

Solar and Wind Forecast Error Reserve Sharing in a Multi-Utility Region

Fredrich Kahrl^{1,2}, Jennie Jorgenson³, James Hyungkwan Kim¹,
Lawryn Kiboma³, Dev Millstein¹, Natalie Mims Frick¹, and Brian
Sergi³

¹ Lawrence Berkeley National Laboratory

² 3rdRail Inc.

³ National Renewable Energy Laboratory

October 2024



Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or The Regents of the University of California.

Ernest Orlando Lawrence Berkeley National Laboratory is an equal opportunity employer.

Copyright Notice

This manuscript has been authored by an author at Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. 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.

Solar and Wind Forecast Error Reserve Sharing in a Multi-Utility Region

Prepared for the
Solar Energy Technologies Office
U.S. Department of Energy

Principal Authors
Fredrich Kahrl
Jennie Jorgenson
James Hyungkwan Kim
Lawryn Kiboma
Dev Millstein
Natalie Mims Frick
Brian Sergi

Ernest Orlando Lawrence Berkeley National Laboratory
1 Cyclotron Road, MS 90R4000
Berkeley CA 94720-8136

October 2024

The work described in this study was funded by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy Strategic Analysis, Solar Energy Technologies Office under Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231.

Acknowledgements

The work described in this study was conducted at Lawrence Berkeley National Laboratory and the National Renewable Energy Laboratory and supported by the U.S. Department of Energy's (DOE's) Office of Energy Efficiency and Renewable Energy Strategic Analysis, Solar Energy Technologies Office under Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231. We especially thank Zachary Goff-Eldredge and Ammar Qusaibaty of DOE for supporting this work.

The authors thank the following experts for reviewing this report (affiliations do not imply that those organizations support or endorse this work):

Clayton Barrows
Ryan Wisler

National Renewable Energy Laboratory
Lawrence Berkeley National Laboratory

Table of Contents

- Table of Contents..... ii
- Table of Figures.....iii
- List of Tablesiii
- Abstractiv
- 1. Introduction 1
- 2. Background 1
- 3. Methodology..... 3
 - 3.1 Reserve Requirement Calculations 6
 - 3.2 Costs, Prices, and Reserve Provision..... 9
- 4. Results 10
- 5. Discussion: Decentralized Sharing of Forecast Error Reserves 16
- 6. Conclusions 18
- 7. References..... 21
- APPENDIX..... 23
 - A.1 ReEDS Inputs Assumptions and Installed Capacity by Scenario 23
 - A.2 Solar and Wind Forecast Error Calculations..... 27
 - A.3 Solar and Wind Forecast Error Distributions..... 29

Table of Figures

Figure 1. Change in reserve costs by reserve type in the medium and high coordination scenarios relative to the low coordination scenario	13
Figure 2. Annual day-ahead reserve provision by reserve type for each balancing area in each resource and coordination scenario.....	14
Figure 3. Share of reserve provision by technology and reserve type for each resource and coordination scenario.....	15
Figure A-1. ATB capital cost assumptions used in ReEDS (2004\$).....	24
Figure A-2. Installed capacity in the BPBS scenario and change in installed capacity in other scenarios relative to the BPBS scenario, Southeast region.....	25
Figure A-3. Installed capacity in the BPBS scenario and change in installed capacity in other scenarios relative to the BPBS scenario, neighboring regions.....	26
Figure A-4. Spatial representation of the Southeast in the ReEDS model.....	26
Figure A-5. Probability density function for load following solar forecast errors.....	29
Figure A-6. Probability density function for load following wind forecast errors	30

List of Tables

Table 1. Utility-scale solar PV and wind shares of total generation capacity, total storage capacity, and utility-scale PV-to-storage ratio for each resource scenario within the study footprint	4
Table 2. Description of coordination scenarios.....	5
Table 3. Average day-ahead load following, regulation, and total reserve requirements as a share of average load	10
Table 4. Total cost of reserves and energy by scenario (billion 2022\$).....	11
Table 5. Summary of total reserve requirement and shared reserve strategies for allocating forecast error reserves	17
Table A-1. Key ReEDS input assumptions for developing scenarios.....	23
Table A-2. Illustration of 15-minute and hourly forecast error calculations	28
Table A-3. Kurtosis and skew statistics for solar forecast errors by scenario.....	29
Table A-4. Kurtosis and skew statistics for wind forecast errors by scenario.....	30
Table A-5. Solar and wind forecast error coverage for different standard deviations and scenarios	31
Table A-6. Standard deviations for solar and wind day-ahead forecast errors by scenario	32

Abstract

As electricity systems transition to higher levels of solar and wind generation, electric system operators will likely need to hold additional reserves to manage solar and wind forecast error. Because solar and wind forecast errors tend to be weakly correlated across space, system operators can reduce their reserve requirements by sharing reserves. This paper examines the benefits of forecast error reserve sharing among balancing areas in the Southeastern United States, in scenarios in which solar and wind generation ranges from 34% to 65% of total generation. It finds that day-ahead forecast error reserve requirements increase linearly with growth in solar and wind generation capacity (6%-10% of total capacity), but that reserve sharing can significantly reduce these requirements (by 6%-29%). It finds that, in economic terms, the value of forecast error reserve sharing (\$0.09-\$1.24 billion per year, \$0.12-\$1.68/MWh of load across scenarios) tends to decline with higher levels of solar and wind generation, due to lower reserve and energy prices. Even with declines in reserve prices, forecast error reserve sharing can still provide substantial value, though with higher levels of solar, wind, and electricity storage this value is increasingly tied to avoiding scarcity prices.

1. Introduction

Higher levels of solar and wind generation may require system operators to hold additional operating reserves to manage solar and wind forecast error. Solar and wind forecast errors are not perfectly correlated across space, which means that the operating reserves required to jointly manage forecast uncertainty across multiple system operators will typically be lower than the sum of reserves required for system operators to do so individually. Sharing of forecast error reserves among multiple system operators may thus lower reserve costs, though doing so requires a formal reserve sharing mechanism and supporting infrastructure.

This paper examines the value of forecast error reserve sharing across multiple utilities in the Southeastern United States, a solar-rich region that does not have a centralized system operator, such as a regional transmission organization (RTO). The paper analyzes how the value of forecast error reserve sharing changes across different levels of solar and wind generation, different levels of energy storage, and different reserve sharing arrangements. It explores how more decentralized forecast error reserve sharing could be implemented in practice. Although the paper focuses on the Southeastern U.S., its conclusions are applicable to any region with multiple system operators and regions without centralized balancing markets.

2. Background

Operating practices for electric utilities and other electricity system operators can vary significantly, but in general most system operators tend to have some form of day-ahead supply-demand scheduling process, an intraday supply adjustment, and real-time balancing. During the day-ahead scheduling process, system operators set generation schedules for different time intervals in the following day, based on load, wind, and solar forecasts. They start up (commit) units as necessary to meet any shortfalls in generation and may shut down (de-commit) units if they are no longer needed to meet demand at lowest cost.

Within the operating day, system operators commit or de-commit additional units as needed to address generation or transmission outages and load, wind, and solar forecast error. Within the operating hour, some system operators also conduct economic dispatch, sending dispatch instructions to generators every 5 to 30 minutes. Within dispatch intervals, they often manage deviations in load or generation from forecasts and dispatch instructions through automated generation control (AGC) systems.

System operators hold operating reserves to ensure that their systems can respond to generation and transmission outages, forecast errors, and load and generation variability within the operating hour. In the U.S., typical kinds of reserves include (a) frequency regulation (secondary control) reserves, which address imbalances within dispatch intervals (seconds to minutes); (b) contingency (tertiary control) reserves, which address imbalances that result from generator and transmission outages (minutes); and (c) load following reserves, which address imbalances resulting from load, solar, and wind forecast error

and load, solar, and wind variability within the operating hour (minutes) (Ela et al., 2011). Higher levels of solar and wind generation may require utilities to hold more operating reserves to manage forecast error (Doherty and O'Malley, 2005; Ortega-Vazquez and Kirschen, 2009). For instance, if a utility forecasts 1,000 MW of solar generation between the hours of 09:00 and 10:00 but during the interval 09:00-09:15 only has 500 MW of solar available, the utility will need to ensure that it can provide 500 MW of additional power (upward ramping capability) during that interval.¹ This may require committing additional part-loaded units day-ahead to ensure that the utility has adequate spinning resources available. Utilities may also hold downward reserves to manage solar and wind under-forecasts. For instance, if a utility forecasts 1,000 MW of solar generation and has 1,500 MW available in real-time, it may want to have the flexibility to reduce other generation and absorb the additional 500 MW of solar rather than curtail it.

Solar, wind, and load forecast errors tend to be weakly correlated across space (Focken et al., 2002; Katzenstein et al., 2010), which means that spreading load following reserve requirements over a larger geographic area, through sharing of reserves, can reduce the total reserve requirement (GE Energy, 2010; EnerNex, 2011; King et al., 2011; E3, 2013; Samaan et al., 2017). Reserve sharing does not decrease the solar, wind, or load forecast error for any individual system operator. Rather, if forecast errors are weakly correlated across different systems, pooling forecast errors decreases the combined forecast error for all system operators.

Multi-utility sharing of reserves to manage solar and wind forecast errors ("forecast error reserves") may thus generate cost savings for utilities, but doing so requires a mechanism to facilitate sharing and deployment of these reserves. For instance, consider two utilities that individually have 500 MW and 1,000 MW forecast error reserve requirements and have a collective 1,300 MW requirement. If the utilities share reserves, they can reduce reserve requirements by 200 MW while maintaining the same level of reliability. However, how should responsibility for the 1,300 MW requirement be shared between the two utilities? The analysis in this paper is based on realistic assumptions about how reserve sharing might be carried out in practice. Section 5 (Discussion) explores how decentralized forecast error reserve sharing by utilities without a central system operator might be implemented in practice.

