

The Performance of U.S. Wind and Solar Generators

Richard Schmalensee*

ABSTRACT

Using data on hourly outputs and spot prices for a sample 25 wind and nine solar generating plants covering all seven U.S. ISOs for 2011 and up to 12 adjacent months, this study examines capacity factors, average output values, and several aspects of intermittency. Most performance measures studied vary substantially within and between ISOs, and some vary substantially over time. Implications for research, market design, and policies to support renewable generation are briefly discussed.

Keywords: wind, solar, renewable, VER, generation

* Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139. Email: rschmal@mit.edu.

The author is indebted to John Beuchler, Paul Edmundson, Isabel Flores, Shannon Hann, Eric Hsia, Craig Kazin, Kevin Kirby, Ryan Leonard, Clyde Loutan, Lisa Morelli, Pamela Quittner, Christopher Russell, Carrie Simpson, Gordon van Welie, and Stephen Whitley for their various roles in providing the data used in this study and for helpful conversations. I am indebted for useful comments to Severin Borenstein, Joseph Eto, Lion Hirth, William Hogan, Paul MacAvoy, Paul Joskow, Matthew White, two excellent referees, and participants in workshops of the MIT Center for Energy and Environmental Policy Research and the Energy Institute @ Haas (Berkeley). Errors and opinions are entirely my own, of course

Published as “The Performance of U.S. Wind and Solar Generators.” Schmalensee, Richard. *Energy Journal* Vol. 37, No. 1 (2016): 123-151. Copyright 2016 by the IAEE. All rights reserved.
DOI: 10.5547/01956574.37.1.rsch

1. Introduction

The rapid growth of wind and solar electricity generation globally in recent decades has been driven in large part by a diverse set of subsidies and regulations. In the U.S. these have involved all levels of government, and they have primarily rewarded investment in generation facilities and megawatt-hours (MWh) generated.¹ The main federal subsidy for solar generation has been a 30% investment tax credit, and both solar and wind facilities have been eligible for accelerated depreciation. These and similar investment-based state and local tax credits clearly provide weaker incentives to minimize production costs than generation-based subsidies.

The main federal subsidy for wind generation has been a \$23/MWh production tax credit (PTC).² At the state level, 29 states and the District of Columbia have Renewable Portfolio Standard (RPS) programs that generally require affected utilities to purchase specified quantities of renewable energy credits (RECs), which are produced in proportion to their outputs by certified generating facilities powered by renewable energy. To be certified in half of these programs, a generator must deliver power in the state in question, and the other 15 programs either limit certification to nearby generators or provide incentives for in-state generation.³ Thus the PTC and each state's RPS program treat all certified wind-generated MWh as equally valuable regardless of when and where they are produced; all certified solar-generated MWh are treated similarly;⁴ and RPS programs constrain the siting of wind and solar generators.

Ignoring the facts that the marginal value of electricity varies substantially over time and space and that sites for wind and solar generators vary on several dimensions of quality simplifies these policies but also reduces their efficiency. It is not clear a priori that this is a bad tradeoff, of course. If the PTC and state RPS program also covered nuclear plants, for instance, it would seem unlikely that site choice would affect performance much or that these baseload plants would be able to be very responsive to changes in the value of electricity over time except, perhaps, via scheduling maintenance outages.

Wind and solar plants are different from nuclear plants, of course. Their performance on several dimensions is affected by the weather and thus by location, as well as by operator decisions. Some places are obviously sunnier than others and some are windier, but other aspects of weather also matter. Since the spot market price of electricity varies over time, the average social value of wind and solar generators are importantly affected by *when* the wind blows and *when* the sun shines at their sites, as well as by when facility maintenance is scheduled. Moreover, the outputs of wind and solar facilities are *intermittent* -- variable over time and imperfectly predictable -- so the observed output of a solar or wind facility over time is a realization of a stochastic process. Both the variability of that process and its correlations with electricity demand and the output of other intermittent facilities will affect the facility's contribution to system capacity and the costs of integrating it and facilities like it.

The importance of the variation over space and time along multiple dimensions of wind and solar generators' performance has implications for academic research, market design, and public policy. Can wind

¹ The definitive source on subsidy programs in the U.S. is the DSIRE web site maintained by North Carolina State University for the U.S. Department of Energy: <http://www.dsireusa.org>. REN21 (2014) provides useful information on subsidy programs and renewable generation in other nations.

² This provision expired at the end of 2013. As this is written, its reinstatement is being debated in Congress.

³ For a detailed discussion of state RPS programs, see Schmalensee (2012).

⁴ As we discuss below, the fixed feed-in tariff programs that are popular outside the U.S. provide a fixed payment (sometimes technology-dependent) for each MWh generated regardless of when and where it is generated and thus share this characteristic.

generators' performance in a single region be treated as representative of the nation as a whole? Should all solar plants receive identical treatment in a regional capacity market? Are the costs of constraining the siting of renewable generators acceptably small? Developing the factual basis necessary to address these and other related questions requires quantitative analysis of the performance of wind and solar generators within and between multiple U.S. regions and over time. This paper presents such an analysis.

When essentially all US generating plants were owned by regulated public utilities, data on their performance were publicly available and were employed in numerous empirical studies. Many wind and solar facilities are now unregulated and operate in competitive markets, however, and data on the operation of those facilities are commonly treated as proprietary and confidential. This study employs confidential data supplied by the seven U.S. Independent System Operators (ISOs) on hourly output and nodal spot prices for a sample of 25 wind and 9 solar (photovoltaic) generating plants across the U.S. for 2011 and up to 12 adjacent months. In order to maintain confidentiality most characteristics of those facilities were not provided. This essay presents an analysis of the performance of those 34 plants. As far as I know, this study is the first to employ data of this sort from all seven U.S. ISOs and thus to examine performance differences both within and between regions as well as, to a limited extent, over time.

Section 2 provides a brief description of the data used in this study. More detail is provided in Sections A.1 and A.2 of the Appendix. Taking advantage of our uniquely rich price data, Section A.3 of the Appendix provides new information that may be of independent interest on the distributions of spot prices over time at individual network nodes and across space within ISOs. Nodal price distributions are generally right-skewed and have considerably fatter tails than normal distributions, and prices often differ substantially within some ISOs.

Section 3 considers the distributions of capacity factors among the generators in our sample and of their *value factors*: the ratios of the average spot-price value of these generators' outputs to the unweighted average spot prices they faced. Variation in cross-section and over time in capacity factors and in wind generators value factors are substantial. Some correlates of the cross-section differences in value factors are discussed, and evidence consistent with their expected decline over time is presented. Finally, this section explores generators' reactions to the negative spot prices that all units outside ISONE faced in 2011 and shows that negative prices generally did *not* induce generators to reduce output.⁵

Section 4 presents evidence on three quantities related to the intermittency of wind and solar generation: capacity factors during hours of high spot prices, hour-to-hour and day-to-day variability in output at the facility and ISO levels, and the incidence of low or zero generation. Capacity factors during high-price periods are found to vary substantially over space and time. Measures of variability are developed to deal with the predictable diurnal changes in the output of solar facilities, and a new measure of the gains from geographic averaging across facilities is presented and employed. Variability differs substantially in cross-section, as do gains from geographic averaging. Spells of zero ISO-level wind output appear much more common in the Northeast than elsewhere. Both inter- and intra-regional differences are important, and some quantities vary substantially over time.

⁵ Generators' bids were constrained to be non-negative in ISONE during the period studied here.

This analysis reveals that there is a great deal of variation around overall and regional means of many performance measures, and some vary substantially over time. Some of that variability reflects the poor design of current U.S. subsidy policies, which also leads to social inefficient operating decisions. Section 5 provides a brief discussion of some implications of our substantive findings for research, market design, and public policy.

2. Data Employed

As Joskow (2011), Borenstein (2012), and others have stressed, in the absence of identifiable externalities, the best measure of the marginal social value of the output of any particular generator is given by the location-specific spot prices that generator faces. Unfortunately, in the U.S. location-specific spot wholesale prices exist only in the regions served by the seven Independent System Operators (ISOs), which manage organized wholesale electricity markets and regional transmission systems.⁶ These systems meet around 2/3 of U.S. electricity demand and serve around 2/3 of U.S. electricity customers.

The sample of wind and solar generation plants analyzed here is accordingly drawn from those systems, and data on their hourly outputs and the corresponding spot prices were kindly provided by all of the seven U.S. ISOs:⁷

The Electric Reliability Council of Texas (**ERCOT**), which serves most of Texas.

ISO-New England (**ISONE**), which serves the six New England States.

The Midcontinent ISO (**MISO**), which serves North Dakota, Minnesota, and Iowa, as well as most of South Dakota, Illinois, and Indiana, and small parts of several adjacent states.

The New York ISO (**NYISO**), which serves New York State

The PJM Interconnection (**PJM**), which serves Pennsylvania, New Jersey, Maryland, Delaware, Virginia, West Virginia, and the District of Columbia, as well as most of Ohio and parts of Illinois, Indiana, and other adjacent states.

The Southwest Power Pool (**SPP**), which serves Nebraska, Kansas, and Oklahoma, as well as parts of Texas, New Mexico, and other adjacent states.

The California ISO (**CAISO**), which serves most of California.

For all but SPP, the spot price data are Locational Marginal Prices (LMPs) or nodal prices for the network nodes at which each generator in the sample is located. These LMPs are defined as the short-run marginal cost of meeting an additional MWh of demand at the node in the transmission system at which the generator is located, taking into account transmission losses, transmission line capacity constraints, and the (as-bid) costs of incremental generation.⁸ The SPP prices are not LMPs since they do not take into account transmission losses, but they are the

⁶ Baker et al (2013) use system lambda, short-run marginal generating cost, to measure the value of (estimated) solar generation in areas without organized markets. As they note, system lambda differs from the location-specific prices used here because it does not take into account transmission congestion or losses and, unlike location-specific spot price, it can never be negative.

⁷ ISONE, MISO, PJM, and SPP have also been certified by the Federal Energy Regulatory Commission (FERC) as Regional Transmission Organizations (RTOs) and thus have somewhat greater responsibilities for system reliability than the other ISOs. Figure 1.2 of MITEI (2011) shows the ISOs' territories as of late 2011. Since the Southeast and much of the West are not served by ISOs, these regions are not covered by this study.

⁸ See Hogan (1992) or Hsu (1997).

spot prices each generator in fact faced. (As this was written SPP was in the process of moving to a full LMP pricing system.)

I asked personnel of each of the ISOs for hourly price and output data for at least three representative wind facilities and three representative solar facilities, geographically dispersed within each system, covering a two-year (17,520-hour) period that included all of 2011. As described in more detail in the Data Appendix, two years of data were provided for all wind plants except those in ERCOT and CAISO. Both ERCOT and MISO provided data on five wind facilities, so the wind sample includes 25 facilities covering all seven U.S. ISOs.

