NBER WORKING PAPER SERIES

THE PERFORMANCE OF U.S. WIND AND SOLAR GENERATING UNITS

Richard Schmalensee

Working Paper 19509
http://www.nber.org/papers/w19509

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2013

I am 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, and participants in workshops of the MIT Center for Energy and Environmental Policy Research and the Energy Insitutute @ Haas (Berkeley). Errors and opinions are entirely my own, of course. I have received no direct support, though I have of course used MIT facilities and computers. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w19509.ack

The Performance of U.S. Wind and Solar Generating Units
Richard Schmalensee
NBER Working Paper No. 19509
October 2013
JEL No. D24,L94,Q42,Q5

ABSTRACT

Government subsidies have driven rapid growth in U.S. wind and solar generation. Using data on hourly outputs and prices for 25 wind and nine solar generating plants, some results of those subsidies are studied in detail: the value of these plants' outputs, the variability of output at plant and regional levels, and the variation in performance among plants and regions. Output from solar plants was about 32% more valuable on average than output from wind plants. Output variability differs substantially among plants and, on some dimensions, among regions. Policy implications of high generation when prices are negative and dramatic differences in capacity factors are discussed.

Richard Schmalensee
MIT Sloan School of Management
100 Main Street, E62-525
Cambridge, MA 02142
and NBER
rschmal@mit.edu

1. Introduction

Even though total electricity generation in the U.S. increased by less than 10% in the decade between 2001 and 2011, generation using solar power more than tripled, and generation using wind power increased almost 18-fold.[2] That growth was driven in large part by a diverse set of subsidies and regulations at the federal, state, and local levels.[3] In fiscal year 2010, for instance, federal subsidies for solar generation exceeded $1 billion, and subsidies for wind generation came to just under $5 billion.[4] In addition, 29 states and the District of Columbia have Renewable Portfolio Standard (RPS) programs that require distribution utilities to purchase specified fractions of electricity from certified generators that use energy that has been defined as renewable,[5] and 15 of these programs have provisions favoring or requiring solar generation.

Together these two technologies accounted for only 3% of U.S. generation in 2011, with wind accounting for 98% of that total, but unless subsidies are cut sharply, most analysts expect their shares to increase significantly in the next several decades.[6] In order to evaluate subsidies for these technologies, it would be useful to know how well the generating plants involved have performed. This is not a simple question, since wind and solar facilities differ in two important ways from the conventional fossil-fuel and nuclear generators that provided 87% of U.S. electricity in 2011.[7]

First, the average performance of both wind and solar generators is importantly affected by the weather and thus by location. Some places are of course sunnier than others and some are windier. But other aspects of weather also matter. For instance, as a number of authors have

---

[2] U.S. Energy Information Administration (2012b), Table 8.2a.
[3] The definitive source on all these programs is the DSIRE web site maintained by North Carolina State University for the U.S. Department of Energy: http://www.dsireusa.org.
[4] U.S. Energy Information Administration (2011), Table ES2.
[5] Most utilities don't actually have to purchase "green" electricity; instead they need to purchase Renewable Energy Credits (called RECs) produced in proportion to their output by certified generators that use renewable energy. For an overview of RPS programs circa 2011, see Schmalensee (2012); for current, detailed descriptions see http://www.dsireusa.org/.
[6] The share of wind and solar generation is already considerably higher in many European nations. In 2011, wind alone accounted for 28% of generation in Denmark, for instance, and that nation has set a goal of 50% of electricity from wind by 2020.[6] (The Official Website of Denmark, http://denmark.dk/en/green-living/wind-energy/, visited June 30, 2013.) And in Germany, solar power alone accounted for 3.2% of generation in 2011, more than 71 times its share of US generation. (Reuters, "German solar power output up 60 pct in 2011," December 29, 2011, available at http://www.reuters.com/article/2011/12/29/germany-solar-idAFL6E7NT1WK20111229?sp=true, visited June 30, 2013.)
[7] Hydroelectric power accounts for most of the generation not provided by wind, solar, and conventional sources, and it is also differs in these two ways from the other conventional generation technologies.

recognized, both the social value and private profitability of wind and solar generators at particular sites relative to conventional sources are importantly affected by *when* the wind blows and *when* the sun shines at those sites. It follows that discussions of "grid parity" based on comparisons of the levelized cost per kilowatt-hour (kwh) of electric energy (generally abbreviated LCOE) from wind or solar power with average wholesale prices have little economic or commercial meaning.[8]