Several previous analyses have explicitly or implicitly examined the benefits of sharing forecast error-related reserves across a region, using different methods and with different assumptions about the level of solar and wind generation. GE Energy (2010) found that, with 10% wind and 1% solar generation, reserve pooling reduced operating costs in the Western Interconnection by \$2.3 billion (\$2.6/MWh of load, 2022\$). Also studying the Western Interconnection, Samaan et al. (2017) found

¹ The examples in this paper are time averaged for simplicity. In this case, the utility would forecast 1,000 MWh of solar energy for one hour (09:00-10:00), or 1,000 MWh/h and an average of 1,000 MW. In the 09:00-09:15 interval, the 500 MW of available solar would be an average over the 15-minute interval (125 MWh of energy). System operators hold operating reserves in units of capacity (MW) and use continuous duration requirements or operating constraints to ensure that reserves will be able to provide sufficient energy (MWh) over that interval. Reserves are a capacity product. A unit providing reserve capacity may never be called upon to provide reserve energy.

that, with 11% and 33% variable generation (8%-22% wind, 3%-6% solar), balancing area consolidation reduced regulation and load following reserve requirements by 58% to 72%. With approximately 10% combined wind and solar generation, E3 (2013) found a \$5 to \$98 million (\$0.02-\$0.34/MWh of load, 2022\$) reduction in flexibility reserve costs from reserve sharing between the California Independent System Operator (CAISO) and PacifiCorp, depending on transfer capability. Jiang et al. (2018), examining reserve requirements and reserve allocation strategies in hypothetical wind-rich systems, found that reserve sharing reduced costs by more than 20%. Focusing on smaller municipal utilities in Florida, Hale and Zhou (2021) found that, with 30% solar generation, collectively procuring operating reserves would reduce reserve requirements for the smallest utility by roughly 50%.

This study makes several contributions. First, it develops an approach for valuing reserve sharing in two-stage (e.g., day-ahead and real-time) production simulation models using model prices (shadow prices) rather than production costs. Model prices capture the relationship between marginal energy and reserve costs and allow reserve costs to be calculated separately from changes in production costs, which is helpful for isolating changes in reserve costs when both dispatch and reserve sharing assumptions change, as will often be the case. Second, whereas the focus in earlier studies was on wind, this study focuses on reserve sharing with higher levels of solar generation (23%-46%), analyzes much higher combined levels of solar and wind generation (34%-65%) than have been previously studied, and includes significant amounts of battery storage (13%-49% of peak load). Third, it evaluates reserve sharing under different levels of operational coordination and with different kinds of reserve sharing arrangements. In regions without centralized system operators, this approach is more realistic than assuming operators either share all or do not share any reserves. Fourth, this study explores how decentralized sharing of reserves to manage solar and wind forecast errors could be implemented in practice.

3. Methodology

This paper focuses on the Southeastern U.S. The Southeast has a large number of distribution utilities but a smaller number of system operators (“balancing area authorities” or “BAAs”), which tend to be larger utilities that manage operations on behalf of smaller ones. Utilities in the region have long traded electricity with one another and many began participating in a 15-minute Southeast Energy Exchange Market (SEEM) in 2022. However, the region does not have an RTO. Instead, each BAA is responsible for real-time energy dispatch and maintaining and deploying contingency, load following, and regulation reserves.

The Southeastern U.S. has abundant solar resources. The analysis examines future (2035) scenarios in which utilities in the region have significantly higher amounts of solar PV generation and electricity storage than they currently have. The resource portfolios used in the analysis were developed using the Regional Energy Deployment System (ReEDS) capacity expansion model (Ho et al., 2021). ReEDS is a U.S.-wide model that includes sub-regional resolution but does not provide detail down to the level of individual utilities. The analysis focuses instead on five Southeast balancing regions, which are

collections of transmission nodes in ReEDS. These balancing regions match approximately to five Southeast BAAs that were initial participants in SEEM: Associated Electric Cooperative Incorporated (AECI), Duke Energy, Kentucky Utilities (KU) and Louisville Gas & Electric (LG&E), Southern Company, and Tennessee Valley Authority (TVA).

The analysis considers resource scenarios that have progressively higher levels of solar PV and electricity storage within the study footprint in 2035 (Table 1). The solar scenarios include base solar PV (BP, where the P refers to PV), medium solar PV (MP), and high solar PV (HP) scenarios. The MP and HP scenarios each have a low storage (LS) and high storage (HS) scenario. This configuration results in five resource scenarios: (1) base solar, base storage (BPBS); (2) medium solar, low storage (MPLS); (3) medium solar, high storage (MPHS); (4) high solar, low storage (HPLS); and (5) high solar, high storage (HPHS). The Appendix (Section A.1) describes the cost and other assumptions behind each scenario and the resulting generation capacity portfolios.

Table 1. Utility-scale solar PV and wind shares of total generation capacity, total storage capacity, and utility-scale PV-to-storage ratio for each resource scenario within the study footprint

Resource Scenario	Description	Solar PV Share of Generation Capacity	Wind Share of Generation Capacity	Storage Capacity (PV-to-Storage Capacity Ratio)
BPBS	Base solar PV, base storage	27%	8%	19 GW (4.0)
MPLS	Medium solar PV, low storage	33%	6%	25 GW (4.0)
MPHS	Medium solar PV, high storage	39%	5%	45 GW (2.9)
HPLS	High solar PV, low storage	42%	12%	57 GW (3.0)
HPHS	High solar PV, high storage	43%	11%	70 GW (2.6)

We initially targeted different levels of solar PV and electricity storage in our study portfolios. However, interactions among resources in the ReEDS modeling made it difficult to isolate solar and storage. Lower cost storage enabled more solar PV capacity in the MPHS and HPHS scenarios, though this complementarity between solar and storage attenuated with higher levels of solar (HPLS vs. HPHS). In the MPLS and MPHS scenarios, lower solar cost assumptions and the resulting increase in solar PV capacity reduced wind capacity relative to the BPBS scenario. Alternatively, the HPLS and HPHS scenarios had significantly higher levels of wind capacity relative to the BPBS scenario, due to the carbon tax that we used in these scenarios to achieve higher levels of solar PV capacity (see Section A.1). These resource interactions are important for understanding the results.

For each resource scenario, the analysis examines three regional coordination scenarios, with progressively higher levels of forecast error reserve sharing and lower hurdle rates among balancing regions (Table 2). The model holds three kinds of reserves: contingency (both spinning and non-spinning), load following (both spinning and non-spinning), and regulation reserves. In the low coordination scenario, balancing regions can only share contingency reserves. In the medium

coordination scenario, balancing regions can share contingency and spinning load following reserves but all reserves that are allocated to a balancing region are required to be held within that region. In the high coordination scenario, balancing regions can share all reserves and reserves can be held anywhere in the Southeast. As an illustration, if the shared (Southeast-wide) reserve requirement in some interval is 1,000 MW and each balancing region is allocated 200 MW of this requirement, in the medium coordination scenario each balancing region will need to hold 200 MW of reserves within its territory; in the high coordination scenario, alternatively, this 1,000 MW could be held anywhere in the Southeast, including all within a single balancing region.

Table 2. Description of coordination scenarios

Coordination Scenario	Reserve Sharing	Reserve Location	Hurdle Rates Day-Ahead (Real-Time)
Low	Contingency reserves	Reserves held within each balancing region	\$10/MWh (\$5/MWh)
Medium	Contingency and spinning load following reserves	Reserves held within each balancing region	\$5/MWh (\$0/MWh)
High	Contingency, all load following, and regulation reserves	Reserves can be held anywhere in the Southeast	\$0/MWh (\$0/MWh)

Notes: Hurdle rates are \$/MWh virtual cost adders that increase the cost of importing and exporting electricity among regions but are not included in cost accounting. The hurdle rates here are in 2004\$, consistent with ReEDS and PLEXOS inputs, and are based on E3 (2011) and Chang et al. (2016). The rationale for using lower hurdle rates in real-time is that real-time transactions can use transmission capacity that has not been reserved day-ahead, in areas with some form of physical transmission rights.

The production simulation modeling uses a two-stage approach. In a day-ahead stage, the model schedules generation to meet hourly demand, using hourly averaged solar and wind profiles that incorporate persistence-based solar and wind forecast errors. In a real-time stage, the model re-dispatches generation to meet 15-minute demand, using 15-minute actual solar and wind profiles and incorporating operating constraints on generators between time periods, including fixed coal commitment and hydro dispatch. This two-stage modeling approach captures the general scheduling and dispatch process used by most system operators. We use the PLEXOS model for production simulation.²

Both the ReEDS and PLEXOS modeling for this study are national in scale. They incorporate all load and generation for the continental U.S. The rationale for using a national model is that changes in system operator coordination within a region are likely to lead to changes in its interactions with neighboring regions. For instance, the development of an RTO-style market in the Southeastern U.S. would lead to changes in intertie flows with the Midcontinent Independent System Operator (MISO), PJM Interconnection (PJM), and the Southwest Power Pool (SPP) on its borders. Similarly, changes in reserves and energy costs resulting from different levels of regional coordination are interactive and

² See <https://www.energyexemplar.com/plexos>.

should be considered in tandem. Changes in energy prices will have a significant impact on the cost of reserves and the value of sharing reserves (Hummon et al., 2013).

3.1 Reserve Requirement Calculations

We incorporate solar and wind forecast errors into load following and regulation reserve requirements, which we calculate exogenously. The model holds hourly load following reserves in day-ahead scheduling and releases them to be dispatched in real-time.³ The model holds contingency and regulation reserves in both day-ahead and real-time. Due to data limitations, we do not account for load forecast errors in determining load following reserves. Load following reserves are thus held exclusively to manage wind and solar forecast error. Regulation reserve requirements include a wind and solar forecast error component and a component tied to load forecast error and variability. PLEXOS allows users to set different required response times for different categories of reserves, which means that resources will need to have adequate ramp capability and residual capacity (headroom or legroom) to provide a given category and level (MW) of reserve. We use a 20-minute response time for all load following reserves and a 5-minute response time for contingency and regulation reserves. Spinning reserves must already be committed but non-spinning reserves can be offline. We require storage to have sufficient state of charge to be able to provide a given level of reserves for one continuous hour.