The solar sample is much less comprehensive. Either because they had no utility-scale solar units or because they had so few that providing data on any of them might compromise confidentiality, ERCOT, MISO, NYISO, and SPP were unable to provide any solar data. ISONE, PJM, and CAISO provided data on three solar photovoltaic plants each, but in each case for fewer than the hoped-for 17,520 hours, as discussed in detail in the Data Appendix.

The reason for requesting data covering a two year period was to enable examination of year-to-year changes in various quantities. For facilities with shorter data series, Early/Late changes reported in various Tables below were computed between 8760-hour early and late periods that overlapped, as described in the Data Appendix. This procedure avoids contamination by seasonality but undoubtedly understates year-to-year variability.

In order to preserve confidentiality, five of the ISOs provided no information on the locations of the facilities their data covered. The two exceptions were ISONE, for which reported differences in location had no obvious relation to any performance measure and thus played no role in this study, and ERCOT, for which, as discussed below, locational differences were very important indeed.

Finally, only CAISO provided nameplate capacities. Instead of these missing figures, I used the largest observed hourly output in the data as the measure of capacity for all facilities. Thus the *capacity factor* (*CF*) for each plant was computed as the ratio of average hourly generation to the maximum hourly generation observed. For the three CAISO wind facilities, maximum observed generation was 110.3%, 103.4%, and 102.0% of nameplate capacity. For the three CAISO solar facilities these numbers are 99.8%, 93.0%, and 108.1%. These data suggest that our *CF* statistics will be close to the true capacity factors and perhaps slightly below them on average. Moreover, since in four of the six CAISO units output above nameplate capacity was observed, it is not obvious that *CF* is inferior to the conventional nameplate-based measure of capacity utilization. Finally, it seems unlikely that any of the sites were so poorly chosen by profit-seeking developers that generation at that site remained substantially below nameplate capacity for all of the more than 13,000 consecutive hours for which we have data.

While this data set does contain an abundance of plant-specific hourly information in multiple regions, it has obvious weaknesses that should be kept in mind in interpreting the results below. We do not observe the vintages of these facilities nor any specific differences in the technologies they embody. We do not have data on all the wind or solar plants in any ISO, nor do we know the locations of the plants we do observe. Thus we do not know whether our sample of plants is truly representative in any sense.

3. Output Value

3.1 Measuring Relative Value

Consider a proposed wind or solar generating facility and set aside for the moment its impact on system-wide capacity, any incremental integration costs, and any other associated externalities.⁹ Let C be that facility's pre-subsidy levelized cost of electric energy (LCOE), let P_h be the nodal spot wholesale price it faces in hour h , and let Q_h be its output during that hour. Then neglecting within-year discounting, that facility will make a positive contribution to social welfare if and only if

$$\Pi = \sum_{h=1}^H (P_h - C) Q_h > 0, \quad (1)$$

where the summation is over the H hours of the year considered.¹⁰

Discussions of renewable generation sometimes compare the LCOEs of wind or solar plants with average market prices or the LCOEs of conventional, dispatchable technologies. But, as several authors have pointed out, this would be appropriate only if price were constant or, as in the case of baseload generators, if output were approximately constant.¹¹ Wind and solar generators do not face constant prices, of course, and they do not produce even approximately constant output. Re-writing (1) shows that what matters in general is the weighted average spot price, using facility-specific outputs as weights:

$$\Pi = (\psi \bar{P} - C) \sum_{h=1}^H Q_h, \quad (2)$$

where \bar{P} is the unweighted average spot price faced by this plant, and ψ is the ratio of the output-weighted average of the spot prices it faces to the unweighted average of those prices:

$$\psi \equiv \frac{\sum_{h=1}^H P_h Q_h}{\sum_{h=1}^H Q_h} / \bar{P}. \quad (3)$$

I will follow Hirth (2013) and call ψ the *value factor* for the facility-year being considered. It is an average relative price, with the average (or baseload) price as numeraire.

If a facility's ψ is expected to be generally above one, it may be a socially attractive investment even if its LCOE is generally somewhat above the average wholesale price. On the other hand, if its ψ is expected to be generally below one, its LCOE would need to be substantially *below* the average spot price for the facility to represent an attractive investment.

To get a more intuitive understanding of ψ , let

$$q_h = Q_h / \bar{Q}, \quad \text{and} \quad p_h = P_h / \bar{P}, \quad (4)$$

⁹ Issues related to system capacity and integration costs are the focus of Section 4. Externalities of particular importance are likely to be pollution-reduction benefits associated with the displacement of conventional generation, and, as Siler-Evans et al (2013) argue, regional differences in conventional generation fleets imply considerable regional differences in those external benefits.

¹⁰ If C were the LCOE taking into account available subsidies, (1) would be an approximate test for profitability. Not all power is sold at the spot price, of course, but with price uncertainty, negotiating a contract to sell at a certain, constant price is likely to reduce, not increase, the investment's expected profitability, since the counterparty will generally need to be compensated for bearing price risk.

¹¹ See Borenstein (2008), Fripp and Wiser (2008), and Lamont (2008) for early recognition of this point. Joskow (2011) and Borenstein (2012) provide clear expositions of it, and Hirth (2013) provides a useful overview of the related literature

where \bar{Q} is the (unweighted) average hourly output. These quantities have unit means by construction. A bit of tedious algebra then shows¹²

$$\psi = 1 + \text{cov}(p_h, q_h) = 1 + \sigma(p_h)\rho(p_h, q_h)\sigma(q_h). \quad (5)$$

That is, ψ is greater than or less than one depending on whether prices and quantities are positively or negatively correlated. Given this correlation, the absolute value of the difference between ψ and one is greater the greater are the standard deviations of the price and quantity relatives defined in (4). If either standard deviation is zero, so that price or quantity is constant, ψ equals one, as (3) shows directly.

There is no reason to expect unit-specific values of ψ to be constant over time. In particular, as Hirth (2013) has argued, if the outputs of different wind or solar facilities are positively correlated, one would expect spot prices during their high-output periods to be decreased by increased wind or solar penetration. Thus one would expect unit-specific ψ s generally to decline over time, as wind and solar penetrations are generally increasing. Hirth (2013) observed such a decline in aggregate wind and solar data for European systems over a multi-year period.¹³

We have at most a two-year sample, however, and changes in the weather between our early and late periods could easily lead to changes in the pattern of outputs that mask unfavorable changes in the pattern of prices. In an attempt to control for such weather changes, I proceeded as follows. Let P^t be the vector of prices in period t , and let Q^t be the corresponding vector of outputs. We can then decompose the change in ψ between two periods, 1 and 2, of equal length as follows, using equation (3) to define the function $\psi(P, Q)$:

$$\begin{aligned} \psi(P^2, Q^2) - \psi(P^1, Q^1) &\equiv \\ \frac{1}{2} \{ &[\psi(P^2, Q^2) - \psi(P^1, Q^2)] + [\psi(P^2, Q^1) - \psi(P^1, Q^1)] \} + \\ \frac{1}{2} \{ &[\psi(P^2, Q^2) - \psi(P^2, Q^1)] + [\psi(P^1, Q^2) - \psi(P^1, Q^1)] \} \equiv \Delta\psi_p + \Delta\psi_q. \end{aligned} \quad (6)$$

The first term in the middle of (6), $\Delta\psi_p$, is the average of the changes in ψ caused by changes in prices between the two periods with the output vector held constant, computed using both periods' output vectors. One would expect $\Delta\psi_p$ to be a more sensitive measure of the adverse impact of increased renewable penetration on prices than the raw change in ψ .

3.2 Capacity Factors and Value Factors

Table 1 provides information on the capacity and value factors of the wind and solar facilities in our sample, and Figures 1 and 2 provide scatter plots of these quantities for wind and solar units, respectively. The capacity factor, CF , of any facility is equal to the product of CF^+ , the average capacity factor in the periods when that facility's generation is positive, and Pr^+ , the fraction of hours in which its generation is positive. Taking logs of this identity and calculating sample variances and covariances for our sample of wind plants for 2011, 80% of the variance in

¹² This is closely related to the main result of Lamont (2008), though he works with system marginal costs instead of nodal spot prices.

¹³ See also Mills and Wiser (2013). Correlation of wind and solar generators' outputs within ISOs is discussed in Section 4. Adding substantial renewable generation to existing systems with low load growth can also be expected to depress average wholesale prices; see Gelabert et al (2011) for an analysis of this effect in Spain.

$\log(CF)$ is contributed by the variance in $\log(CF^+)$, and the covariance between the two components adds another 12%. For solar facilities, in contrast, the two variance terms are roughly equal, and the covariance *subtracts* 20%.

****Insert Table 1 and Figure 1 and 2 near here****

The average wind and solar CF values in Table 1 are in line with reported national averages.¹⁴ But their 2011 ranges are impressive, as are their 2011 coefficients of variation.¹⁵ Moreover, even though wind plants in our sample have roughly twice average the capacity factor of solar plants, the two distributions have a sizeable overlap: five of the wind plants had 2011 capacity factors below the highest CF among the solar plants, while all three CAISO solar facilities had capacity factors above the lowest wind plant CF .

ISO-specific average CF s also had a substantial range: from 0.23 (ISONE) to 0.42 (SPP). In a standard analysis-of-variance decomposition, between-ISO variation accounted for 69% of the total variance of wind plant CF s.¹⁶ For solar plants, the average CF in CAISO (0.24) was more than twice the averages in PJM (0.09) and ISONE (0.11), and between-ISO variation accounted for 93% of the total 2011 variance in solar plant CF s.

The early/late coefficients of variation shown in Table 1 were computed as the ratios of the standard deviations of changes between early and late 8760-hour periods (some of which overlap, as the Data Appendix indicates) and the corresponding 2011 means. The early/late coefficients of variation and the 2011 coefficients of variation are thus not strictly comparable. Nonetheless, Table 1 would seem to imply that in this sample the cross-section differences in capacity factors of both wind and solar plants are more important than the variation between the two (adjacent or overlapping) years compared, much more important in most cases.¹⁷

The average value factors shown in Table 1 indicate that on average, relative to unweighted average spot prices, an MWh from a solar facility was worth about 32% more than an MWh from a wind facility, reflecting the fact that spot prices are generally higher in the daytime than at night. But there is considerable variation in this ratio at the individual facility level. The mean of the 225 (= 9x25) ratios of a solar plant ψ to a wind plant ψ in our sample was 1.37, with a standard deviation of 0.37 and an inter-quartile range of 1.21 to 1.41.

As Figure 1 indicates, only two of the wind facilities in our sample had value factors much above one in 2011.¹⁸ These were the two coastal facilities in ERCOT. At the other extreme, the lowest value factor was computed for one of the western ERCOT plants. Overall, there is a weak negative relation between ψ and CF ($\rho = -0.25$), which is visible in Figure 1. The relation is considerably stronger if the three ERCOT outliers just mentioned are removed ($\rho = -0.53$). Table 1 shows that for wind generators the early/late variation in ψ , discussed in more detail below, is roughly comparable to the cross-section variation in 2011 and that ψ is somewhat less variable than

¹⁴ See, e.g., Joskow (2011)

¹⁵ As the efficiency of wind turbines has improved over time, some of the variation in CF among wind facilities likely reflects unobservable vintage effects rather than locational differences. (Vintage effects seem a priori less likely to be substantial for solar generators.) On the other hand, as we note below, inter-ISO variation exceeds intra-ISO variation, for which vintage effects seem likely to be more substantial. This argues that the observed variation in CF is primarily due to locational differences. The early/late variation in CF shown in Table 1 is clearly uncontaminated by vintage effects.