The second difference between wind and solar generators and those powered by nuclear energy or fossil fuel is that the outputs of wind and solar facilities are *intermittent*: variable over time and imperfectly predictable. Thus the observed output of a solar or wind facility over time is a realization of a stochastic process, and the variability of that process matters. If renewable generating facilities A and B have the same average output and the same correlation with variable load (electricity demand less other intermittent generation) but B's output is more variable over time and/or is less easy to predict, integrating B into an electric power system will impose larger integration costs on the rest of the system than integrating A, all else equal.[9] Another reason why discussions of "grid parity" have little meaning is that they ignore this important fact. In order to accommodate high levels of intermittent generation, a power system will generally need to alter its operations, perhaps by finding ways to make demand more responsive to system conditions, and/or to deploy a more flexible fossil-fueled generation fleet.[10] It may require more backup capacity to handle times when intermittent generators' output is low or even zero if those times occur when system demand is high.

When essentially all US generating units were owned by regulated public utilities, data on the performance of those units were publicly available and were employed in numerous empirical

[8] 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

[9] There is an extensive engineering/economic literature on the costs of integrating intermittent resources, much of it based on simulation of particular regional systems. Recent contributions are provided and discussed by Baker et al (2013), Hirth (2013), Holttinen et al (2011), Kopsakangas-Savolainen and Svento (2012), Mills and Wiser (2013), Pérez-Arriaga (2010), and Ueckerdt et al (2013).

[10] Increasing the frequency of spot price determination has been advanced as a means of facilitating the integration of intermittent generators and is being employed for that purpose in some regions. An interesting research question, raised in personal communication by William Hogan, is whether in the long run very frequent scheduling and spot price determination (setting both every five minutes, say) would provide sufficient incentives for generator flexibility so as to eliminate the external costs of intermittency as a policy or market design problem.

studies.[11] 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. I was nonetheless able to obtain hourly output and nodal spot price data for a sample of 25 wind and 9 solar (photovoltaic) generating plants across the U.S. for 2011 and up to 12 adjacent months, but 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. The average value of the electricity generated by these plants and its correlates are examined, as well as the variability of output over time at the plant and regional levels and the variation among these plants and regions along important performance dimensions.

Almost all previous studies of wind or solar power have used data on wind energy or insolation along with engineering information to infer electrical output.[12] This often assumes ideal performance and always eliminates variation due to differences in the design and construction of individual facilities and in operations and maintenance practices. Moreover, the data considered here relate to sites actually chosen and developed by profit-seeking enterprises, not the larger set of theoretically possible sites often analyzed.

Section 2 provides a brief description of the data used in this study. More detail is provided in the Data Appendix, along with summary statistics on output change distributions and on price distributions over time at individual nodes and across space within ISOs. Output change distributions do not depart for from normality, particularly at the daily time-scale. 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 value of output from the wind and solar generating facilities in our sample relative to the average spot prices they face and analyzes some correlates of cross-section and inter-temporal differences in those relative values. Having data on individual facilities, I am also able to explore their reaction to the negative spot prices that all units outside ISONE faced in 2011. Differences in the timing of generators' outputs imply that on average in our sample a kilowatt-hour from a solar generator is about 32% more valuable, relative to average wholesale spot prices, than a kilowatt-hour from a wind generator, broadly in line with

[11] E.g., Joskow and Schmalensee (1987).

[12] The two exceptions of which I am aware use actual generation at the regional (Green and Vasilakos (2012)) or national (Hirth (2013)) level.

the results of earlier studies (see Hirth (2013)). But in contrast to conventional fossil-fueled generators, the relative value of output from wind and solar facilities varies substantially among facilities and over time. Because location is important and weather changes, the cross-section and intertemporal variation along the dimensions of wind and solar performance documented here is in principle as important as the corresponding averages.

Section 4 presents evidence on two aspects of the intermittency of outputs from wind and solar plants: hour-to-hour and day-to-day variability in output and the incidence of low or zero generation. The predictable diurnal changes in the output of solar facilities pose particular challenges, and measures of variability are proposed to deal with them. In addition to analysis at the facility level, evidence is presented on the extent to which geographic averaging across facilities can serve to smooth wind and solar output and on the relation between the incidence of low or zero generation at the plant and regional levels. As with the analysis of value, a complex pattern is revealed, with substantial cross-section variation.