Following the approach in EnerNex (2011), we calculate load following and regulation reserve requirements related to wind and solar forecast errors using five steps: (1) calculate solar and wind forecast errors for each solar and wind generation decile; (2) calculate the standard deviation of forecast errors for each generation decile; (3) fit a function to forecast error standard deviation and generation decile data; (4) use the function in (3) to calculate hourly and 15-minute forecast error standard deviations; (5) use the forecast error standard deviations in (4) to calculate reserve requirements.

For step (1), we calculate solar forecast errors using a clear-sky persistence forecast and wind forecast errors using a simple persistence forecast. We use different forecast error calculations for load following reserves, which are only held in the day-ahead stage, and regulation reserves, which are held in both day-ahead and real-time stages. For load following reserves, we calculate forecast errors as the maximum absolute difference between the 15-minute generation profile in the next hour and the hourly average generation profile in the previous hour. This approach allows us to incorporate sub-hourly variability and uncertainty into forecast errors in a system that has hourly day-ahead scheduling. We use hourly rather than daily differences to calculate persistence forecast errors because these better reflect forecasting accuracy that should be achievable with modern forecasting systems (see Section A.2). For regulation, we calculate forecast errors as the difference between 15-minute generation profiles. We generate solar and wind profiles at the ReEDS transmission node level, using

³ These reserves are thus mainly used for day-ahead commitment and are more akin to commitment reserves. The real-time model solves hourly and includes a lookahead period, which means that these reserves do not likely need to be maintained in real-time dispatch.

the National Renewable Energy Laboratory's System Advisor Model (Blair et al., 2018) via the Renewable Energy Potential (reV) model (Lopez et al., 2024), but we aggregate profiles and calculate forecast errors to the balancing region level.

Step (2) is straightforward.

For step (3), we use a second order polynomial to fit forecast error standard deviations and generation (G).

$$\sigma = \alpha + \beta_1 G^2 + \beta_2 G + \varepsilon \quad (1)$$

where σ is a vector of solar or wind forecast error standard deviations by generation decile and G is a vector of solar or wind generation for each decile, and parameters α , β_1 , β_2 , and ε are coefficients estimated from the regression. The second order polynomial captures the nonlinear relationship between solar and wind generation and the standard deviation of their forecast errors. Forecast error standard deviation increases until the fourth or fifth generation decline, after which it begins to decrease.

We then calculate solar and wind forecast error standard deviations for each time period t (hourly or 15-minute) using the results of the regression in Equation 1.

$$\sigma_t = \alpha + \beta_1 g_t^2 + \beta_2 g_t \quad (2)$$

where σ_t is the hourly or 15-minute standard deviation of solar or wind generation forecast errors, β_1 and β_2 are the coefficients estimated in the previous regression, and g_t is solar or wind generation. The resulting solar and wind forecast errors are almost entirely uncorrelated with each other and we assume that they are from independent distributions. The load following reserve requirement in each hour ($LF.RR_h$) is

$$LF.RR_h = \theta \sqrt{s.\sigma_h^2 + w.\sigma_h^2} \quad (3)$$

where θ is a parameter for the number of standard deviations of forecast errors that system operators hold load following reserves to cover, $s.\sigma_h$ is the standard deviation of solar forecast errors in hour h and $w.\sigma_h$ is the standard deviation of wind forecast errors in hour h . We use a value of $\theta = 1$ for spinning load following reserves and a value of $\theta = 2$ for non-spinning load following reserves. The model only holds upward spinning load following reserves, which assumes that real-time curtailment will be the lowest cost strategy to manage under-forecast risk. Each type of reserve (spinning, non-spinning, regulation) has different commitment, response time, and duration requirements in PLEXOS. The regulation reserve requirement uses the same general approach but includes an additional requirement of 1% of load in each period. Real-time regulation reserve requirements for each 15-minute interval ($RG.RR_t$) will thus be

$$RG.RR_t = \theta \sqrt{\left(\frac{1\% \times L_t}{\theta}\right)^2 + s.\sigma_t^2 + w.\sigma_t^2} \quad (4)$$

where θ is a parameter for the number of standard deviations of forecast errors that system operators hold regulation reserves to cover, L_t is 15-minute load, $s.$ σ_t is the standard deviation of solar forecast errors in a 15-minute interval, and $w.$ σ_h is the standard deviation of wind forecast errors in that interval.

We calculate day-ahead hourly regulation requirements ($RG.$ RR_h) as the maximum of 15-minute reserve requirements in each hour. For regulation reserves we use a value of $\theta = 3$. These θ values are consistent with EnerNex (2011) but in practice will be system specific (Krad et al., 2016). As a default, PLEXOS holds bidirectional regulation reserves. Exploring whether separate upward and downward reserves would affect the results is a topic for future research.

We use the same method for calculating reserve requirements across the regional coordination scenarios in Table 2. In scenarios in which reserve sharing is not allowed, we calculate forecast errors using each individual balancing region's wind and solar profiles. In scenarios in which reserve sharing is allowed – spinning load following in the medium coordination scenario and all load following and regulation in the high coordination scenario – we calculate forecast errors using the regional sum of solar generation and wind generation, which results in lower forecast errors and reserve requirements. In the medium coordination scenario, the pooled spinning load following reserve requirement is allocated to balancing regions based on their shares of hourly wind and solar generation.

In practice, sharing of load following and regulation reserves will require limits due to transmission constraints. For instance, if a utility is relying on 200 MW of shared forecast error reserves but only has 100 MW of available import capacity, the utility may face an energy shortfall if it needs to deploy the reserves. We use zonal import rather than network-wide limits, assuming that the latter would be overly complex to implement. Generically, the shared reserve requirement in each balancing region i for time interval t (SRR_{it}) will be

$$SRR_{it} = \max (RR_{it} - TX_i, RR_{Rt}) \quad (5)$$

where RR_{it} is the individual reserve requirement for balancing region i in time interval t , TX_i is the total amount of import transmission capacity into balancing region i , and RR_{Rt} is the shared reserve requirement for the region.

As is generally the case (Hodge et al., 2012; Mauch et al., 2013), our solar and wind forecast errors are skewed and leptokurtic (see Section A.3). In our case, logit and other transformations of the original profile data did not produce normally distributed forecast errors for either solar or wind, or by extension their joint distribution. The approach described in this section does not require a normality assumption, but if forecast errors are not normally distributed it is less straightforward to translate standard deviations into forecast error coverage because the joint distribution of solar and wind forecast errors is unknown. However, coverage can be roughly approximated from the individual distributions: one standard deviation covers 58%-70% (over-under forecast) of forecast errors for solar

and 82%-76% (over-under forecast) for wind, two standard deviations covers 83%-93% and 96%-90%, and three standard deviations covers 94%-98% and 99.4%-97% (see Section A.3). Forecast error coverage of the joint solar-wind forecast error distribution will be between these values.

3.2 Costs, Prices, and Reserve Provision

We calculate the value of reserve sharing using reserve prices (shadow price constraints from the model) rather than production costs. Using prices rather than production costs better reflects the opportunity cost of reserves and resource sharing and allows us to capture the significant changes in dispatch between the Southeast and MISO, PJM, and SPP that occur as a result of lower hurdle rates. We use a two-settlement approach to incorporate day-ahead and real-time energy and reserve prices, so that the cost of supply for product i (CS_i) is

$$CS_i = \sum_n \sum_t Q_{int}^{DA} \times P_{int}^{DA} + (Q_{int}^{RT} - Q_{int}^{DA}) \times P_{int}^{RT} \quad (6)$$

where Q_{int} is the quantity of product i supplied at transmission node n at time t , P_{it} is the price of product i at node n at time t , and the DA and RT superscripts refer to day-ahead and real-time stages. The i products include energy and each reserve type. Because load following reserves are not maintained in real-time, P^{RT} is zero and these reserves are only valued at the day-ahead price. Market costs based on market prices are not a net cost because they include net revenues to generation owners (inframarginal rents). For reserves, they are akin to the perspective of a system operator that procures reserves through a uniform price auction.

We use scarcity prices for reserve shortages in PLEXOS that reflect the value of lost load, rising from \$2,500/MW-h (non-spinning reserve), to \$5,000/MW-h (spinning reserve), to \$7,500/MW-h (regulation reserve), and finally to \$10,000/MWh (lost load).⁴ Since PLEXOS co-optimizes energy and reserves, reserve scarcity prices also affect energy prices. Scarcity prices have a significant impact on the results: a \$5,000/MWh price for 10 hours per year is equivalent to a \$50/MWh price for 1,000 hours per year. To assess the effects of scarcity pricing, we include post-processed scenarios in which reserve and energy prices are capped at \$1,000/MW-h or MWh, respectively.⁵

PLEXOS allows most resources, including wind and solar generation but not distributed PV or nuclear generation, to provide reserves as long as they satisfy rules for availability and duration. We do not consider scenarios in which solar and wind generation are unable to provide reserves. In our study year, 2035, with much higher levels of solar and wind generation online, it is unlikely that solar and wind generators would not be able to do so (Kim et al., 2023).

⁴ These prices are in 2004\$, consistent with the other inputs in ReEDS and PLEXOS. MW-h refers to a MW of reserves held for a duration of one hour. Lost load is in units of energy (MWh) here because it is the absence of reserves rather than a reserve shortage.

⁵ This energy cap was also implemented in 2004\$, though the results are in 2022\$.

4. Results

Table 3 shows the average day-ahead load following and regulation reserve requirements as a share of average annual load in each scenario. Across the resource scenarios (from BPBS to HPHS), higher levels of solar generation significantly increase reserve requirements. However, per unit of installed capacity, the standard deviation of solar and wind load following forecast errors is relatively constant across resource scenarios, at 2 MW/100 MW (solar) and 6 MW/100 MW (wind). Day-ahead forecast error reserve requirements thus scale linearly with installed solar and wind capacity, at 8%-10% of total capacity in the low and medium coordination scenarios and 6%-7% in the high coordination scenarios.

