidence that the relatively high solar penetrations in CAISO and ISO-NE led to increased costs of forecast errors in these regions. In CAISO and ISO-NE the cost of NAM forecasting errors was approximately $1/MWh (averaged over years 2017 – 2019). During that same period, the costs bounced around both sides of $0.0/MWh in SPP, PJM, and ERCOT. This indicates that there was no, or very little, cost to

NAM forecast errors in these low penetration regions, but modest costs in ISO-NE and CAISO. CAISO, however, has much greater solar penetration than ISO-NE (~20% versus ~4% in 2019). Thus, it is perhaps surprising that the costs of forecast errors are the same in these two regions. One explanation could be CAISO's participation in the Western Energy Imbalance Market. Participation in this regional market means that CAISO can readily call on resources outside of its boundaries to help balance solar forecast errors. It is possible that this regionalization helps to reduce the costs of forecast errors in CAISO and helped minimize differences in the cost of forecast errors between CAISO and ISO-NE. However, the impact of regionalization on the cost of forecast errors is a complex topic and will need to be addressed in future research.

We did find that the errors in solar forecasts that CAISO faced in 2019 were of smaller magnitude than the overall net load forecast errors. In other words, there are many sources of uncertainty in RT net load, and solar is not to the point yet where it is dominating this uncertainty. In fact, we saw only a low level of correlation between RT and DA price differences with net load forecast errors. This indicates that there are other processes that cause swings in RT prices (relative to DA) beyond changes to net load. One can easily imagine that over the next decade solar forecast errors will begin to drive net load forecast errors, and may expand the total size of these errors. In this case, net load forecast errors will move beyond the magnitude that we observed historically. A modeling effort would be required to investigate the potential costs of solar forecast errors in this future state of the system. It would also be interesting to explore how expanded reliance on battery storage may impact forecast error cost in the future.

Finally, we examined the question of whether it was more attractive for solar plants to participate in the DA market, despite the costs of forecast errors, versus just selling actual generation into the RT market. We found that there was indeed a small incentive for most solar plants to participate in the DA markets. When using NAM forecasting, the national average premium for participating in the DA market ranged from −0.5 to 5.2 $/MWh, depending on the year. The important drivers of the DA premium were regional and year-to-year variations in the systematic DA to RT price differences, factors that had little to do with solar and more to do with macro factors (e.g., gas prices). This lack of sensitivity to forecast error costs may change of course, if the costs of forecast errors increase in the future.

Data availability

Data and code are available upon request of the corresponding author, however hourly pricing data was purchased and the license does not allow us to share hourly pricing data directly.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The work described in this study was supported by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy under Lawrence Berkeley National Laboratory contract no. DE-AC02-05CH11231. The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges, that the U.S. Government retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes. The authors wish to thank Michele Boyd and Tassos Golnas (EERE) for helpful guidance and comments on the scope of the analysis. We wish to thank Will Gorman (LBNL) for development of SAM analysis scripts.

Appendix A

Fig. 1A below shows the average, across all plants in the study, of the slope derived by a linear fit of RT – DA nodal prices and NAM or persistence forecast errors. The errors were normalized by plant capacity, thus the units of the slope are simply $/MWh. A value of the slope of $10/MWh would imply that on average, the difference between RT and DA prices increased by $10/MWh if the output of a solar was over forecasted by 100% of a plant's capacity. The relationship was fit separately each year based on hourly data.

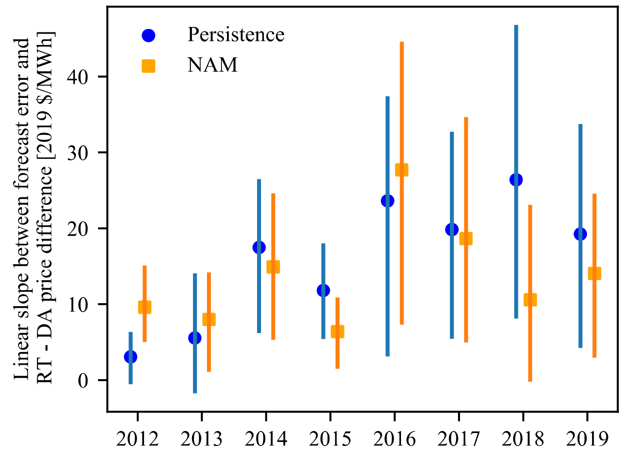

Fig. 1A. Average slope based on the linear fit of the RT-DA price difference to NAM or persistence errors. Note that errors here are normalized by plant capacity (the units are $/MWh as opposed to $/MWh^2, as was the case in Fig. 8 in the main text). The intervals around the mean show the range of slopes across the 20th to 80th percentile of all plants.

References

Antonanzas, J., Pozo-Vázquez, D., Fernandez-Jimenez, L., Martinez-de-Pison, F., 2017. The value of day-ahead forecasting for photovoltaics in the Spanish electricity market. Solar Energy 158, 140–146.

Blair, N. J., N. A. DiOrio, J. M. Freeman, P. Gilman, S. Janzou, T. W. Neises and M. J. Wagner (2018). System Advisor Model (SAM) General Description (Version 2017.9.5), National Renewable Energy Lab.(NREL), Golden, CO (United States).

Bolinger, M., J. Seel, D. Robson and C. Warner (2020). Utility-Scale Solar Data Update: 2020 Edition, Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States).

Chaturvedi, D., Isha, I., 2016. Solar power forecasting: A review. International Journal of Computer Applications 145 (6), 28–50.

Cirés, E., Marcos, J., de la Parra, I., García, M., Marroyo, L., 2019. The potential of forecasting in reducing the LCOE in PV plants under ramp-rate restrictions. Energy 188, 116053.