¹⁶ Average CF s and ψ s for each ISO are given in Schmalensee (2013, Table 2). While the null hypothesis of no differences among ISO means is not especially plausible, it may be useful to keep in mind that the 5% critical levels of the standard F-test of that hypothesis correspond to between-ISO shares of 47% for the wind sample and 95% for the very small solar sample. The 1% critical value for the wind sample corresponds to a between-ISO share of 57%.

¹⁷ These statistics are likely to understate the variation in wind facility capacity factors over longer periods, both because early and late periods overlap for 8 units and because wind energy apparently changes significantly over multi-year periods: see Gunturu and Schlosser (2011) for a discussion and references.

¹⁸ Another (in ISONE) had $\psi = 1.003$, and two others (one in ISONE and one in NYISO) had $\psi > 0.98$.

CF in both time-series and cross-section dimensions. ISO-average ψ s for wind facilities ranged only from 0.84 (CAISO) to 0.99 (ISONE), and between-ISO variation accounted for only 11% of the total 2011 variance in these quantities.

All of the solar facilities have $\psi > 1$, and there is considerably less cross-section variation than for wind generators. There is a weak negative relation between ψ and CF ($\rho = -0.45$) among solar plants that is visible in Figure 2. This reflects the existence of three distinct groups of facilities. The upper-left group in Figure 2 all had high values of ψ and low values of CF in 2011; three were in PJM and one was in ISONE. The two plants with low values of ψ in the center of the Figure were both in ISONE. Finally, the three facilities in the right-most cluster were, as one might have expected, all in CAISO. They had slightly higher values of Pr^+ than the other facilities and much higher values of CF^+ than all but one other facility. Between-ISO variation accounted for 68% of the (small) 2011 total variance in solar plant ψ s.

Discussions of the location of wind and solar generators typically focus on the average energy density of wind and the amount of insolation per unit area.¹⁹ Such discussions suggest the desirability of choosing locations to maximize total output per unit of capacity, measured here by CF . But to maximize *value* per unit of capacity, one would want to choose a location that maximized the value-adjusted capacity factor: $VCF \equiv \psi * CF$. Table 1 gives summary statistics for VCF for wind and solar units.

For wind units, VCF has a slightly greater range and slightly more variability than CF . The generator with the greatest VCF in 2011 was the right-most of the two high- ψ facilities in Figure 1, one of the Coastal ERCOT facilities, not one of the four generators that had higher CF s. On the other hand, broadly consistent with the findings of Lewis (2010) for Michigan, only for CAISO did the facility with the highest CF in the region not also have the highest VCF . Decomposing the variance of $\log(VCF)$ for wind plants as above reveals that variations in $\log(CF)$ and $\log(\psi)$ are of roughly equal importance, with the substantial negative covariance reducing the variance of $\log(VCF)$ by 37%.

As Figure 2 might suggest, in our small sample of solar generators the variance of $\log(VCF)$ mainly reflects the variance of $\log(CF)$, with the covariance term reducing the variance of $\log(VCF)$ by 9%. The three CAISO plants with the largest capacity factors also have the largest values of value-adjusted capacity factors, and within each of the three regions the plant with the highest CF also has the highest VCF . For wind facilities the cross-section variation in VCF is somewhat more important than the early/late variation; for solar facilities it is much more important, reflecting the large differences between CAISO on the one hand and PJM and ISONE on the other.

The data presented here raise an obvious question: why were some wind and solar facilities built on sites that produced much lower capacity factors than were attained elsewhere? Three reasons seem plausible. First, since most historical wind data were collected at ground level or slightly above, well below the hub height of wind turbines,²⁰ it is thus possible that poor data led to poor siting decisions for some early wind plants. Second,

¹⁹ See, for instance, the location-specific information provided by the National Renewable Energy Laboratory: http://www.nrel.gov/renewable_resources/

²⁰ Gunturu and Schlosser (2011) discuss the generally available wind energy data. Note that this (potential) problem does not arise for solar facilities, and I am told that wind farm financing is now contingent on the supply of multiple years of hub-height data.

wholesale electricity prices tend to be higher in ISONE and NYISO than in the Midwest or California,²¹ so that wind and solar facilities may be less uneconomic in the former regions than their capacity factors would suggest. Finally, and probably most importantly, as noted in the Section 1, state RPS programs constrain the siting of wind and solar generators. The rationale is generally economic development, but the result is clearly sub-optimal generator siting from the point of view of the nation as a whole.

****Insert Table 2 near here****

Table 2 decomposes the early/late changes in value factors per equation (6). While only 13 of the 25 wind-plant ψ s fell, the average change was negative. What is striking is that none of the $\Delta\psi_q$ terms were negative; the decline in the average value of ψ is entirely due to the 14 negative values of $\Delta\psi_p$. For our nine solar facilities, $\Delta\psi$ was positive on average and for most facilities. Table 2 reveals, however, that this result was due to the 7 positive values of $\Delta\psi_q$; the late periods were apparently a bit sunnier than the early ones. Seven of nine values of $\Delta\psi_p$ were negative, and the average of this quantity for these solar units was roughly equal to the average for wind generators. Thus a general movement of relative spot prices against wind and solar generators is visible.

Analysis at the ISO level suggests, however, that this price movement may not have been driven primarily by increased penetration of wind and solar generation. I compared within-ISO averages of $\Delta\psi_p$ to changes in the corresponding wind or solar shares of total generation between 2010 and 2012 and, to allow for nonlinear effects, to changes in the squares of those shares.²² All generation shares increased, but changes in shares or their squares were uncorrelated with $\Delta\psi_p$ averages for both solar or wind. All the solar shares were so small (the largest was 0.7 percent for CAISO in 2012) that it is not surprising that the small observed changes had no apparent effect on prices. The wind shares were considerably larger, however, and some increases were substantial. The share of wind in generation in SPP, for instance rose from 4.9 percent to 9.4 percent and, while the average $\Delta\psi_p$ was negative for SPP, it was smaller in absolute value than the average declines in PJM, NYISO, and MISO, in all of which both the initial wind share and its change were notably smaller.

It is not clear that much weight can be placed on these negative results, however. The $\Delta\psi_p$ s were calculated using at most 24 months of data centered on 2011, not the 2010-2012 period for which we have state-level generation data. Moreover, the boundaries of five of the seven ISOs did not correspond exactly to state boundaries. Finally, increased concentration of wind or solar facilities in particular regions within an ISO (e.g., solar facilities in New Jersey) could lead to systematic declines in observed nodal spot prices in those regions, even if wind or solar penetration did not change much at the ISO level (as it did not within PJM, though the average solar $\Delta\psi_p$ was much the largest in absolute value).

3.3 Some Possible Correlates of Value Factors

Table 3 sheds some light on why output from some facilities is more valuable on average than output from others. Solar generators do not produce at night, for instance, when prices are generally low. In contrast, the first line of

²¹ The FERC provides information on average prices at selected points within the ISOs: <http://www.ferc.gov/market-oversight/mkt-electric/overview.asp>.

²² I used state-level generation data from the Energy Information Administration's *State Electricity Profiles*, e.g., <http://www.eia.gov/electricity/state/california/index.cfm>. For PJM, SPP, and MISO, I used only those states entirely in the ISO. I used California and Texas data for CAISO and ERCOT, respectively, even though parts of those states are not formally within those ISOs. Thus, except for ISONE and NYISO I have necessarily measured ISO-level generation changes with error.

Table 3 reveals that wind generators on average produced 15% more per hour during the low-load night hours (defined here and below as 10:00 pm through 6:00 am) than at other times. While the ratio of generation at night to generation at other times appears stable over time for individual facilities, at least within our short sample period, there is considerable variation among facilities, and this ratio is highly negatively correlated with 2011 value factors. All but three wind plants generated more on average at night than at other times; two of these were the coastal ERCOT facilities with the highest 2011 value factors in the sample. At the other extreme, the plant with the highest ratio of generation at night to generation at other times was the western ERCOT facility with the lowest 2011 value of ψ .

****Insert Table 3 near here****

In the U.S., system peak loads generally occur in the summer. Table 3 reveals that, as one would expect, all solar units had higher average generation in the summer (defined here as June through August) than in other months, while wind facilities on average produced about 20% less. Four wind facilities were more productive in the summer than in other months, however: one in ISONE and all three in CAISO.²³ Differences in the ratio of generation in the summer to generation at other times were essentially uncorrelated with differences in ψ among wind facilities or among solar facilities, however.

Wind or solar generators add more to system reliability the more they produce during periods when the system is short of capacity and incremental generation is particularly valuable. To explore this aspect of performance and its relation to value factors, we study generation at times of high spot prices. While peak system-level demand will cause high spot prices, such prices can also occur with moderate loads and planned or unplanned outages of generation or transmission capacity. I found a spot price level for each facility such that it was exceeded for about 100 hours in 2011 and studied generation during those peak-price hours.²⁴

Table 3 gives the ratio of average generation during those hours to average generation at other times. While the ranges of this statistic for both wind and solar plants are impressive, so is the contrast between wind and solar plants. Wind facilities generated on average 28% less during peak-price hours than other hours during 2011, while solar facilities generated 58% more on average. While all solar plants generated more during peak than off-peak hours in 2011, only four out of 25 wind plants did: the two coastal facilities in ERCOT, again, and two of the three ISONE generators. Within the solar sample and, especially, the wind sample, the ratio of peak to off-peak generation was positively correlated with plant-specific value factors in 2011. For both technologies, this ratio exhibits substantial cross-section variation and even greater inter-temporal variability.

3.4 Negative Spot Prices

All generators outside ISONE faced negative spot prices for at least 18 hours during 2011.²⁵ These prices often reflected transmission congestion, and Texas has for some time lacked adequate transmission capacity between its wind generators in the west and its load centers in the east. Thus it is not surprising that one of the western ERCOT units had the highest incidence sample of negative prices in this sample in 2011: 1542 hours, or just under 18% of the year. The other two western ERCOT units faced negative prices in 857 and 529 hours during 2011. But other

²³ The CAISO result is consistent with Fripp and Wisler (2008), who find that winds in coastal passes in California are stronger in the summer than in other seasons.

²⁴ The range was 95 to 104 hours.

areas also experienced a high incidence of negative prices. Nine other units in the sample, including all six CAISO wind and solar units, faced negative spot prices during more than 500 hours in 2011.