Table 3. Average day-ahead load following, regulation, and total reserve requirements as a share of average load

Coordination Scenario	Resource Scenario				
	BPBS	MPLS	MPHS	HPLS	HPHS
Load Following (Spinning + Non-Spinning Reserves)					
Low	10%	11%	13%	21%	20%
Medium	9%	10%	12%	19%	18%
High	8%	8%	9%	15%	15%
Regulation (Bidirectional)					
Low	3%	3%	4%	5%	5%
Medium	3%	3%	4%	5%	5%
High	2%	2%	2%	3%	3%
Total (Load Following + Regulation)					
Low	14%	14%	16%	26%	25%
Medium	13%	13%	15%	24%	23%
High	10%	10%	12%	19%	18%

Notes: Load following reserves are only held day-ahead; regulation reserves are held day-ahead and in real-time. Day-ahead and real-time regulation reserve requirements are similar, as a share of load. For simplicity, we do not show real-time regulation and total reserves in this table. For load following reserves, non-spinning reserves are two times as large as spinning reserves; to calculate the share of spinning reserves divide the total by three.

Across coordination scenarios (in each resource scenario column, from Low to High rows in Table 3), reserve sharing reduces average day-ahead load following and regulation reserve requirements, with larger reductions from the medium coordination scenarios (spinning load following reserve sharing) to the high coordination scenarios (spinning, non-spinning load following and regulation reserve sharing). As a reminder, the model also holds spinning/non-spinning contingency reserves (3%/3% of load in each period), which are not shown in Table 3 but will add 6% to the total numbers.

Reductions in reserve requirements from reserve sharing are higher in scenarios with higher levels of solar and wind generation, though percentage reductions are comparable across scenarios. For instance, in the BPBS scenario, load following and regulation reserve requirements in the high coordination scenario (10% of load) are around 4 percentage points and 29% lower than in the low coordination scenario (14% of load), whereas in the HPHS scenario reserve requirements in the high

coordination scenario (18% of load) are 7 percentage points and 28% lower than in the low coordination scenario (25% of load). With an average Southeast-wide load of 85 GW, each percentage point reduction in the reserve share of load is equivalent to an 850 MW reduction in required reserves.

For each resource and coordination scenario, Table 4 shows the total annual cost of reserves (load following, regulation, contingency) at uncapped reserve prices and with reserve prices capped at \$1,000/MW-h (in parentheses). For reference, Table 4 also shows the total uncapped and capped (in parentheses) cost of energy by scenario. In the medium coordination scenarios, annual reserve cost savings relative to the low coordination scenarios range from \$0.09 billion (\$0.12/MWh of load, MPLS scenario) to \$0.47 billion (\$0.64/MWh of load, BPBS scenario).⁶ In the high coordination scenarios, cost savings relative to the low coordination scenarios range from \$0.16 billion (\$0.21/MWh of load, MPHS scenario) to \$1.24 billion (\$1.68/MWh of load, BPBS scenario).

Table 4. Total cost of reserves and energy by scenario (billion 2022\$)

Coordination Scenario	Resource Scenario				
	BPBS	MPLS	MPHS	HPLS	HPHS
Total Reserve Cost (Billion 2022\$)					
Low	\$1.93 (\$0.67)	\$1.16 (\$0.41)	\$0.53 (\$0.27)	\$0.90 (\$0.36)	\$0.48 (\$0.22)
Medium	\$1.46 (\$0.45)	\$1.07 (\$0.36)	\$0.31 (\$0.13)	\$0.56 (\$0.19)	\$0.23 (\$0.08)
High	\$0.69 (\$0.22)	\$0.70 (\$0.19)	\$0.37 (\$0.09)	\$0.28 (\$0.10)	\$0.28 (\$0.06)
Total Energy Cost (Billion 2022\$)					
Low	\$38.97 (\$24.95)	\$29.28 (\$21.24)	\$20.78 (\$19.45)	\$10.20 (\$9.42)	\$10.60 (\$9.27)
Medium	\$36.72 (\$23.33)	\$31.16 (\$21.38)	\$27.11 (\$20.01)	\$8.69 (\$8.06)	\$8.65 (\$7.94)
High	\$31.23 (\$21.98)	\$30.19 (\$19.61)	\$31.06 (\$19.77)	\$9.51 (\$6.59)	\$14.29 (\$7.63)

Notes: Energy costs for each time interval are calculated as $(L^{DA} - DPV^{DA} + SC^{DA}) \times P^{DA} + ([L^{RT} - L^{DA}] - [DPV^{RT} - DPV^{DA}] + [SC^{RT} - SC^{DA}]) \times P^{RT}$, where L is load, DPV is distributed PV, SC is storage charging, and the DA and RT superscripts represent day-ahead and real-time. The \$1,000/MW-h cap is in 2004\$, consistent with model inputs. All costs in the table are in 2022\$, inflated from 2004\$ using the U.S. Bureau of Labor Statistics consumer price index.

The results in Table 4 capture the interplay of two forces: (1) reductions in reserve requirements from reserve sharing in the higher coordination scenarios, and (2) lower energy and reserve prices and thus a lower cost of reserves that result from higher levels of low marginal cost solar and wind generation and electricity storage. Since energy and reserve procurement are co-optimized, lower energy prices

⁶ We report \$/MWh costs and cost savings in terms of real-time end use load (excluding storage charging), which is 742 TWh across all scenarios.


translate into lower opportunity costs for providing reserves instead of energy and thus lower reserve prices.

In most scenarios, higher levels of coordination lower reserve costs. Scarcity prices play a role in the two scenarios where higher levels of coordination increase reserve costs (MPHS Medium to High coordination, HPHS Medium to High coordination in Table 4). If reserve prices are capped at \$1,000/MW-h (values in parentheses in Table 4) the cost of reserves falls with higher levels of coordination in all scenarios.

Across scenarios, reserve costs in a limited number of hours have a significant impact on the results, particularly in the high coordination scenarios. In the low and medium coordination scenarios, the highest 100 total reserve cost hours (out of 8,760 annual hours) account for 62% to 91% (low) and 62% to 93% (medium) of total reserve costs across resource scenarios. In the high coordination scenarios, they account for 98% to 99% of total costs. Most of this difference between scenarios is due to differences in spinning reserve (load following and contingency) costs. Because all spinning reserves are shared in both the medium and high coordination scenarios, this suggests that the ability to hold reserves anywhere in the region in the high coordination scenarios reduces spinning reserve prices to zero or very low levels in most hours.

Despite lower reserve prices and lower total reserve costs in the MP and HP scenarios, reserve cost savings from reserve sharing are still substantial in some scenarios. For instance, reserve cost savings in the MPLS and HPLS high coordination scenarios relative to the low coordination scenarios are \$460 million (\$0.61/MWh of load) and \$620 million (\$0.84/MWh) per year, respectively (Table 4). This suggests that, even as wholesale prices decline, reserve sharing may still provide significant operational benefits. The importance of scarcity prices for the results in all scenarios suggests that the value of reserve savings will depend on operating conditions and constraints, but also that there may be a large economic upside to reserve sharing by reducing exposure to the physical scarcity events that lead to scarcity prices.

Reserve costs by reserve category provide additional insight on the results. Figure 1 shows changes in reserve costs in the medium and high coordination scenarios relative to the low coordination scenario in each resource scenario. In the MPLS and MPHS medium coordination scenarios, spinning contingency and regulation reserve costs increase relative to the low coordination scenarios but then decline in the high coordination scenarios. Across scenarios, the largest share of total reserve cost savings is from non-spinning load following reserves, even in the medium coordination scenarios where these reserves are not shared across balancing areas. In the high coordination scenarios, the cost and prices of all non-spinning reserves (load following, contingency) fall to zero.



Notes: SCT is spinning contingency reserves, NSCT is non-spinning contingency reserves, REG is regulation reserves, SLF is spinning load following reserves, NSLF is non-spinning load following reserves. LC, MC, and HC are the low, medium, and high coordination scenarios.


Figure 1. Change in reserve costs by reserve type in the medium and high coordination scenarios relative to the low coordination scenario

The importance of non-spinning reserves in the results was unexpected. Production cost modeling will often ignore non-spinning reserves because they are difficult to meaningfully incorporate. As long as the system has adequate capacity and planned outages do not occur during periods of system stress, non-spinning reserve prices should, in principle, be zero. However, if energy and reserve constraints are nested, changes to the availability of spinning reserves may impact prices for non-spinning reserves. Lower non-spinning prices and costs in the medium and high coordination scenarios may thus reflect reserve sharing (high coordination scenario), the ability to hold reserves anywhere in the region (high coordination scenario), lower energy prices (high and medium coordination scenarios), or reduced pressure on non-spinning reserves due to sharing of other reserves. These different drivers are difficult to disentangle in more complex models.

In the low and medium coordination scenarios, non-spinning load following reserve prices are only positive during three to four hours of the evening (19:00-22:00) in summer months (June to September) when demand is high and solar PV generation drops off. In all scenarios, these hours are the periods when operating reserve margins (available generation minus load) are lowest. This suggests that the benefits of sharing non-spinning load following reserves are related to relieving system capacity constraints. We discuss the implications of this result further in the conclusions.

Changes in reserve holding across balancing regions in the high coordination scenarios are significant. Figure 2 shows reserve holdings by resource scenario and by reserve type for each of the five balancing

regions. In the medium coordination scenarios, reserve holdings for all balancing regions decline as a result of reserve sharing. In the high coordination scenarios, reserve holdings decline in most balancing regions but increase in the SOCO and, to a lesser extent LGEE, balancing regions as PLEXOS shifts reserves to regions where doing so will further minimize regional operating costs. Most of these changes are in load following reserves and are likely related to the capacity mix of different balancing areas. The substantial decline in reserves in AECI is notable because AECI is the smallest balancing region (12 GW peak) and has the highest share of solar and wind generation (87%-98%) in the study area. This suggests that smaller balancing regions can benefit not just by pooling forecast error reserves (still holding most reserves locally), but also by being able to hold a significant portion of reserves to manage forecast error in other regions. Deeper reserve sharing like this may be difficult to achieve without a centralized system operator.