De Giorgi, M.G., Congedo, P.M., Malvoni, M., Laforgia, D., 2015. Error analysis of hybrid photovoltaic power forecasting models: A case study of mediterranean climate. Energy conversion and management 100, 117–130.

EIA (2014). Today In Energy, 2013: Average wholesale eletricity prices at major trading locations. U.S. Energy Information Adminsitration (EIA). https://www.eia.gov/todayinenergy/detail.php?id=14511.

EIA (2015). Today In Energy, 2014: Average wholesale electricity prices at major trading locations. U.S. Energy Information Administration (EIA). https://www.eia.gov/todayinenergy/detail.php?id=19531.

EIA (2016). Today In Energy, 2015: Average wholesale electricity prices at major trading locations. U.S. Energy Information Administration (EIA). https://www.eia.gov/todayinenergy/detail.php?id=24492#.

EIA (2017). Today In Energy, 2016: Average wholesale electricity prices at major trading locations. U.S. Energy Information Administration (EIA). https://www.eia.gov/todayinenergy/detail.php?id=29512.

EIA (2019). Today In Energy, 2018: Average wholesale electricity prices at major trading locations. U.S. Energy Information Administration (EIA). https://www.eia.gov/todayinenergy/detail.php?id=37912.

EIA (2020). Today In Energy, 2019: Average wholesale electricity prices at major trading locations. U.S. Energy Information Administration (EIA). https://www.eia.gov/todayinenergy/detail.php?id=42456.

EIA (2021). Electricity Generation by Energy Source/Activity. U.S. Energy Information Administration (EIA). https://www.eia.gov/international/data/world/electricity/electricity-generation.

Energy Information Administration (2020a). Annual Electric Generator Report EIA-860.

Energy Information Administration (2020b). Power Plant Operations Report EIA-923.

Gala, Y., Fernández, Á., Díaz, J., Dorronsoro, J.R., 2016. Hybrid machine learning forecasting of solar radiation values. Neurocomputing 176, 48–59.
Gürtler, M., Paulsen, T., 2018. The effect of wind and solar power forecasts on day-ahead and intraday electricity prices in Germany. Energy Economics 75, 150–162.
Hong, T., 2015. Crystal ball lessons in predictive analytics. EnergyBiz 12 (2), 35–37.
Kaur, A., Nonnenmacher, L., Pedro, H.T., Coimbra, C.F., 2016. Benefits of solar forecasting for energy imbalance markets. Renewable energy 86, 819–830.
Kraas, B., Schroedter-Homscheidt, M., Madlener, R., 2013. Economic merits of a state-of-the-art concentrating solar power forecasting system for participation in the Spanish electricity market. Solar Energy 93, 244–255.
Kulakov, S., Ziel, F., 2021. The impact of renewable energy forecasts on intraday electricity prices. Economics of Energy & Environmental. Policy. 10 (1) https://doi.org/10.5547/2160-589010.5547/2160-5890.10.110.5547/2160-5890.10.1.skul.
Lorenz, E., Heinemann, D., Wickramarathne, H., Beyer, H., Bofinger, S., 2007. Forecast of ensemble power production by grid-connected PV systems. 20th European PV Conference.
Luoma, J., Mathiesen, P., Kleissl, J., 2014. Forecast value considering energy pricing in California. Applied Energy 125, 230–237.
Martín, L., Zarzalejo, L.F., Polo, J., Navarro, A., Marchante, R., Cony, M., 2010. Prediction of global solar irradiance based on time series analysis: Application to solar thermal power plants energy production planning. Solar Energy 84 (10), 1772–1781.
Martinez-Anido, C.B., Botor, B., Florita, A.R., Draxl, C., Lu, S., Hamann, H.F., Hodge, B.-M., 2016. The value of day-ahead solar power forecasting improvement. Solar Energy 129, 192–203.
Millstein, D., Wiser, R., Mills, A.D., Bolinger, M., Seel, J., Jeong, S., 2021. Solar and wind grid system value in the United States: the effect of transmission congestion, generation profiles, and curtailment. Joule 5, 1749–1775.
NCEP (2020). North American Mesoscale Forecast System (NAM) [12 km], NCEI DSI 6173, gov.noaa.ncdc:C00630.
Ruhnau, O., P. Hennig and R. Madlener (2015). Economic implications of enhanced forecast accuracy: the case of photovoltaic feed-in forecasts.
Sengupta, M., Xie, Y., Lopez, A., Habte, A., Maclaurin, G., Shelby, J., 2018. The National Solar Radiation Data Base (NSRDB). Renewable and Sustainable Energy Reviews 89, 51–60.
Sobri, S., Koohi-Kamali, S., Rahim, N.A., 2018. Solar photovoltaic generation forecasting methods: A review. Energy Conversion and Management 156, 459–497.
Wiser, R., Millstein, D., Bolinger, M., Jeong, S., Mills, A., 2020a. The hidden value of large-rotor, tall-tower wind turbines in the United States. Wind Engineering.
Wiser, R. H., M. Bolinger, B. Hoen, D. Millstein, J. Rand, G. L. Barbose, N. R. Darghouth, W. Gorman, S. Jeong and A. D. Mills (2020). Wind Technology Data and Trends: Land-Based Focus, 2020 Update, Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States).
Woo, C.-K., Moore, J., Schneiderman, B., Ho, T., Olson, A., Alagappan, L., Chawla, K., Toyama, N., Zarnikau, J., 2016. Merit-order effects of renewable energy and price divergence in California's day-ahead and real-time electricity markets. Energy Policy 92, 299–312.
Wu, J., Botterud, A., Mills, A., Zhou, Z., Hodge, B.-M., Heaney, M., 2015. Integrating solar PV (photovoltaics) in utility system operations: Analytical framework and Arizona case study. Energy 85, 1–9.