During the period covered by our data, ISONE alone did not allow generators to bid negative prices. While negative spot prices have nonetheless arisen in ISONE during episodes of severe congestion, such episodes have been rare, and there are no negative prices in any of our ISONE data. During 2011, all six ISONE wind and solar facilities faced spot prices that were exactly zero during 47 hours. (One faced zero spot prices in two other hours.) At the ISONE interior hub, day-ahead prices were positive in all those hours, and the spot prices were unaffected by congestion. It thus appears that the spot market cleared at the lowest price bid, zero, because the system load was unexpectedly low, not because of congestion.²⁶ Because of ISONE's unusual market structure during this period, the analysis in the remainder of this section excludes the facilities in that ISO.

****Insert Table 4 near here****

Table 4 shows that on average wind facilities outside ISONE had positive outputs during 92% of the hours when the spot price they faced was negative. Moreover, wind plants produced 49% more on average during those hours than at other times. It seems likely that wind generation in negative-price periods would have been even higher in the absence of ISO curtailment actions, particularly in ERCOT.²⁷ As a mechanical matter, this explains the negative correlation in 2011 between the incidence of negative prices and units' value factors: output was generally high when the spot price was negative, so the more frequent were negative prices, the lower the average spot price received. Table 4 reveals that the incidence of negative prices varied considerably among the non-ISONE plants in our sample as well as between early and late periods.

Negative prices at wind plant nodes are sometimes described as a nighttime phenomenon, since, as noted above, wind output is higher on average at night when demand is generally lower. It is indeed true that wind plants were more likely to face negative prices at night, but the difference is less pronounced than sometimes suggested. The eight night-time hours between 10:00 pm and 6:00 am accounted for 52% of the negative prices in the 2011 sample. Since there are twice as many daytime as nighttime hours, nighttime hours were $2.2 \{=.52/[(1-.52)/2]\}$ times as likely to have a negative price than daytime hours.

Table 4 shows that the solar plants in our PJM and CAISO samples faced negative spot prices roughly as often on average as wind plants in 2011. This average hides a large inter-ISO difference, however: facilities in PJM averaged 37 hours with negative prices, while facilities in CAISO averaged 729 such hours – more than all but the three western ERCOT plants in our 22-plant, ISONE-excluded sample of wind facilities. Solar facilities produced 62% less on average when prices were negative than at other times, thus accounting in part for their generally higher value factors than wind generators.

²⁵ See Huntowski et al (2012) for a general discussion of negative spot prices in wholesale electricity markets and some historical data.

²⁶ Thirty-three of the zero-price hours occurred between 2:00 am and 7:00 am, when loads are usually low. Another 12 occurred during other times in the August 28 – 30 period, when Tropical Storm Irene was moving north through New England, causing outages and reducing air conditioning demand.

²⁷ All ISOs reported that they curtailed wind output under some conditions, though clearly not whenever spot prices turned negative. The specific conditions and the frequency of curtailment seem to have varied substantially among the seven ISOs. Information on curtailments in a few regions during 2011 is provided in U.S. Department of Energy (2012, pp. 42-43), which suggests that curtailments were particularly important in ERCOT. See also the discussion of zero-output hours in Section 4.3, below.

A major reason for this difference is that, as just noted, negative prices are more likely to occur at night, when solar facilities can't generate. Negative prices were slightly more likely to occur at night for solar plants than for wind plants. In 2011, 57% of the hours when solar plants faced negative prices occurred at night, so that solar plants were $2.7 \{=.57/[(1-.57)/2]\}$ times more likely to face a negative price in a nighttime hour than in a daytime hour. The last two lines in Table 4 present the incidence of negative prices in daytime hours for the solar facilities in our PJM and CAISO sample and the fraction of daytime hours with negative prices in which those facilities had positive generation. It is clear that PJM and CAISO solar facilities generally continued to generate during the day even when facing negative prices.²⁸

Electricity generation in the face of negative prices is an unintended consequence of governments' subsidy policies discussed in Section 1. As long as the production tax credit (PTC) plus the spot price is positive it is profitable to run a wind generator and sell the output on the spot market. If the generator's output is instead sold under a long-term contract with a positive per-MWh price, as is relatively common, the incentive to operate the generator in the face of negative spot prices is even stronger. Thus as wind generation increases in importance, one can expect the incidence of significantly negative spot prices to increase.²⁹

To explain the failure of solar plants to shut down when the sun is shining but the spot price is negative one must look to state RPS programs, which treat all MWhs generated as equivalent, regardless of when they are generated. The outputs of wind and solar generators are accordingly often sold under long-term contracts with positive per-MWh prices that do not depend on the current spot price. Even if power is sold on the spot market, the value of the associated RECs does not depend on when they were produced. Thus state RPS programs provide solar facilities an incentive to generate even when the value of incremental electricity (i.e., the spot price) is negative, and they provide wind facilities an additional incentive to do so.

4. System Capacity and Integration

4.1 System Capacity

In theory, estimating the contribution a particular actual generating facility's contribution to a power system's load carrying capacity requires a good deal of system- and facility-specific information. In practice, however, ISO capacity markets and system planners often rely on capacity factors during peak periods, sometimes averaged across an entire system.³⁰

****Insert Table 5 near here****

Table 5 provides information on capacity factors during the peak-price periods defined in Section 3.3. Comparing the 2011 *CF* means in Table 5 with those in Table 1 reveals that, as Table 3 showed, wind units

²⁸ I also examined the incidence of negative prices during prime solar generating hours, taken to be from 9:00 am to 3:00 pm. One PJM facility experienced negative prices in 6 prime hours during 2011 and generated during all of them; another PJM plant generated in the single prime hour in which it faced a negative price; and the third PJM solar facility did not face a negative price during any prime hour in 2011. In contrast, the three CAISO facilities saw negative prices in over 3% of all prime hours in 2011, and all generated in at least 97% of those hours.

²⁹ Huntowski et al (2012) find that negative prices have become more common since 2006 at several locations, though their data do not show a 2009-11 trend despite large increases in wind capacity over that period. In early-to-late comparisons in our sample, the incidence of negative prices declined for 14 of 25 wind units, with major declines in western ERCOT and CAISO and substantial increases in MISO. Among solar units there were major declines in CAISO and small increases elsewhere. Since negative prices often reflect transmission congestion, the most plausible explanation for substantial declines in their incidence is expansion of transmission capacity.

³⁰ See Abdullah et al (2014) on computing facility-specific estimates of effective load carrying capacity and NERC (2011) and Rogers and Porter (2012) on the simpler methods often used in practice.

generally produce less in peak-price periods than at other times, while solar units generally produce more. The range of capacity factors for wind in Table 5 is wider than in Table 1, with ERCOT facilities generally high and NYSIO facilities generally low. The range of capacity factors for solar in Table 5 is notably narrower than in Table 1, with CAISO, as expected, well above the others.

The final two columns in Table 5 show considerable cross-section variability, though less than in Table 1, and a good deal more early-late variability than in Table 1. This suggests the need to base estimates of contributions to system capacity on data covering at least several years.

The second line of Table 5 shows that 80% of the variance among wind 2011 peak-price capacity factors was within, not between, ISOs. Thus the use of ISO-wide wind plant averages for planning or capacity market purposes is likely to entail significant mis-characterization of individual generators. In contrast, about half the early-to-late period variation for wind units is between ISOs, pointing to the importance of ISO-wide changes. The variance decomposition results for our small solar sample reflect the very high 2011 peak-period capacity factors for CAISO generators and the fact that all early-to-late changes were negative for those generators and positive for all others. Both wind and solar results presumably reflect both year-to-year ISO-wide changes in the timing of peak-price periods and in wind and insolation patterns.

4.2 Changes in Output

The more variable are the outputs of wind and solar generation, all else equal, the more ramping capability hydroelectric and fossil-fueled generators must be able to provide.³¹ The PTC and state RPS programs provide no incentive for generators to select sites with low variability of wind or insolation, and some RPS siting constraints may work in the opposite direction.

A natural measure of short-run output variability for wind generators is the standard deviation of hour-to-hour changes, scaled by each facility's average hourly output to enable comparisons among facilities of different sizes:³²

$$Vh \equiv \sigma(Q_h - Q_{h-1}) / \bar{Q}_{2011}, \quad (7a)$$

Where $\sigma(\cdot)$ denotes the standard deviation of the quantity in parentheses, Q_h is output in some hour h , and \bar{Q}_{2011} is average hourly generation in 2011.³³ It is also of interest to consider day-to-day variability by focusing on differences between output in the same hours in adjacent days:

$$Vd \equiv \sigma(Q_h - Q_{h-24}) / \bar{Q}_{2011}. \quad (7b)$$

These measures make little sense for solar plants, however, since solar generation is so strongly affected by predictable diurnal variations in insolation. Knowing that output is stable at zero during the night is not very

³¹ There is a huge engineering/economic literature on the costs of integrating intermittent resources, much of it based on simulations of particular regional systems. See Baker et al (2013), Hirth (2013), Holtinen et al (2011), and Mills and Wiser (2013).

³² I make no claim that the measures employed here are in any sense optimal, but they do seem easy to interpret. For alternative measures see, e.g., Ela and O'Malley (2012).

³³ The obvious alternative would be to scale by maximum observed output, our proxy for capacity. But maximum observed output is by definition an outlier, and average output seems a better measure of facility importance. All of the measures developed here are negatively correlated with capacity factors; when scaled by maximum observed output the correlations are positive. To see how this can arise mechanically, suppose a facility's output is M with probability A/M and zero otherwise, so that mean output is A . Then it is easy to show that the standard deviation of the difference between independent draws from this distribution is $\sqrt{2A(M+A)}$. Dividing by A yields an expression that is decreasing in the capacity factor A/M , while dividing by M yields an expression that is increasing in A/M .

informative, and system operators can easily predict the substantial output variation during the day that follows from diurnal changes in insolation – though at high enough penetration even perfectly predictable output changes may be large enough to present significant operational challenges.

An hour-to-hour measure that does not suffer from the diurnal change problem, at least during daylight hours, is the standard deviation of differences between output in each hour and the average of the outputs in the preceding and following hours, normalized by average hourly output during 2011 as above:³⁴

$$Vht \equiv \sigma \left(Q_h - \frac{Q_{h-1} + Q_{h+1}}{2} \right) / \bar{Q}_{2011}. \quad (7c)$$

A comparison of Vh and Vht for wind facilities provides some information on the extent to which hour-to-hour changes reflect short-term trends. The day-to-day analog of Vht is

$$Vdt \equiv \sigma \left(Q_h - \frac{Q_{h-24} + Q_{h+24}}{2} \right) / \bar{Q}_{2011}. \quad (7d)$$

Vdt should be free any problems posed by predictable seasonal changes in the case of solar plants, and a comparison of Vdt with Vd sheds some light on the importance of multi-day trends for wind plants.

In order to eliminate the influence of the predictable night-time stability of solar output at zero, the samples used to compute Vht and Vdt for solar facilities were limited to hours in which either the facility's generation was positive or its generation in both comparison hours was positive. This effectively treats other instances of zero generation as predictable. Statistics for wind plants were computed using data for all hours.