Notes: SCT is spinning contingency reserves, NSCT is non-spinning contingency reserves, REG is regulation reserves, SLF is spinning load following reserves, NSLF is non-spinning load following reserves. LC, MC, and HC are the low, medium, and high coordination scenarios. A TW-h is a terawatt of reserve capacity scheduled for one hour. Annual reserve provision here is the sum of hourly reserve capacity over the year.

Figure 2. Annual day-ahead reserve provision by reserve type for each balancing area in each resource and coordination scenario

The supply of reserves by different generation technologies changes across both resource and coordination scenarios (Figure 3). Storage displaces coal and gas generation for spinning and regulation

reserves beginning in the MPHS scenario. Gas generation provides most non-spinning reserves in all scenarios. The two largest changes in reserve supply from the low and medium to high coordination scenarios are (1) a shift from combined cycle gas (GCC) to gas combustion turbines (GCT) for non-spinning load following reserves in all resource scenarios, and (2) a shift from hydropower to battery storage for spinning reserves. The former may be the result of PLEXOS freeing up GCC capacity to provide energy. The latter is likely due to our assumption that hydropower has higher variable operating and maintenance costs than battery storage.



Notes: Resource technologies: B is battery, 2 to 10 is battery duration, PHS is pumped hydropower storage, OFWD is offshore wind, ONWD is onshore wind, HYD is hydropower, BIO is biomass power, OGS is oil and gas steam, GCC is gas combined cycle, GCT is gas combustion turbine. Reserve categories: SCT is spinning contingency reserves, NSCT is non-spinning contingency reserves, REG is regulation reserves, SLF is spinning load following reserves, NSLF is non-spinning load following reserves. LC, MC, and HC are the low, medium, and high coordination scenarios.

Figure 3. Share of reserve provision by technology and reserve type for each resource and coordination scenario

Electricity storage clearly plays an important role in the results. Battery storage accounts for nearly all spinning and regulation reserve provision beginning in the MPHS scenarios, and increasing battery storage capacity (MPLS to MPHS scenario, HPLS to HPHS scenario) increases the share of battery storage in reserve provision. Reserve costs decline from low storage (LS) to high storage (HS) scenarios in nearly all scenarios (Table 4) and are about the same in the one exception (HPLS to HPHS). Declining

reserve costs also suggest that the value of reserve sharing should decline with higher levels of storage, which is indeed the case in all but one scenario (MPHS low to medium coordination has higher reserve cost savings than MPLS low to medium coordination). This result also suggests that, in shifting to electricity systems with higher amounts of storage, the value of reserve sharing will increasingly depend on storage opportunity costs that are not easily captured in production cost models. Storage also presents new challenges for reserve modeling, because high levels of storage shift the reserve scheduling problem from unit commitment toward state-of-charge management, requiring more coordination between stages in multi-stage models.

With the transmission capacity assumptions used in this study, transmission limits have a minimal impact on reserve sharing. The most important impact of transmission limits is for sharing of non-spinning load following reserves, because they are the largest reserve requirement. Similarly, because of the two-settlement approach described in Section 3.2, real-time prices have a limited impact on the results. For this kind of analysis, including intraday adjustment dynamics may not be worth the effort, as in most cases it does not meaningfully change the results.

5. Discussion: Decentralized Sharing of Forecast Error Reserves

The analysis in the previous sections describes the effects of forecast error reserve sharing, but it does not provide more detailed insight on how reserve sharing could be implemented in practice. More decentralized forecast error reserve sharing, such as in the medium coordination scenarios here, may not require a centralized system operator such as an RTO. However, day-ahead forecast error reserve sharing without a centralized system operator requires supporting rules, processes, and software. There are multiple approaches to designing a decentralized, multi-utility program for sharing forecast error reserves, which could draw on the design of contingency reserve sharing arrangements – the Northwest Power Pool’s Reserve Sharing Program, for instance (NWPP, 2023).

In practice, the most straightforward approach would likely consist of four elements: (a) calculating a total forecast error reserve requirement for the multi-utility region, which will be lower than the sum of utilities’ individual requirements; (b) allocating the total reserve requirement to individual utilities based on their day-ahead hourly solar and wind forecasts; (c) establishing reserve deployment obligations for each utility but allowing utilities to voluntarily request deployment; and (d) settling reserve deployment (energy) at a market-based or common administratively set price. These four elements are consistent with the approach used in this paper.

Forecast error reserve requirements that are pooled across utilities are lower because forecast errors are weakly correlated among utilities, but pooling does not change the forecast error reserve needs for individual utilities. This means that if the utilities choose to pool reserves, they must also share reserve energy. Utilities will need others in the pool to deploy reserve energy in some cases to manage their forecast errors. Although the mechanics of this kind of reserve sharing are often assumed away in

modeling analyses, exploring how they might work in practice is important for gauging their institutional and political feasibility.

There are several ways to allocate pooled reserve requirements and reserve deployment obligations among utilities. For instance, consider again the scenario described in Section 2 (Background): in some time interval, Utility A has a 500 MW forecast error reserve requirement and a day-ahead solar forecast of 2,500 MW, Utility B has a 1,000 MW forecast error reserve requirement and a day-ahead solar forecast of 5,000 MW, and both utilities have a pooled 1,300 MW forecast error reserve requirement. The shared reserve in this case would be 200 MW ($= 1,500 \text{ MW} - 1,300 \text{ MW}$). Each utility would need to plan to deploy some of this shared reserve to the other, to cover any shortfalls in forecast error reserves. In each interval, the maximum deployment obligation for each utility would be the difference between the other utility’s individual and shared reserve requirements, to ensure that each utility has sufficient resources available to maintain the same level of forecast error coverage.

The two utilities could allocate the total reserve requirement (1,300 MW) or the shared reserve (200 MW), with different implications for reserve requirements and reserve deployment obligations (summarized in Table 5). Allocating the total reserve requirement (1,300 MW) to both utilities based on forecasted solar generation would result in a 433 MW ($= 1,300 \text{ MW} \times 2500 \text{ MW}/7500 \text{ MW}$) reserve requirement for Utility A and an 867 MW ($= 1,300 \text{ MW} \times 5000 \text{ MW}/7500 \text{ MW}$) requirement for Utility B. These requirements imply a 133 MW ($= 1,000 \text{ MW} - 867 \text{ MW}$) maximum deployment obligation for Utility A, to cover the difference between Utility B’s individual and shared reserve requirements, and a 67 MW ($= 500 \text{ MW} - 433 \text{ MW}$) deployment obligation for Utility B, to cover Utility A. Alternatively, allocating the shared reserve (200 MW) based on forecasted solar generation would result in a 67 MW deployment obligation ($= 200 \text{ MW} \times 2500 \text{ MW}/7500 \text{ MW}$) for Utility A and a 133 MW ($= 200 \text{ MW} \times 5000 \text{ MW}/7500 \text{ MW}$) obligation for Utility B. These obligations imply a reserve requirement of 367 MW ($= 500 \text{ MW} - 133 \text{ MW}$) for Utility A and 933 MW ($= 1,000 \text{ MW} - 67 \text{ MW}$) for Utility B. In summary, while both approaches benefit all utilities, the first benefits utilities with more wind and solar generation and the second benefits utilities with less.

Table 5. Summary of total reserve requirement and shared reserve strategies for allocating forecast error reserves

	Allocate total reserve requirement		Allocate shared reserve	
	Reserve requirement	Maximum reserve deployment obligation	Reserve requirement	Maximum reserve deployment obligation
Utility A (2,500 MW solar forecast)	433 MW	133 MW	367 MW	67 MW
Utility B (5,000 MW solar forecast)	867 MW	67 MW	933 MW	133 MW

With more than two entities in a forecast error reserve sharing pool, the allocation math quickly becomes complex and there is a need for clear rules and automated reserve setting and communication. The approach advocated above (“(b) allocating the total reserve requirement to

individual utilities based on their day-ahead hourly solar and wind forecasts”) prioritizes simplicity. Clear rules and automation determine what happens during lower probability events when larger forecast errors occur in tandem across utilities and, relatedly, the extent to which utilities want to share non-spinning reserves.

In the medium coordination scenario, we limit forecast error reserve sharing to spinning reserves. Non-spinning reserves create additional complexity because they involve coordination of planned outage schedules and resource adequacy. However, as this analysis has shown, sharing non-spinning forecast errors can create significant value. As long as reserve obligations are clear in advance and there are clear rules and mechanisms for implementing these obligations, decentralized sharing of non-spinning forecast error reserves should also be feasible. The two other elements in the high coordination scenario – sharing regulation reserves and optimizing the location of reserves across a larger region – require more centralized system operations.

Implementation of decentralized forecast error reserve sharing likely requires a dedicated organization, to document and maintain rules, develop and maintain software, ensure transmission availability, and settle accounts among pool participants. For the Southeast U.S., for instance, this third-party organization could be the entity charged with managing market software and settlement for SEEM. However, pool members could also undertake this task on a rotating basis. In most countries, including the U.S., multi-state or multi-province reserve sharing agreements would require the political support of local regulators and governments.

6. Conclusions

Growth in solar and wind generation may require utilities and other system operators to hold additional reserves to manage forecast error. This study examined the value of forecast error reserve sharing among utilities in the Southeastern U.S. in scenarios with higher amounts of solar and wind generation. It found that average day-ahead forecast error reserve requirements increased significantly with higher levels of solar and wind generation, from 14% of load (23% solar, 10% wind) to 25% of load (46% solar, 19% wind). Average reserve requirements increased linearly with growth in solar and wind generation capacity, at 6%-10% of total capacity depending on reserve sharing assumptions. Sharing forecast error reserves across a region significantly reduced the amount of reserves that individual system operators needed to hold to manage forecast error. Day-ahead load following and regulation reserve requirements declined by 6%-8% (1-2 percentage points, as a share of load) in scenarios in which balancing regions shared spinning load following reserves and by 25%-29% (by 4-7 percentage points, as a share of load) in scenarios in which they shared spinning and non-spinning load following reserves and regulation reserves.

The study found that annual reserve cost savings from reserve sharing ranged from \$0.09 to \$1.24 billion (\$0.12-\$1.68/MWh, 2022\$), depending on resource portfolios and the extent of reserve sharing. Lower energy and reserve prices in scenarios with higher levels of solar, wind, and storage tended to

reduce the value (reserve cost savings) from reserve sharing. More frequent scarcity pricing tended to increase the value of reserve sharing. Across scenarios, the highest 100 price hours accounted for 62%-99% of total reserve costs, underscoring the importance of scarcity prices in the results. Annual cost savings in the high solar low storage scenario (\$620 million, \$0.84/MWh, 64% solar and wind generation) scenario suggests that savings from forecast error reserve sharing may still be substantial even with relatively high levels of solar and wind generation.

With higher levels of variable generation and electricity storage, capacity and energy adequacy begin to be a larger consideration in reserve costs than generator operating constraints and production costs. The importance of scarcity prices and non-spinning reserves in our results underscores this point. If a system has adequate capacity, reserve prices will be low and the value of reserve sharing will be low. If it does not, reserve prices and costs will be very high in hours in which the system is capacity constrained. The importance of scarcity costs in this study is consistent with other studies (Levin and Botterud, 2015; Vijay et al., 2017; Frew et al., 2021). It suggests the need to incorporate scarcity pricing in modeling the benefits of reserve sharing and other coordination arrangements among system operators. It also suggests that only focusing on operational benefits may not capture the most important values of coordination.

Capacity adequacy tradeoffs—in this case between reserve sharing and building additional local resources—are generally better captured in capacity expansion models than production simulation models. This will be particularly true for more conservative system operators that wish to hold more non-spinning reserves to cover lower probability forecast error events (higher values of σ , see Section 2). However, in practice, system operators, load serving entities, generators, and other resource planners do not presently undertake joint capacity expansion planning to ensure economically efficient levels of forecast error-related reserves. In addition, capacity expansion models have only more recently begun to tackle the endogeneity challenges of forecast error reserves (reserves are a function of resource decisions), and more work in this area is needed to determine where more detail is important, and potentially to develop new heuristics. Future work could explore the design, implementation, and modeling dimensions of joint capacity expansion that supports forecast error reserve sharing.

The approach used in this study – national modeling with endogenous energy-reserve interactions – highlights the difficulties of meaningfully isolating geographic areas and markets in this kind of analysis. The market coordination required to enable sharing of load following reserves will tend to lead to larger changes in dispatch and wholesale energy prices, both within and across regions. Reserve prices and opportunity costs are strongly influenced by energy prices. Both of these realities suggest that evaluating the value of reserve sharing by looking at changes in local production costs will provide incomplete, and potentially misleading, results. It may not always be practical to do more comprehensive analysis, but it is important to be cognizant of potential interactions with electricity systems and markets outside of the boundaries of more narrowly defined analysis. Alternatively, however, more complex models may also not provide clear intuition on how reserve sharing affects system operations or on the tradeoffs and risks for the design of reserve sharing mechanisms. Simpler

models, with a small number of generators, transmission constraints (if any), and regions, can help to provide intuition on how changes in assumptions affect the results and can be a useful complement to more complex models.

Although the focus in this discussion on decentralized reserve sharing has focused on U.S. utilities, this study may be relevant for U.S. RTOs, provincial and national grid companies in China, transmission system operators (TSOs) in Europe, state and national load dispatch centers in India, and other system operators in other countries. In the U.S., for instance, how should PJM incorporate wind forecast uncertainty in MISO into its day-ahead scheduling process? Can multiple regional system operators reduce the amount of capacity and energy held in reserve day-ahead to manage forecast error by pooling forecast errors across multiple regions? The results of this study suggest that, even with declining wholesale energy and reserve prices, there may be significant benefits to doing so.

7. References

- Blair, N., N. DiOrio, J. Freeman, P. Gilman, S. Janzou, T. Neises, M. Wagner. 2018. System Advisor Model (SAM) General Description (Version 2017.9.5). NREL/TP-6A20-70414. Golden, NREL.
- Chang, J.W., J.P. Pfeifenberger, C. Onur, A. Mariko, G. Aydin, K. Van Horn, P. Cahill, L. Regan, C. McIntyre. 2016. *Senate Bill 350 Study Volume V: Production Cost Analysis*. Boston: The Brattle Group
- Doherty, R., M. O'Malley. 2005. "A new approach to quantify reserve demand in systems with significant installed wind capacity." *IEEE Transactions on Power Systems* 20(2): 587-595.
- Ela, E., M. Milligan, B. Kirby. 2011. *Operating Reserves and Variable Generation*. U.S. National Renewable Energy Laboratory (NREL) Technical Report NREL/TP-5500-51978. Golden: NREL.
- Energy and Environmental Economics (E3). 2011. *WECC EDT Phase 2 EIM Benefits Analysis & Results*. San Francisco: E3.
- E3. 2013. *PacifiCorp-ISO Energy Imbalance Market Benefits*. San Francisco: E3.
- EnerNex. 2011. *Eastern Wind Integration and Transmission Study*. NREL Subcontract Report NREL/SR-5500-47078. Golden: NREL.
- Focken, U., M. Lange, K. Mönnich, H.P. Waldl, H.G. Beyer, A. Luig. 2002. "Short-term prediction of the aggregated power output of wind farms—a statistical analysis of the reduction of the prediction error by spatial smoothing effects." *Journal of Wind Engineering and Industrial Aerodynamics* 90: 231-246.
- Frew, B., G. Brinkman, P. Denholm, V. Narwade, G. Stephen, A. Bloom, J. Lau. 2021. "Impact of operating reserve rules on electricity prices with high penetrations of renewable energy." *Energy Policy* 156, 112443.
- GE Energy. 2010. *Western Wind and Solar Integration Study*. NREL Subcontract Report NREL/SR-550-47434. Golden: NREL.
- Hale, E., E. Zhou. 2021. *Absorbing the Sun: Operational Practices and Balancing Reserves in Florida's Municipal Utilities*. NREL/TP-6A20-79385. Golden: NREL.
- Ho, J., J. Becker, M. Brown, P. Brown, I. Chernyakhovskiy, S. Cohen, W. Cole, S. Corcoran, K. Eurek, W. Frazier, P. Gagnon, N. Gates, D. Greer, P. Jadun, S. Khanal, S. Machen, M. Macmillan, T. Mai, M. Mowers, C. Murphy, A. Rose, A. Schleifer, B. Sergi, D. Steinberg, Y. Sun, E. Zhou. 2021. *Regional Energy Deployment System (ReEDS) Model Documentation: Version 2020*. NREL Technical Report NREL/TP- 6A20-78195. Golden, CO: NREL.
- Hodge, B.-M., D. Lew, M. Milligan, H. Holttinen, S. Sillanpää, E. Gómez-Lázaro, R. Scharff, L. Söder, X. G. Larsén and G. Giebel, D. Flynn, J. Dobschinski. 2012. *Wind Power Forecasting Error Distributions: An International Comparison*. NREL/CP-5500-56130. Golden, CO: NREL.
- Hummon, M., P. Denholm, J. Jorgenson, D. Palchak, B. Kirby, O. Ma. 2013. *Fundamental Drivers of the Cost and Price of Operating Reserves*. NREL/TP-6A20-58491. Golden, CO: NREL.

- Jiang, H., J. Xu, Y. Sun, S. Liao, D. Ke, Y. Jiang, B. Tang. 2018. "Dynamic reserve demand estimation model and cost-effectivity oriented reserve allocation strategy for multi-area system integrated with wind power." *IET Generation, Transmission & Distribution* 12(7): 1606-1620.
- Katzenstein, W., E. Fertig, J. Apt. 2010. "The variability of interconnected wind plants." *Energy Policy* 38(8): 4400-4410.
- Kim, J.H., F. Kahrl, A. Mills, R. Wisler, C. Montañés, W. Gorman. 2023. "Economic evaluation of variable renewable energy participation in U.S. ancillary services markets." *Utilities Policy* 82: 101578.
- King, J., B. Kirby, M. Milligan, S. Beuning. 2011. "Flexibility Reserve Reductions from an Energy Imbalance Market with High Levels of Wind Energy in the Western Interconnection." NREL Technical Report NREL/TP-5500-52330. Golden, CO: NREL.
- Krad, I., E. Ibanez, W. Gao. 2016. "A Comprehensive Comparison of Current Operating Reserve Methodologies." *Proceedings of the IEEE Power Engineering Society Transmission and Distribution Conference and Exposition*.
- Levin, T., A. Botterud. 2015. "Electricity market design for generator revenue sufficiency with increased variable generation." *Energy Policy* 87: 392-406.
- Lopez, A., P. Pinchuk, M. Gleason, W. Cole. 2024. *Solar Photovoltaics and Land-Based Wind Technical Potential and Supply Curves for the Contiguous United States: 2023 Edition*. NREL Technical Report NREL/TP-6A20-87843. Golden, CO: NREL.
- Mauch, B., J. Apt, P.M.S. Carvalho, P. Jaramillo. 2013. "What day-ahead reserves are needed in electric grids with high levels of wind power?" *Environmental Research Letters* 8: 034013.
- Northwest Power Pool (NWPP). 2023. *Northwest Power Pool Reserve Sharing Program Documentation*. Approved 2-2-2023.
- Ortega-Vazquez, M.A., D.S. Kirschen. 2009. "Estimating the Spinning Reserve Requirements in Systems With Significant Wind Power Generation Penetration." *IEEE Transactions on Power Systems* 24(1): 114-124.
- Samaan, N.A., Y.V. Makarov, T.B. Nguyen, R. Diao. 2017. "Balancing Authority Cooperation Concepts to Reduce Variable Generation Integration Costs in the Western Interconnection: Consolidating Balancing Authorities and Sharing Balancing Reserves." In Du, P., R. Baldick, A. Tuohy (Eds.), *Integration of Large-Scale Renewable Energy into Bulk Power Systems*. Springer.
- Vijay, A., N. Fouquet, I. Staffell, A. Hawkes. 2017. "The value of electricity and reserve services in low carbon electricity systems." *Applied Energy* 201: 111-123.