****Insert Table 6 near here****

This difference in sample selection may have something to do with the somewhat surprising finding, presented in Table 6, that solar facilities showed higher average values of Vht and Vdt than wind facilities in 2011. But there are only nine facilities in the solar sample, and, as Table 6 also shows, the corresponding coefficients of variation were much larger for solar plants than for wind plants, which themselves showed substantial variation, particularly in cross-section. Once again, averages do not have much information regarding the performance of any particular generator in any particular year.

Across all wind plants, all pairwise correlations between these measures exceeded 0.72, and the correlations between Vh and Vht and between Vd and Vdt exceeded 0.99. Not only are some locations windier than others on average, it would seem that some have more *variable* wind than others, using almost any measure of variability. The correlation between Vht and Vdt across the small sample of solar plants was 0.99; sunshine is similarly more variable at both hourly and daily time-scales at some locations than at others. Locations with less variability are clearly more socially desirable, all else equal, but, as noted above, current subsidy regimes do not favor them. Moreover, the early/late coefficients of variation in Table 6 indicate that site-specific variabilities change over time, particularly for solar plants, so there is little reason to expect the relative variability of generation at different sites to remain constant.

³⁴ While hourly data may be adequate to capture most variability in wind power, the output of individual photovoltaic facilities can change from minute to minute as clouds pass over.

Excluding two outliers on each side, all values of Vht for wind plants were between 31% and 41% below the corresponding Vh values. Thus a substantial fraction of facility-specific hour-to-hour variability was generally associated with short-term trends or “ramping” episodes. But the variability around those trends was still substantial. For 21 of the 25 plants in this sample, Vht exceeded 0.15 in 2011. Trend episodes extending over multiple days seem to have been considerably less important than shorter multi-hour trends: again dropping four outliers, all values of Vdt for wind plants were between 14% and 19% below the corresponding Vd values.

For all plants, Vd is substantially larger than Vh , and Vdt is substantially larger than Vht . Day-ahead forecasting is more difficult than hour-ahead forecasting, so that the difference between the hourly and daily measures in Table 6 is likely to under-state the difference in the importance of unforecastable changes between these two time-scales. On the other hand, given output changes of equal magnitude, it is easier to adjust to a day-ahead than an hour-ahead change.

Looking at ISO-specific averages reveals that an important reason for the pronounced wind/solar differences in Table 6 is the very high variability of the output from solar plants in ISONE as compared to plants in PJM or CAISO.³⁵ This also drives the high inter-ISO share of the total variance of the solar measures. Focusing on wind plants, ISONE and NYISO stand out as having plants with high variability, while ERCOT and CAISO’s wind plants tend to have low average variability. It seems clear that ISO-level differences are important for all wind-plant variability measures and account for the bulk of the total sample variance for Vd and Vdt for wind plants.

All the within-ISO pairwise correlations among wind plant outputs are positive, suggesting the likelihood of higher aggregate variability than if those outputs were statistically independent.³⁶ The individual correlations differed substantially: eight of the 10 correlations between the outputs of the five ERCOT wind facilities in this sample were below 0.2, for instance, while all three of the SPP correlations were above 0.5, as were all three of the NYISO correlations. These differences among ISOs likely reflect differences in plants’ geographic dispersion within each ISO, of course, as well as any differences in the coherence of regional weather patterns.

The variability of the total output of all the wind or solar plants in any ISO depends on the number of plants involved and their relative scales, as well as the properties of the stochastic process that generates their outputs, which in turn may depend importantly on where facilities are sited. Existing support policies provide no incentives for generators to choose sites with an eye to reducing their impact on ISO-level output variability.

To focus the properties of the relevant stochastic processes in our sample it is necessary to remove the effects of differences in facility scale. Section A.4 of the Appendix derives a scale-adjusted measure, R , of the ISO-specific gains from geographic averaging implied by our data.³⁷ With positive correlations among plants’ outputs, one expects R generally to be between zero and one, with higher values indicating greater realized gain from geographic averaging. With plants of roughly equal size, $R = 1$ corresponds to the case of uncorrelated outputs, and in this case aggregate variability using the measures presented here should fall roughly as the square root of the

³⁵ Recall that ISONE also stood out in Section 3.2 as having low wind and solar capacity factors. ISO-specific averages of the V and R measures in Table 6 are given in Schmalensee (2013).

³⁶ Solar plant outputs within each ISO are of course highly correlated because of the diurnal cycle.

³⁷ Obviously this measure is at best a rough approximation, which it is hoped will shed some light on differences among ISOs. More refined analyses employing more comprehensive information include Katzenstein et al (2010), Fertig et al (2012), and Huang et al (2013).

number of plants. At the other extreme, $R = 0$ corresponds to the case of perfect correlation among generator outputs, and in this case aggregate variability should remain roughly constant when new plants come on line.

The last section of Table 6 presents information on these measures for the variability measures defined by equations (7).³⁸ Given the high correlations among the plant-level variability measures noted above, one might also expect high correlations across ISOs for the R s corresponding to wind plant variability measures, and the correlations for wind plants in fact range from 0.62 (R_{ht} and R_{dt}) and 0.99 (R_d and R_{dt}). ISONE would seem to have gained less from averaging than the other ISOs: it accounts for all of the wind-plant minimum R 's in Table 6, as well as the minimum R for solar facility V_{hts} . At the other extreme, PJM accounts for all the maximum values for wind plants shown in the Table. On the other hand, PJM seems to have gained distinctly less from geographic averaging among solar plants than the other two ISOs. The tiny within-ISO sample sizes and our lack of knowledge about the geographic dispersion of the plants in those samples counsel against taking these differences too seriously, of course.

Finally, except for CAISO solar plants, the gains from geographic averaging are substantially lower in all cases at the day-ahead than at the hour-ahead time-scale. This suggests that intra-ISO changes in wind and solar energy from day to day are more highly correlated than changes from hour to hour. In fact, among wind plants the assumption of statistical independence seems a good approximation for V_{ht} and V_h except for ISONE, since the corresponding R s all exceed 0.80.

4.3 No or Low Generation

A rough indicator of the need to provide backup capacity for wind or solar facilities explored in this section is the frequency with which these intermittent generators unexpectedly produce little or no output at the regional level. These two technologies have quite different behaviors on this dimension: wind generators fairly often (as quantified below) produce zero output, while solar facilities rarely produce zero except at night.³⁹ I deal first with wind and then with solar.

Every wind plant had at least 100 hours with zero output in 2011, and all but two (one in CAISO and one on the Gulf coast in ERCOT) had at least 300.⁴⁰ The average across all units was 948 hours, just under 11% of the year, and the maximum was 1879 hours (in ISONE). The available evidence suggests that most zero-output hours occurred because there was simply not enough wind, not because of ISO curtailment orders. Curtailment orders would seem extremely unlikely when the spot price is positive, since with marginal cost of wind generation effectively zero, a positive price signals that it is in both the generator's and the system's interest for output to be positive. In 2011, for all but 4 facilities (2 western ERCOT plants, and one each in MISO and CAISO), more than

³⁸ The statistics shown were computed using all available data for each ISO, not just data for 2011. I initially computed the ISO-level variability measures for solar generation used in Table 6 by restricting the sample to hours when either actual total ISO generation was positive or total ISO generation was positive in both comparison hours and compared it with plant-level measures computed similarly. This yielded anomalous results for ISONE, however. Two ISONE solar plants had positive outputs for many fewer hours than the third, so that the sample used to measure ISO-level variability had many hours in which output from those two plants were zero and aggregate variability was accordingly low. (These two facilities had positive output for fewer than 20% of the hours in 2011, while the third ISONE plant and all other plants in the sample had positive output for at least 45% of those same hours.) To remove the effect of these differences, I recomputed all plant-specific measures used in the solar comparisons in Table 6 using the same set of hours as for the ISO measures. ISO and plant samples of hours were somewhat different for the other two ISOs as well, but the same recomputation produced only tiny changes in their numbers.

³⁹ This, of course, reflects the fact that wind generators need a certain minimum wind speed to produce electricity (they also produce zero when the wind speed exceeds a safe maximum), while photovoltaic facilities produce something whenever there is sunlight.

⁴⁰ More detail on differences between ISOs related to plant level zero-output hours are given in Schmalensee (2013).

97% of hours with zero generation occurred when the spot price was positive. The average of this percentage across all units in the sample was 96.3%. Thus the vast majority of zero-output hours were very unlikely to have been the result of curtailment orders.⁴¹ In addition, most hours with zero facility generation were part of spells of zero generation lasting at least three hours: the corresponding percentage averaged 83% across all plants and was above 70% for all but one plant (in CAISO).

****Insert Table 7 near here****

Table 7 is concerned with episodes of zero output at all sample plants in each ISO.⁴² Ideally, the likelihood that a new wind facility would have zero output at the same time as other wind facilities in the same ISO would be a consideration in site selection, but existing subsidy schemes provide no incentives for this to occur.

In 2011 only ERCOT had no hours in which none of the wind plants in the sample were generating, and only the last three ISOs listed in Table 7 had fewer than 70 such hours. These three also had much lower ratios of the number of such hours to the average number of hours a plant in the ISO had zero output, as the second column in Table 7 shows.⁴³ Because we have only a small sample of the wind facilities in each ISO, it is not necessarily the case that *no* wind plants in the ISO had positive output when all our sample plants produced zero. But if the plants in this sample are reasonably dispersed within each ISO, as I had requested them to be, periods when none of them are generating are likely to be periods in which total wind generation in the ISO is at least low and perhaps close to zero. As Table 7 shows, the bulk of ISO-wide zero generation hours in 2011 occurred in the daytime (defined as above as 6:00 am to 10:00 pm), when wind generation tends to be somewhat lower on average than at night.

Apart from CAISO, a large fraction of the hours of ISO-wide zero generation from wind facilities in this sample in 2011 were parts of spells lasting at least three hours, though the importance of such spells was substantially lower at the ISO level than at the plant level in all cases. The three northeastern ISOs – NYISO, ISONE, and PJM – stand out not only in terms of the incidence of hours with zero output from all sample plants but also in terms of the length of zero-output spells. All had spells of at least 13 consecutive hours of zero output from all sample plants, while the maximum spell length in the other ISOs was 6 hours.⁴⁴

Geographic averaging will be more effective in reducing or eliminating periods of ISO-wide zero or low generation if instances of zero generation at the plant level are independent events than if they are positively related. If hours with zero generation at the plant level were independent events, the probability of an hour with zero generation at the ISO level would simply be the product of the corresponding plants' probabilities. The last column

⁴¹ Since the production tax credit for wind power in 2011 was \$23/MWh, as long as the spot price a generator faced exceeded -\$23, it was in the generator's interest for its output to be positive. (Many facilities received additional compensation under state RPS and other programs or sold under long-term contracts with positive prices and so would find production profitable at even lower spot prices.) For all but one of the four facilities just mentioned in the text, at least 97% of zero-output hours had a spot price above -\$23. The remaining unit, in MISO, faced such a price during 93% of the hours when it produced no output.

⁴² See Gunturu and Schlosser (2011) for one study of regional calm periods in the U.S. with little or no wind generation possible and a discussion of other such studies.

⁴³ In 2011 these averages ranged from 532 hours in ERCOT to 1,424 hours in ISONE.

⁴⁴ Some of the difference in the incidence of ISO-wide zero-output hours between MISO and ERCOT on the one hand and the northeastern ISOs on the others necessarily reflects the fact that these two ISOs contributed 5 plants each to our sample, while the others contributed only three each. To get a rough sense of how much of the measured ISO-level performance difference reflected this sample size difference, I examined sub-samples of three facilities from MISO and ERCOT. The MISO sub-sample, chosen at random, had 28 zero-output hours in 2011, 57% of zero-output hours in spells of at least three hours, and a maximum zero-output spell in all the data lasting 8 hours. The ERCOT sub-sample, consisting of two plants in the west and one on the coast, had only 2 zero-output hours in all the data, neither of which occurred in 2011. Thus it does not seem likely that the difference between the number of units sampled in MISO and ERCOT on the one hand and NYISO, ISONE, and PJM on the other is the main source of the substantial performance differences shown in Table 7.

in Table 7 reports the ratio of the actual frequency of zero-generation hours at the ISO level to the product of the plant-specific frequencies, using all available data for each ISO because of the rarity of ISO-level zeros. All ratios are substantially above one. The second-last column uses the normal approximation to the binomial distribution to construct a test statistic that would have a standard normal distribution under the null hypothesis of independence. All the resulting Z-statistics are large enough to reject that null hypothesis at any conceivable significance level. For ERCOT, which experienced only one hour of ISO-wide zero wind generation in the entire sample, the exact probability that no such hours would occur under independence was 0.9956, so that the probability of observing one or more hours under that null hypothesis is well under 0.01.⁴⁵ Thus the last two columns of Table 7 show that there is a statistically significant and quantitatively important regional low-wind effect, so that geographic averaging of wind output is generally likely to be noticeably less effective at reducing this dimension of variability than if plant-level zero-generation hours occurred independently.

Solar plants of course have long periods with zero region-wide generation every day, but these periods are almost completely predictable. It might be a troublesome surprise if a solar facility produced nothing in some hour despite having produced positive output in the hour before and the hour after, but such events are extremely rare, accounting for less than 0.3% of the relevant hours for the average plant. Zero generation was more frequent when generation was positive in the same hour in the two adjacent days, but even these events only occurred in 3.4% of the relevant hours for the average plant. Rather than devote attention to events that are so rare even at the plant level, it seemed more useful to focus on hours of low, though generally positive, generation relative to the two adjacent hours or to the same hour in the two adjacent days, defining “low generation” as an output less than half the mean of the two comparison hours.⁴⁶

It turns out that even incidents of low generation relative to adjacent hours are not common at the plant level – the average in 2011 was 133 hours, 1.5% of the hours in the year. Even though the tests for independence described above strongly reject the null hypothesis, ISO-level incidents of low generation relative to adjacent hours are extremely rare: CAISO had no such hours in 2011, ISONE had one, and PJM had only 14.

Hours with low plant-level generation relative to the same hour in adjacent days were much more common in 2011; the average was 583 hours. Independence is again strongly rejected, and the incidence of ISO-level hours with low generation relative to adjacent days ranged from 18 (CAISO) to 473 (PJM). One can think of these events as reflecting days that are much cloudier than adjacent days, and it is no surprise that the probability that such days occur for any one plant has an apparently important regional component – or that such events are much less common in California than in the East.

5. Some Implications

This study has used a unique dataset to investigate many aspects of the performance of individual wind and solar generators across the U.S. The message that emerges most clearly is the general importance of variability within and between ISOs and, for some measures, over time. While data covering a more comprehensive sample of plants

⁴⁵ As mentioned above, data in U.S. Department of Energy (2012) suggests that curtailment was particularly frequent in ERCOT, and a significant fraction of the zero-output hours in two of the five ERCOT units occurred when the spot price was negative, suggesting the possibility of frequent curtailments. Thus the low incidence of ERCOT-wide zero output hours is somewhat surprising.

⁴⁶ Further details on instances of low generation are provided in Schmalensee (2013).

and more years would be desirable, it seems unlikely that such a dataset would fail to deliver this same message. In any case, individual facilities in our sample often depart dramatically from overall averages. Thus, for instance, most wind plants generate less in the summer than other seasons, but not those in our CAISO sample. Most wind value factors are less than those of constant-output baseload plants, but not those of the two coastal plants in ERCOT. For some other technologies, historical averages may be good predictors of the performance of individual new facilities, but that is clearly not true for wind or solar generation. Regions differ, site selection within regions seems to be important, and performance often varies over time.

Serious students of electric power systems already know that no two regional systems are alike because of differences in generation fleets, market rules, and load profiles. This study underlines the fact that no two wind or solar generators are alike either. While average performance is certainly of some interest, actual performance is likely to differ considerably between and within regions, and regional averages can change over time. Similarly, treating all wind units or all solar units identically for capacity market or planning purposes may mask substantial differences and lead to perverse outcomes. For instance, Roberts and Porter (2012) reported that both MISO and ERCOT simply used fixed percentages of nameplate capacity for planning purposes, even though the peak-price period capacity factors in our sample units varied from 0.15 to 0.31 in MISO and from 0.08 to 0.46 in ERCOT. Similarly, despite the large differences in the benefits that ISOs have realized from geographic averaging, Roberts and Porter (2012) report that only CAISO attempted to reward individual facilities for contributing to system diversity.

Two of the patterns observed in Sections 3 and 4 have clear implications for public policy. First, when spot prices are negative and wind or solar plants are able to generate, they generally do so. Encouraging renewable generation when its marginal value to the electric grid is negative obviously raises costs to society, but that is what both the federal PTC and state RPS programs do by rewarding all covered generation equally, regardless of when and where it occurs. In regions with organized wholesale markets, it would provide superior incentives to pay output subsidies only when the spot price is positive or even to pay them on top of (or even make them proportional to) the spot price – i.e., to adopt some version of what has been called a premium-based feed-in tariff (PFIT).⁴⁷ In regions that have not yet adopted organized wholesale markets and thus lack nodal spot prices, it would be better to base subsidies on system lambda than to ignore changes in the value of electricity over time.⁴⁸

Moreover, Schmidt et al (2013) argue that a PFIT would provide an incentive to site incremental wind generators where the covariance with the spot price would be high, all else equal, incentives that current U.S. subsidy regimes and flat feed-in tariffs do not provide. In a simulation study of the Austrian market they find that as a consequence of this incentive, under a PFIT scheme the spatial diversification of wind facilities would be greater, the correlation of wind output with system marginal costs would be higher, and the variance of aggregate wind output would be lower than under a flat feed-in tariff. Thus they find that a PFIT would help address some of the intermittency issues considered in Section 4.

⁴⁷ See Schmidt et al (2013) on this policy approach. For the sake of completeness it should also be noted that the production and investment tax credits are substantially less efficient than direct subsidies because firms without substantial taxable income must engage in tax equity financing in thin markets with high transactions costs (Bipartisan Policy Center, 2011).

⁴⁸ See Baker et al (2013).

Finally, a second empirical pattern with clear implications for public policy is the huge observed regional differences in facility performance – most clearly the capacity factor differences presented in Table 2 and depicted in Figures 1 and 2. One reason why wind or solar plants are sometimes built on sites that will produce poor performance is that site choice is constrained by state RPS programs that limit the locations of facilities that can be used to satisfy utilities’ renewable energy requirements, often because of a desire to create in-state jobs. Since wind and solar generation are very capital-intensive technologies, it is unlikely that these limits can ever have much impact on any state’s employment. But it is certain that for the nation as a whole it would be more efficient to generate electricity from solar power in CAISO than in ISONE, and it would be more efficient to generate electricity from wind in SPP than in NYISO. A national RPS program (or feed-in-tariff) would give a much higher return per dollar spent than a collection of state programs that restrict generator siting.

References

- Abdullah, Md A., Kashem M. Muttaqi, Ashish P. Agalgaonkar, and Danny Sutanto (2014). “A Noniterative Method to Estimate Load Carrying Capacity of Generating Units in a Renewable Energy Rich Power Grid.” *IEEE Transactions on Sustainable Energy* 5(3): 854-865.
- Baker, Erin, Meredith Fowle, Derek Lemoine, and Stanley S. Reynolds (2013). “The Economics of Solar Electricity.” University of California, Energy Institute at Haas, Working Paper 240 (April). Available at http://ei.haas.berkeley.edu/pdf/working_papers/WP240.pdf.
- Bipartisan Policy Center (2011), “Issue Brief: Reassessing Renewable Energy Subsidies.” March. Available at http://bipartisanpolicy.org/sites/default/files/BPC_RE%20Issue%20Brief_3-22.pdf.
- Borenstein, Severin (2008). “The Market Value and Cost of Solar Photovoltaic Electricity Production.” Center for the Study of Energy Markets, University of California Energy Institute, Working Paper 176 (January). Available at <http://www.ucei.berkeley.edu/PDF/csemwp176.pdf>.
- _____ (2012). “The Private and Public Economics of Renewable Electricity Generation.” *Journal of Economic Perspectives* 26(1): 67-92.
- Ela, Erik, and Mark O’Malley (2012). “Studying the Variability and Uncertainty Impacts of Variable Generation at Multiple Timescales.” *IEEE Transactions on Power Systems* 27(3): 1324-33.
- Fertig, Emily, Jay Apt, Paulina Jaramillo, and Warren Katzenstein (2012). “The Effect of Long-Distance Interconnection on Wind Power Variability.” *Environmental Research Letters* 7(3): 034017-22.
- Fripp, Matthias, and Ryan W. Wisser (2008). “Effects of Temporal Wind Patterns on the Value of Wind-Generated Electricity in California and the Northwest.” *IEEE Transactions on Power Systems* 23(2): 477-485.
- Gelabert, Liliana, Xavier Labandeira, and Pedro Linares (2011). “An Ex-Post Analysis of the Effect of Renewables and Cogeneration on Spanish Electricity Prices.” *Energy Economics* 33 (Supplement 1): S59-S65.
- Gunturu, Udaya Bashir, and C. Adam Schlosser (2011). “Characterization of Wind Power Resource in the United States and its Intermittency” MIT Joint Program on the Science and Policy of Global Change, Report No. 209 (December). Available at <http://globalchange.mit.edu/research/publications/reports/all>.
- Hirth, Lion (2013). “The Market Value of Variable Renewables.” *Energy Economics* 38: 218-236.