APPENDIX

This appendix covers:

- ReEDS input assumptions and installed capacity by scenario
- Solar and wind forecast error calculations
- Solar and wind forecast error distributions

A.1 ReEDS Inputs Assumptions and Installed Capacity by Scenario

We used the ReEDS capacity expansion model to develop resource portfolios for our five resource scenarios:

- Base solar PV base storage (BPBS)
- Medium solar PV low storage (MPLS)
- Medium solar PV high storage (MPHS)
- High solar PV low storage (HPLS)
- High solar PV high storage (HPHS)

To develop the portfolios, we adjusted five key assumptions, shown in Table A-1. Costs in ReEDS were based on NREL’s Annual Technology Baseline (ATB) 2022 projections. The high storage (HS) scenarios used more aggressive (“advanced”) cost estimates from ATB. The high solar (HP) scenarios used a carbon tax, beginning with \$46/tCO₂ in 2022 and increasing to \$88/tCO₂ by 2035 (2004\$). The carbon tax was applied to investment decisions in ReEDS but was not carried over to PLEXOS modeling.

Table A-1. Key ReEDS input assumptions for developing scenarios

Assumption	Resource Scenario				
	BPBS	MPLS	MPHS	HPLS	HPHS
Battery costs	conservative	Conservative	Advanced	conservative	advanced
Carbon tax	none	None	None	yes	yes
Coal retirements	accelerated	Accelerated	Accelerated	n/a*	n/a*
DPV adoption	mid case	low-cost case	low-cost case	low-cost case	mid case
Utility-scale PV costs	moderate	Advanced	Advanced	moderate	advanced

We used assumptions on distributed PV (DPV) adoption and utility-scale PV costs to adjust and differentiate the higher solar scenarios (MP and HP). DPV adoption was based on projections from NREL’s dGen model.⁷ Adoption scenarios varied with the cost of DPV, which were related to cost assumptions for utility-scale PV. For DPV adoption, we used low-cost cases for all higher solar scenarios except for HPHS, for which we use a mid-case assumption to enable more utility-scale PV and better differentiate it from HPLS. For the HPLS case, we used moderate PV costs to better differentiate it from HPHS but used advanced PV costs in all other solar scenarios. In all scenarios, we assumed accelerated

⁷ See <https://www.nrel.gov/analysis/dgen/>.

retirements of coal units, consistent with trends within the Southeast region. Figure A-1 shows the ATB capital cost envelopes used in the analysis.

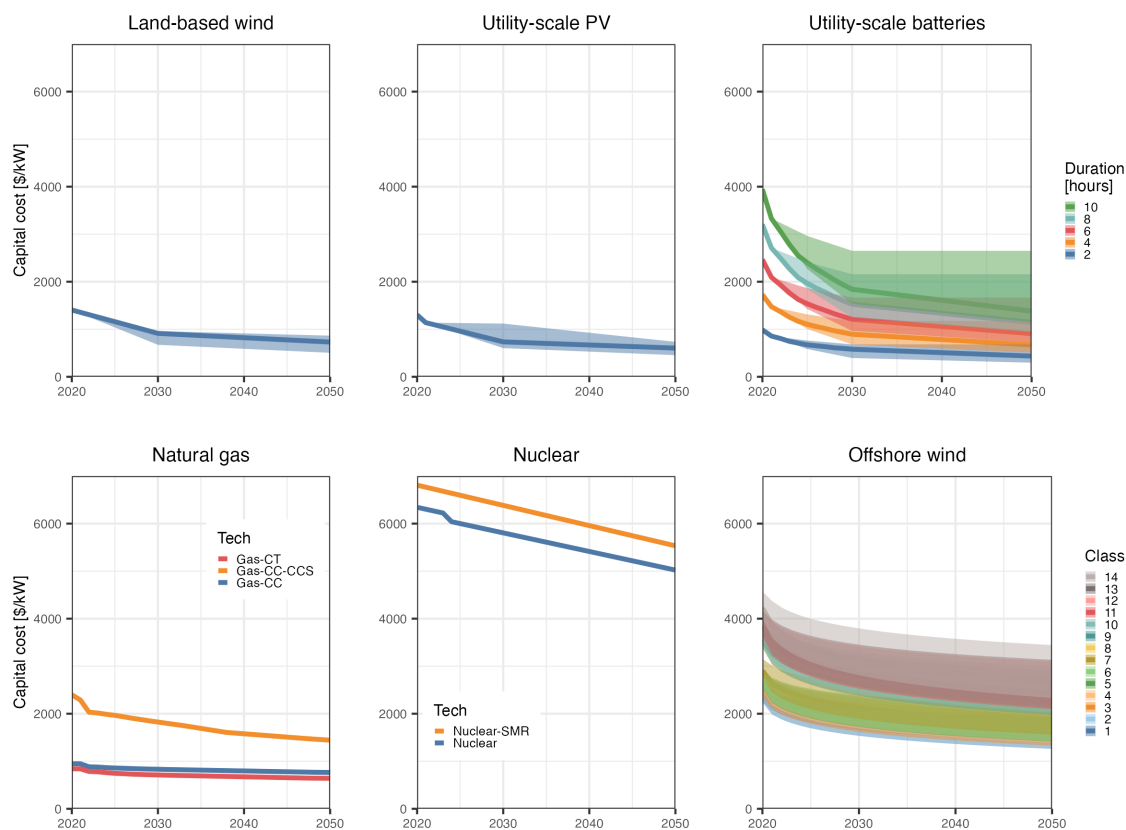


Figure A-1. ATB capital cost assumptions used in ReEDS (2004\$)

All scenarios included representation of Inflation Reduction Act (IRA) incentives. Utility PV projects are now eligible for up to a \$28.60/MWh production tax credit (PTC) and can choose between the investment tax credit (ITC) and PTC. We assumed that solar projects choose the PTC given its higher lifetime value.⁸ Incorporating IRA incentives into the model significantly increased the amount of “baseline” solar deployment in ReEDS. For instance, in the BPBS scenario, solar generation in ReEDS increased from 12% to 22% of total generation in 2035 after modeling IRA provisions. Not all IRA incentives were included in the ReEDS model (e.g., the 45V tax credit for hydrogen).

In the ReEDS model, we assumed that each balancing region must meet its own demand plus planning reserve requirements, rather than allowing for planning reserve sharing. However, we allowed balancing regions to use imports to meet resource adequacy needs, subject to hurdle rates. We used a hurdle rate of \$7.5/MWh (2004\$) to capture transmission friction between balancing areas.

⁸ For more detail on the representation of IRA provisions in ReEDS and the impact on renewable deployment see <https://www.nrel.gov/docs/fy23osti/85242.pdf>.

Figure A-2 shows the ReEDS portfolios that result from these assumptions. It shows installed capacity for the Southeast region in the BPBS and changes in installed capacity in the higher solar and storage scenarios. Although the changes in installed capacity in the higher solar and storage scenarios were largely additional solar and battery storage and retirements of coal generation, Figure A-2 also illustrates the interactions among solar, wind, and storage. In the MP scenarios, solar and wind were substitutes: higher levels of solar capacity reduced wind capacity relative to the BPBS scenario. In the HP scenarios, solar and wind were complements: the carbon tax needed to achieve higher levels of solar also led to higher levels of wind.




Figure A-2. Installed capacity in the BPBS scenario and change in installed capacity in other scenarios relative to the BPBS scenario, Southeast region

ReEDS is a national model. For consistency, we imposed similar assumptions on the rest of the U.S. As a result, changes in installed capacity in neighboring regions (Florida, MISO, PJM, SPP) resembled those in the Southeast region (Figure A-3). Figure A-4 provides a visualization of the spatial resolution of the zones used to represent transmission and load balancing in ReEDS for the Southeast.

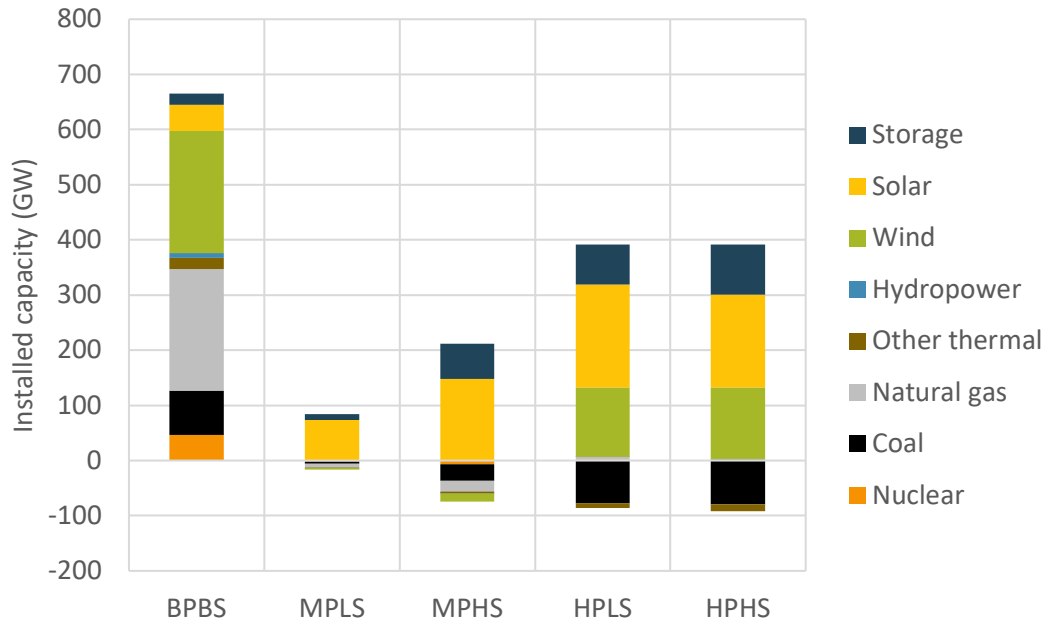



Figure A-3. Installed capacity in the BPBS scenario and change in installed capacity in other scenarios relative to the BPBS scenario, neighboring regions



Note: Numbered areas are ReEDS transmission/load balancing zones and colored regions are the Southeast balancing regions in focus for this study.