- Hogan, William W. (1992). "Contract Networks for Electric Power Transmission." *Journal of Regulatory Economics* 4(2): 211-242.
- Holtinen, Hannele, Peter Meibom, Antje Orths, Bernhard Lange, Mark O'Malley, John Olav Tande, Ana Estanqueiro, Emilio Gomez, Lennart Söder, Goran Strbac, J. Charles Smith, and Frans van Hulle (2011). "Impacts of Large Amounts of Wind Power on Design and Operation of Power Systems." *Wind Energy* 14(2): 179-92.
- Hsu, Michael (1997). "An Introduction to the Pricing of Electric Power Transmission." *Utilities Policy* 6(3): 257-270.
- Huang, Junling, Xi Lu, and Michael B. McElroy (2014). "Meteorologically Defined Limits to Reduction in the Variability of Outputs from a Coupled Wind Farm System in the Central US." *Renewable Energy* 62: 331-340.
- Huntowski, Frank, Aaron Patterson, and Michael Schnitzer (2012). "Negative Electricity Prices and the Production Tax Credit." The NorthBridge Group, report (September). Available at http://www.nbgroup.com/publications/Negative_Electricity_Prices_and_the_Production_Tax_Credit.pdf.
- Jónsson, Tryggvi, Pierre Pinson, and Henrik Madsen (2010). "On the Market Impact of Wind Energy Forecasts." *Energy Economics* 32(2): 313-320.
- Joskow, Paul L. (2011). "Comparing the Costs of Intermittent and Dispatchable Electricity Generating Technologies." *American Economic Review* 100(3): 238-241.
- Katzenstein, Warren, Emily Fertig, and Jay Apt (2010). "The Variability of Interconnected Wind Plants." *Energy Policy* 38: 4400-4410.
- Lamont, Alan D. (2008). "Assessing the Long-Term System Value of Intermittent Electric Generation Technologies." *Energy Economics* 30: 1208-1231.
- Lewis, Geoffrey McD (2010). "Estimating the Value of Wind Energy Using Electricity Locational Market Price." *Energy Policy* 38(7): 3221-3231.
- Mills, Andrew, and Ryan Wiser (2013). "Changes in the Economic Value of Variable Generation at High Penetration Levels: A Pilot Case Study of California." Lawrence Berkeley National Laboratory, conference paper (March). Available at <http://emp.lbl.gov/sites/all/files/lbnl-6182e.pdf>.
- MIT Energy Initiative (MITEI) (2011). *The Future of the Electric Grid*. Cambridge: MITEI. Available at <http://mitei.mit.edu/publications/reports-studies/future-electric-grid>.
- North American Electric Reliability Corporation (NERC) (2011). *Methods to Model and Calculate Capacity Contributions of Variable Generation for Resource Adequacy Planning*. Princeton: NERC (March). Available at <http://www.nerc.com/files/ivgtf1-2.pdf>.
- REN21 (2014), *Renewables 2014 Global Status Report*. Paris: REN21 Secretariat. Available at http://www.ren21.net/Portals/0/documents/Resources/GSR/2014/GSR2014_full%20report_low%20res.pdf.
- Rogers, J., and K. Porter (2012). "Summary of Time Period-Based and Other Approximation Methods for Determining the Capacity Value of Wind and Solar in the United States." National Research Laboratory, Subcontract Report NREL/SR-5500-54338. Available at <http://www.nrel.gov/docs/fy12osti/54338.pdf>

- Schmalensee, Richard (2012). "Evaluating Policies to Increase the Generation of Electricity from Renewable Energy." *Review of Environmental Economics and Policy* 6(1): 45-64.
- _____ (2013). "The Performance of U.S. Wind and Solar Generating Plants." MIT Center for Energy and Environmental Policy Research Working Paper 2013-012. Available at <http://web.mit.edu/ceep/www/publications/workingpapers/2013-012.pdf>
- Schmidt, J., G. Lehecka, V. Gass, and E. Schmid (2013). "Where the Wind Blows: Assessing the Effect of Fixed and Premium Based Feed-In Tariffs on the Spatial Diversification of Wind Turbines." *Energy Economics* 40: 269-276.
- Siler-Evans, Kyle, Inêz Lima Azevedo, M. Granger Morgan, and Jay Apt (2013). "Regional Variations in the Health, Environmental, and Climate Benefits of Wind and Solar Generation." *Proceedings of the National Academy of Sciences* 110(29): 11768-11773.
- U.S. Department of Energy (2012). *2011 Wind Technologies Market Report*. Washington, D.C. (August). Available at http://www1.eere.energy.gov/wind/pdfs/2011_wind_technologies_market_report.pdf.

Appendix

A.1. Wind Generators

Two full years of data (17,520 hours) were provided for the following, enabling examination of changes between two adjacent 8760-hour periods:

- ISONE: 3 wind facilities, 5/26/2010 – 5/24/2012.⁴⁹
- MISO: 5 wind facilities, 5/1/2010 – 4/29/2012.
- NYISO: 3 wind facilities, 6/1/2010 – 5/30/2012.
- PJM: 3 wind facilities, 6/2/2010 – 5/31/2012.
- SPP: 3 wind facilities, 5/1/2010 – 4/29/2012.

Data for the five ERCOT facilities began on 12/1/2010, when ERCOT switched to LMP pricing, and ran through 6/4/2012, for 13,247 hours. (This is not a multiple of 24 because this period includes two transitions to daylight savings time and only one transition from that regime.) This means that 4,273 hours are included in both the early and late 8,760-hour samples used to examine year-to-year changes. Three of the five ERCOT units were identified as being in the West; the remaining two were identified as being on the Gulf Coast.

Data for the three CAISO wind units ran from 8/1/2010 through 6/30/2012. One hour was missing in 2011 and two were missing in 2012, for a total of 16,797 hours. The early and late periods used to examine year-to-year changes thus had 723 hours in common.

A.2. Solar Generators

Data for two of the ISONE solar units ran from 4/2/2010 through 3/31/2012, a total of 17,520 hours. Data for the third unit also ran through 3/31/2012 but did not begin until 10/2/2010, for a total of 13,128 hours and an early/late overlap of 4,392 hours. The calculations reported in Tables 8 and 10 in the text employed this shorter sample period for all ISONE solar units.

⁴⁹ Recall that 2012 was a leap year.

Data on one of the PJM solar units ran from 10/1/2010 through 9/29/2012, for a total of 17,520 hours. Data for the other two ran from 8/10/2010 through 5/31/2012, a total of 15,864 hours and an early/late period overlap of 1,656 hours. The calculations reported in tables 8 and 10 in the text employed the 10/1/2010 – 5/31/2012 period common to all these units, covering 14,616 hours.

The CAISO solar data covered the period 8/3/2010 – 6/30/2012, with three hours missing, for a total of 16,749 hours and an early/late period overlap of 771 hours. A close inspection of these data revealed a number of obvious errors. (1) One facility showed identical small, positive outputs for all night hours. Setting all these to zero decreased 2011 generation by 0.545%. (2) Another facility had many nights with small, identical, *negative* generation shown for all hours. Setting all negative reported outputs to zero for this facility increased 2011 generation by 0.467%. (3) The third CAISO solar facility had 29 instances of negative reported generation. These were set to zero. In addition, the raw data showed generation equal to 88% of the maximum observed generation for one stretch of 66 consecutive hours. Rather than lose these days, I replaced the apparently erroneous numbers with actual generation in the corresponding hours of adjacent days. The net effect was to reduce 2011 generation by 2.33%.

A.3. Price Distributions

While the main focus of this study is on the variation of generation plant output and its relation to nodal spot price, our data set provides considerable new information on nodal spot prices in U.S. ISOs. Table A.1 provides summary information on those prices in 2011. Average prices (across all nodes within an ISO) varied by almost a factor of two across the ISOs. More surprising, perhaps, are the differences in variability, as measured by the average of node-specific coefficients of variation. Some of this no doubt reflects substantial regional differences in caps on energy prices: the maximum prices observed in ERCOT (\$3510) and CAISO (\$2297) in 2011 were well above those in the other ISOs. But this is not the whole story: the maximum observed price in MISO (\$591) was among the lower ones, but the corresponding coefficient of variation was substantially above all except those for ERCOT and CAISO.

****Insert Table A.1 near here****

All but one of the nodal price distributions (a wind plant in PJM) were right-skewed according to the standard scaled-third-moment measure, and the mean exceeded the median for all 34 distributions. Because these distributions are heavy-tailed (as discussed below), it seemed useful to present a measure of skewness involving the tails. Letting μ be the sample mean of a nodal price distribution and F^{-1} its empirical inverse distribution function, the *rskew* statistics summarized in Table A.1 are defined as follows:

$$rskew \equiv \left[F^{-1}(.975) - \mu \right] / \left[\mu - F^{-1}(.025) \right]. \quad (A.1)$$

Qualitatively, *rskew* gives the ratio of the length of the right-hand tail to the length of the left-hand tail. The *rskew* statistic exceeded one for all but one of the facilities in this sample (the exception was a wind plant in MISO). The differences in *rskew* among ISOs in Table A.1 may reflect real ISO-level differences in market behavior: *rskew* was

less than 2.0 for all of the 16 nodes in the first four ISOs listed in Table A.1 and above 2.0 for 16 of the 18 nodes in the last three ISOs.⁵⁰

Visual inspection of time-series plots at individual nodes gives the impression of many small deviations from the mean, coupled with a few quite large deviations. The standard scaled-fourth-moment measure of kurtosis confirms this impression; all 34 price distributions are leptokurtic, with heavier-than-gaussian tails.⁵¹ In an attempt to provide a more intuitive measure of tail heaviness, Table A.1 presents ISO-specific averages of *ckurt*, defined as follows:

$$ckurt \equiv \frac{\sigma}{\frac{F^{-1}(.8413) - F^{-1}(.1587)}{2}}, \quad (\text{A.2})$$

where σ is the sample standard deviation, and F^{-1} is as above. For a normal distribution, the difference in the denominator in this equation would be two standard deviations, so one can think of the ratio in the denominator as an estimator of the standard deviation based on the center of the data and the assumption of normality. If the distribution has heavier-than-gaussian tails, however, the sample standard deviation, based on all the data, will be larger than this ratio. Thus differences between *ckurt* and one gives an indication of the extent to which tail heaviness affects that standard deviation. Table A.1 indicates that the price distributions in ERCOT, NYISO, and CAISO depart substantially more from normality than those in the other ISOs, a result consistent with that obtained using the standard scaled-fourth-moment measure.⁵²

The next two columns in Table A.1 provide average serial correlation coefficients for each ISO. The first-order serial correlation coefficients provide a measure of the smoothness of the price series over time. Differences among ISOs are substantial, with ISONE and CAISO being the extreme cases. Correlations between the prices in hours h and $h-24$ reflect the day-to-day variability in prices, taking out diurnal effects. These correlations are all lower than the first-order correlations, with the day-to-day correlation in CAISO lower by more than a factor of two than that in any other ISO. These statistics indicate that time-of-day pricing is not a good approximation to true dynamic pricing.

The final three columns in Table A.1 give an indication of variability of prices over space in the various ISOs, reflecting differences in the ISOs' geographic scope, in the detailed topography of their transmission systems, load centers, and generator locations, as well as a host of other factors. If there were no transmission losses or capacity constraints, prices would generally be equal at all nodes within each ISO. Table A.1 shows clearly that this was not even approximately true. While prices are on average highly correlated within ISONE, they are not at all highly correlated within MISO. Similarly, only in ERCOT was the range of prices less than \$1.00/MWh close to half the time, while in PJM the price range exceeded \$10.00/MWh more than half the time. In general, if retail prices were to reflect marginal costs, they would need to vary substantially over space, as well as over time.

⁵⁰ Different measures of skewness give different results, however. According to the standard scaled-third-moment measure, the distributions in NYISO and ERCOT were, on average, much more skewed than those in the other ISOs. And according to the (mean-median)/ σ measure, the distributions in SPP, ISONE, and PJM stand out as more skewed on average than the others.

⁵¹ Jónsson et al (2010) find that spot price distributions in Denmark are also positively skewed and leptokurtic. In this data set, output change distributions showed smaller departures from normality; see Schmalensee (2013) for details.

⁵² It is worth noting the difference between the maximum and minimum prices faced by generators in 2011 was more than twice as large in ERCOT and CAISO than in any other ISO. NYISO did not stand out on this measure.

A.4. Estimating Gains from Geographic Averaging

In an ISO with N plants, let Q_h^i be the output of plant i in hour h , and let A^i be that plant's average output in 2011.

Let \widetilde{Q}^T be a synthetic total output series computed by rescaling all plants' outputs to have the same 2011 average value as plant number one and adding the results:

$$\widetilde{Q}_h^T = Q_h^1 + \sum_{i=2}^N (A^1 / A^i) Q_h^i, \quad \text{for all } h. \quad (\text{A.3})$$

One can then compute each of the four variability measures defined by equations (7) for each ISO's \widetilde{Q}^T , using NA^1 as the scaling factor, and compare it to the values it would have attained had all the Q^i been statistically independent or if they had been perfectly correlated. Consider any statistic Z^i that, like the quantities for which standard deviations are computed in equations (7), is a linear function of elements of the Q^i . Then the corresponding statistic based on \widetilde{Q}^T , \widetilde{Z}^T , will be the function of the Z^i given by equation (A.3). That is,

$$\widetilde{Z}_h^T = Z_h^1 + \sum_{i=2}^N (A^1 / A^i) Z_h^i, \quad \text{for all } h. \quad (\text{A.4})$$

If the Z^i were uncorrelated, the variance of \widetilde{Z}^T would be the sum of the variances of the terms on the right of (A.4).

It follows that in this first polar case the aggregate variability measure based on \widetilde{Z}^T would be given by

$$\begin{aligned} \left. \frac{\sigma(\widetilde{Z}^T)}{NA^1} \right|_{\rho=0} &= \frac{1}{NA^1} \left[\sigma^2(Z^1) + \sum_{i=2}^N (A^1 / A^i)^2 \sigma^2(Z^i) \right]^{1/2} \\ &= \frac{1}{\sqrt{N}} \left\{ \frac{1}{N} \sum_{i=1}^N \left[\frac{\sigma(Z^i)}{A^i} \right]^2 \right\}^{1/2}. \end{aligned} \quad (\text{A.5a})$$

The quantities in square brackets on the right are just the plant-level variability measures based on the Z^i . The fraction on the right of (A.5a) is the root-mean-square of these plant-level statistics, divided by the square root of the number of plants. In the other polar case of perfect pairwise correlation, it is only slightly more complicated to show that the aggregate variability measure is simply the arithmetic mean of the plant-level measures:

$$\left. \frac{\sigma(\widetilde{Z}^T)}{NA^1} \right|_{\rho=1} = \frac{1}{N} \sum_{i=1}^N \frac{\sigma(Z^i)}{A^i}. \quad (\text{A.5b})$$

A natural measure of the potential value of geographic averaging for the variability measure based on \widetilde{Z}^T is the location of the actual statistic based on that measure on the interval defined by the two polar cases just described:

$$R(\widetilde{Z}^T) = \left[\frac{\sigma(\widetilde{Z}^T)}{NA^1} \right]_{\rho=1} - \frac{\sigma(\widetilde{Z}^T)}{NA^1} \left/ \left[\frac{\sigma(\widetilde{Z}^T)}{NA^1} \right]_{\rho=1} - \frac{\sigma(\widetilde{Z}^T)}{NA^1} \right|_{\rho=0}. \quad (\text{A.6})$$

This is the measure of gains from geographic averaging reported in Table 6 for each of the V measures defined in Equations (7).

Table 1: Summary Measures of Generator Performance

	In 2011		Coefficient of Variation	
	Mean	Range	In 2011	Early/Late*
<u>Wind Generators</u>				
CF	31.2	18.8 - 42.9	0.22	0.13
ψ	0.88	0.39 - 1.14	0.17	0.10
$VCF \equiv \psi * CF$	27.2	14.3 - 44.0	0.26	0.15
<u>Solar Generators</u>				
CF	14.1	6.91 - 25.0	0.52	0.18
ψ	1.16	1.08 - 1.23	0.04	0.07
$VCF \equiv \psi * CF$	16.1	8.16 - 28.2	0.51	0.17

*The standard deviation of early-period to late-period changes, divided by the 2011 mean.

Table 2: Decomposing Changes in ψ *

	Wind Facilities		Solar Facilities	
	Number < 0	Mean Change	Number < 0	Mean Change
$\Delta\psi$	13/25	-0.022	3/9	0.030
$\Delta\psi_p$	14/25	-0.024	7/9	-0.019
$\Delta\psi_a$	0/25	0.002	2/9	0.050

*Decomposition of changes between early and late periods follows (6) in the text.

Table 3: Possible Correlates of Value Factors: Output Ratios

	In 2011		Coefficient of Variation		Correlation with ψ , 2011
	Mean	Range	In 2011	Early/Late [#]	
<u>Wind</u>					
Night/Other	1.15	0.69-1.54	0.16	0.05	-0.73
Summer/Other	0.81	0.38-2.08	0.51	0.16**	-0.15
Peak/Other	0.73	0.21-1.33	0.34	0.30	0.78
<u>Solar</u>					
Summer/Other	1.68	1.13-4.37	0.61	0.17	0.18
Peak/Other	1.58	1.10-2.29	0.25	0.50	0.58

**Excludes ERCOT units, for which we have only three summer months.

#The standard deviation of early-period to late-period changes, divided by the 2011 mean.

Table 4: Negative Prices and Value Factors*

	In 2011		Coefficient of Variation		Correlation with ψ , 2011
	Mean	Range	In 2011	Early/Late [#]	
Output Ratios					
Wind: P < 0/Other	1.49	0.29-2.76	0.34	0.21	-0.51
Solar: P < 0/Other	0.38	0.12-0.59	0.57	0.35	-0.84
Frequencies					
<u>Wind:</u>					
P < 0	0.04	0.002-0.18	1.08	0.73	-0.72
Q > 0 if P < 0	0.92	0.60-1.00	0.11	0.12	-0.06
<u>Solar:</u>					
P < 0	0.04	0.004-0.08	0.99	0.43	-0.95
Q > 0 if P < 0	0.32	0.16-0.43	0.34	0.39	-0.77
P < 0 & Day [†]	0.03	0.001-0.06	1.02	0.56	-0.94
Q > 0 if P < 0 & Day [†]	0.83	0.71-0.93	0.13	0.18	-0.92

*Computed excluding all ISONE plants, which never faced negative prices

[†]Computed excluding observations between 10:00 pm and 6:00 am.

[#]The ratio of the standard deviation of early-period to late-period changes, divided by the 2011 mean.

Table 5: Generator Capacity Factors During Peak-Price Hours

	In 2011		Coefficient of Variation	
	Mean	Range	In 2011	Early/Late*
Wind	0.22	0.08 - 0.46	0.28	0.40
Inter-ISO %**			20	49
Solar	0.20	0.12 - 0.30	0.34	0.54
Inter-ISO %**			77	92

*The standard deviation of early-period to late-period changes, divided by the 2011 mean.

**The between-ISO share of the corresponding total variance.

Table 6: Plant-Level Output Variability Measures and Gains from Geographic Averaging

	Wind Facilities				Solar Facilities	
	<i>Vh</i>	<i>Vd</i>	<i>Vht</i>	<i>Vdt</i>	<i>Vht</i>	<i>Vdt</i>
2011:						
Mean	0.32	1.12	0.21	0.94	0.68	1.61
Minimum	0.16	0.79	0.09	0.60	0.25	0.43
Maximum	0.48	1.52	0.33	1.30	1.64	3.96
σ/μ	0.22	0.17	0.27	0.18	0.68	0.71
Inter-ISO % of σ^2 *	37.3	70.6	42.5	68.9	65.5	74.6
Early/Late σ/μ	0.11	0.12	0.13	0.13	0.25	0.23
Averaging:**						
Mean <i>R</i>	0.85	0.57	0.91	0.60	0.66	0.59
Minimum <i>R</i>	0.61	0.17	0.67	0.18	0.47	0.28
Maximum <i>R</i>	0.96	0.79	0.99	0.81	0.86	0.80

*The percentage of total variance of each measure attributable to differences between ISOs.

**The next three rows report on *Rs* computed for all seven ISOs.

Table 7: Hours with No Wind Generation From All Plants in Each ISO, 2011

ISO (# plants)	Total Hours	/Plant Average	% in Day	% in 3+ Hour Spells	Longest Spell**	Plants Independent?***	
						Z-Statistic	Probability Ratio
NYISO (3)	178	0.16	80	76	19	56.6	11.3
ISONE (3)	173	0.12	76	71	15	27.9	4.0
PJM (3)	106	0.08	89	68	13	19.9	4.2
SPP (3)	74	0.12	92	47	6	51.0	19.3
CAISO (3)	16	0.02	100	21	4	13.4	11.3
MISO (5)	11	0.01	100	45	4	29.8	42.4
ERCOT (5)	0	0.00	–	*	1	*	226.4

*One hour without wind generation was observed in the entire sample: May 11, 2012, 12:00 - 1:00 pm. Under independence the probability of one or more such hours would be .0044.

**Statistics in the columns below incorporate all available data, not just 2011 data.

Table A.1: Summary Statistics of 2011 Hourly Nodal Prices

ISO (# of Nodes)	Average Across Nodes Within Each ISO							Percentage of Hours with Price Range	
	Price, \$/MWh	Coeff. of Variation	rskew*	ckurt*	Correlation of $P^i(h)$ with			< \$1	> \$10
					$P^i(h-1)$	$P^i(h-24)$	$P^i(h)$		
NYISO (3)	38.51	0.80	1.74	3.40	0.56	0.22	0.86	25.5	7.5
SPP (3)	28.72	0.55	1.41	1.36	0.67	0.46	0.88	34.7	12.8
MISO (5)	26.45	1.06	1.23	2.25	0.55	0.26	0.59	0.9	48.2
ERCOT (5)	41.26	3.94	1.25	12.0	0.73	0.40	0.90	49.7	37.3
ISONE (6)	45.57	0.54	2.83	1.96	0.82	0.50	0.99	3.2	4.7
PJM (6)	42.17	0.81	3.17	2.72	0.73	0.42	0.70	0.4	55.4
CAISO (6)	31.10	1.58	2.15	3.51	0.29	0.09	0.84	22.8	5.7

*See text for definitions.




Figure 2. Relative Value (ψ) v. Capacity Factor (CF): Solar, 2011