Figure A-4. Spatial representation of the Southeast in the ReEDS model

A.2 Solar and Wind Forecast Error Calculations

Persistence forecasts imply that the system operator predicts that solar clear-sky ratios and wind generation in the next period will be the same as in this period. For day-ahead persistence forecasts, the strict interpretation of a persistence forecast implies that the system operator forecasts energy (E) at time t the following day ($d+1$) to be either the same as energy at time t today (d). However, this approach ($E_{d+1,t} = E_{d,t}$) tends to lead to forecast errors that are unrealistically large for a 2035 analysis. For instance, in the HPHS scenario a simple day-ahead persistence forecast would have a mean average percentage error (MAPE) of 41% (8.4% normalized by solar installed capacity). For comparison, the California Independent System Operator's (CAISO's) capacity normalized, monthly day-ahead solar forecast error between 2021 and 2023 was 2-4%.⁹ An alternative approach is to use hourly persistence forecasts for the next day ($E_{d+1,t} = E_{d+1,t-1}$) as a proxy for the day-ahead forecast. Considering improvements in forecasting that could be possible by 2035, we used the hourly ($E_{d+1,t} = E_{d+1,t-1}$) rather than the daily ($E_{d+1,t} = E_{d,t}$) persistence approach in this analysis.

For day-ahead forecasts with an hourly timestep, hourly average energy (MWh/h) forecast in the next hourly timestep (e.g., hour ending [HE] 14:00) will be hourly average energy in the previous hourly timestep (e.g., HE 13:00). For real-time forecasts with a 15-minute timestep, average 15-minute energy (MW) in the next 15-minute timestep (e.g., period ending 14:30) will be the average 15-minute energy in the previous 15-minute timestep (e.g., period ending 14:15).

Forecast error is defined as the difference between an observed (actual) value and a forecast value. For 15-minute forecast errors, we calculated forecast errors as the difference between a value in time $t+1$ (actual) and the value at time t (forecast). For hourly forecast errors, we incorporated intra-hour variability by calculating forecast error as the maximum difference between 15-minute values in hour $h+1$ (actuals) and the value in hour h (forecast). Table A-2 provides an illustrative example for both 15-minute and hourly forecast error calculations.

⁹ See CAISO, 2023, "Market Performance and Planning Forum," <https://www.caiso.com/Documents/Presentation-MarketPerformancePlanningForum-Jun29-2023.pdf>.

Table A-2. Illustration of 15-minute and hourly forecast error calculations

Time	A. 15-Minute Wind	B. Hourly Wind	C. 15-Minute Persistence Forecast	D. Hourly Persistence Forecast	E. 15-Minute Forecast Error (A – C)	F. Hourly Forecast Difference (A – D)	G. Hourly Forecast Error (max[F])
13:00	7966	8661	7314	6336	652	1630	3019
13:15	8429	8661	7966	6336	463	2093	
13:30	8892	8661	8429	6336	463	2556	
13:45	9356	8661	8892	6336	463	3019	
14:00	9819	9939	9356	8661	463	1158	1398
14:15	9899	9939	9819	8661	80	1238	
14:30	9979	9939	9899	8661	80	1318	
14:45	10059	9939	9979	8661	80	1398	

Notes: Data are from the aggregated Duke balancing region wind profiles used in this study. Hourly persistence forecasts for 13:00-14:00 are from the previous period. Hourly forecast errors (G) are the maximum difference (F) between an hourly persistence forecast (D) and 15-minute observed values (A) in an hour.

Reserves for managing forecast error between day-ahead scheduling and real-time dispatch (load following reserves) are distinct from those required to manage forecast error within the dispatch interval (regulation reserves). For load following reserves, the system operator will need to make sure it has sufficient generation available in reserve to address the difference between its day-ahead forecast (e.g., 8,661 MW in hour beginning [H.B.] 14:00) and its actual 15-minute dispatch needs. For instance, the system operator expects to have 8,661 MW of wind in each period in H.B. 14:00 but has 10,059 MW in period 14:45-15:00, requiring a maximum downward reserve of 1,398 MW to be able to accommodate higher-than forecasted wind generation during H.B. 14:00. During real-time operations, the system operator expects to have 9,979 MW of wind during 14:45-15:00 but actually has 10,059 MW, requiring 80 MW of downward regulation reserve.

For wind, our forecasts were based on profile (generation) data. For solar, our forecasts were based on clear-sky persistence ratios. Our load following solar forecast errors in hour h+1 ($LF.SFE_{h+1}$) were calculated as

$$LF.SFE_{h+1} = \max\left(\frac{FS_{t,h+1}}{FSC_{t,h+1}} - \frac{HS_h}{HSC_h}\right) \times \max(FSC_{t,h+1})$$

where $FS_{t,h+1}$ is the 15-minute solar profile in 15-minute interval t in hour h+1, $FSC_{t,h+1}$ is the 15-minute solar clear-sky profile in 15-minute interval t in hour h+1, HS_h is the hourly average solar profile in hour h, and HSC_h is the hourly average clear-sky solar profile in hour h. Clear-sky ratios must be multiplied by the clear-sky value to ensure that the final value is in MW terms. We used the maximum 15-minute clear-sky value across the 15-minute t intervals in hour h+1 for this normalization.

Persistence solar forecast errors have issues in the first and last periods each day when the solar profiles move from a zero value to a positive value. These values can be smoothed, but in the interest of simplicity we truncated solar forecast error values in the first and last period each day.

A.3 Solar and Wind Forecast Error Distributions

The resulting forecast error distributions were moderately skewed and leptokurtic. The solar forecast error distributions are positively skewed (bias toward over-forecast) and the wind forecast error distributions are negatively skewed (bias toward under-forecast), but both values are within the range of -1 and -0.5 and +1 and +0.5 that would be considered moderately skewed. The high excess kurtosis values for both solar and wind forecast error distributions indicate they are significantly more peaked than a normal distribution would be. The figures and tables below show day-ahead (load following) solar and wind forecast error probability density functions and skew and kurtosis statistics by resource scenario.




Figure A-5. Probability density function for load following solar forecast errors

Table A-3. Kurtosis and skew statistics for solar forecast errors by scenario

	Scenario				
	BPBS	MPLS	MPHS	HPLS	HPHS
Excess kurtosis	4.48	3.86	4.87	4.93	4.88
Skew	0.77	0.69	0.80	0.79	0.77




Figure A-6. Probability density function for load following wind forecast errors

Table A-4. Kurtosis and skew statistics for wind forecast errors by scenario

	Scenario				
	BPBS	MPLS	MPHS	HPLS	HPHS
Excess kurtosis	3.15	3.43	3.45	3.31	3.47
Skew	-0.62	-0.65	-0.66	-0.71	-0.71

Logit and log transformations of the original solar and wind data did not allow us to reject the null hypothesis that the forecast error data are normally distributed. However, normality is not strictly necessary for the method used in this study. As an alternative, we looked at the approximate coverage for standard deviations in the individual solar and wind forecast error distributions, as a proxy for forecast error coverage in their joint distribution.

Because the forecast error data are skewed, coverage for under-forecasts (actual solar < forecast, negative values) and over-forecasts (actual solar > forecast, positive values) must be normalized by the zero point on the cumulative distribution function. For under-forecasts, the forecast error coverage of s standard deviations will be

$$P(s \times \sigma \leq x \leq 0) = \frac{CDF(0) - CDF(s \times \sigma)}{CDF(0)}$$

For over-forecasts, it will be

$$P(0 \leq x \leq s \times \sigma) = 1 - \frac{1 - CDF(0) - (1 - CDF(s \times \sigma))}{1 - CDF(0)}$$

Table A-5 shows the forecast error coverage of different standard deviations in each scenario. For reference, 1, 2, and 3 standard deviations in a normal distribution will cover 68%, 95%, and 99.7% of values. In general, forecast error coverage for under-forecasts with our approach will be on par or higher than a normal distribution for one standard deviation, but lower for two and three. It will likely be lower for all over-forecasts.

Table A-5. Solar and wind forecast error coverage for different standard deviations and scenarios

Standard Deviations	Scenario				
	BPBS	MPLS	MPHS	HPLS	HPHS
Solar over-forecast					
1	0.58	0.58	0.59	0.59	0.59
2	0.84	0.85	0.84	0.84	0.83
3	0.944	0.951	0.942	0.944	0.944
Solar under-forecast					
1	0.68	0.68	0.69	0.69	0.70
2	0.92	0.92	0.92	0.93	0.93
3	0.982	0.987	0.985	0.985	0.985
Wind over-forecast					
1	0.80	0.79	0.80	0.81	0.82
2	0.96	0.96	0.96	0.96	0.96
3	0.994	0.994	0.994	0.993	0.992
Wind under-forecast					
1	0.77	0.77	0.77	0.76	0.76
2	0.93	0.93	0.93	0.90	0.91
3	0.973	0.971	0.971	0.967	0.970

In the analysis, we assumed that system operators cover three standard deviations of load following forecast errors ($\sigma = 1$ spinning, $\sigma = 2$ non-spinning). With higher levels of solar and wind, the reserves required to cover larger shares of forecast errors are substantial. Table A-6 shows standard deviations of solar and wind load following forecast errors, and a joint standard deviation assuming errors add in quadrature, as an illustration. Actual reserve requirements in our analysis will differ from these because we assume that errors and reserve requirements are a function of solar and wind generation level, but this still gives a sense of scale. In the HP scenario, each standard deviation of forecast error coverage adds roughly 4-5 GW of reserves. This suggests that some form of reserve sharing, whether formal or emergency, will be necessary to deal with less frequent, larger magnitude forecast errors.

Table A-6. Standard deviations for solar and wind day-ahead forecast errors by scenario

Resource	Scenario				
	BPBS	MPLS	MPHS	HPLS	HPHS
Solar	1790	2365	2897	3668	3766
Wind	1454	1187	1171	2824	2603
Joint	2306	2646	3125	4629	4578