Section 5 provides some concluding observations. An important policy implication of the results presented here is that the subsidy schemes used in the U.S. to encourage deployment of wind and solar power lead to economically inefficient generator behavior, particularly of wind generators. Section 5 discusses alternative approaches.

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. Absent contracting problems, those prices are also the best measure of the private value of the generator's outputs as well. 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.[13] These systems meet around 2/3 of U.S. electricity demand and serve around 2/3 of U.S. electricity customers.

[13] 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. Unlike location-specific spot price, system lambda can never be negative.

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:[14]

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.

Note that the Southeast and much of the West are not served by ISOs and are thus not covered by this study.

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 Megawatt-hour (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.[15] The SPP prices are not LMPs since they do not take into account transmission losses, but they are the spot prices each generator in fact faced. (As this was written SPP was in the process of moving to a full LMP pricing system.)

[14] 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. The boundaries of these territories change from time to time. A current map has generally been available from the ISO/RTO Council: http://www.isorto.org.

[15] See Hogan (1992) or Hsu (1997).

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 grid-connected 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 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, because the locations of the facilities in our sample were not randomly determined but were chosen by profit-seeking enterprises, it seems unlikely that any of the sites were so poorly chosen 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 is superior in some ways to others that have been employed in related studies, particularly in providing 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 a truly random sample in any sense.

3. Output Value

Section 3.1 discusses the measurement of the value of output from wind and solar plants, and section 3.2 presents data on output value and capacity utilization in our sample. Sections 3.3 and 3.4 consider output patterns and responses to negative prices that might be thought to affect output value.

3.1. Measuring Relative Value

Consider a generating facility that is too small to affect spot market prices. Let C be that facility's levelized cost of electric energy (LCOE), taking into account any available subsidies, 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 and assuming all sales occur at spot prices, that facility will make a positive economic profit in any particular year if and only if

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

where the summation is over the H hours of the year considered.[16] 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.

As noted above, discussions of renewable generation often assert or imply that if the subsidy-free LCOE of a wind or solar plant were less than the corresponding average spot price, that unit would be a profitable investment even without government support. But, as many authors have noted, this is not in general correct. It would be correct if price were constant, or, as in the case of baseload generators, it would be approximately correct if output were approximately constant.

But wind and solar generators do not face constant prices, and they do not produce even approximately constant output. In this case, what matters is not the unweighted average spot price but the weighted average spot price, using facility-specific outputs as weights:

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

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 faced to the unweighted average of those prices:

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

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 an 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.

---

[16] To calculate the social value of the facility, one would need to use the subsidy-free LCOE and take all relevant externalities into account, including those associated with alternatives to the facility being considered.

To get a more intuitive understanding of ψ, let

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

where $\bar{Q}$ is the (unweighted) average hourly output. These quantities have unit means by construction. A bit of tedious algebra then shows[17]

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

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 quality 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 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.

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 (3) to define the function $\psi\left(P,Q\right)$:

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

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

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 factor, as defined in Section 1, for the wind and solar facilities in our sample. *CF* for any facility is equal to the product of two quantities also shown in Table 1: CF^+, the average capacity factor in the periods when that facility's generation is positive, and *Pr(Q>0)*, 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 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%.

The average wind and solar *CF* values in Table 1 are in line with reported national averages.[18] But except for the probability of positive generation for wind facilities, the 2011 ranges of all three quantities just discussed seem impressive for both wind and solar, as do their 2011 coefficients of variation. And 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*.

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 overlatp, 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 imply that in this sample the cross-section differences in the first three quantities shown for both wind and solar plants are more important than the variation

[18] See, e.g., Joskow (2011)

between the two (adjacent or overlapping) years compared, much more important in most cases.[19]

Table 1 also provides information on the facilities' value factors. Perhaps the most striking result in Table 1 is the substantial difference between the average value factors for wind and solar units. Table 1 indicates that on average, relative to unweighted average spot prices, a kwh from a solar facility was worth about 32% more than a kwh from a wind facility. This result mainly reflects the strong, but not perfect, diurnal correlation between solar radiation and demand for electricity. This result implies that, all else equal, solar generators could attain commercial viability (or "grid parity") at a substantially higher LCOE than wind generators.

As Figure 1 indicates, only two of the wind facilities had value factors much above one in 2011.[20] 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 (ρ = -0.25), which is visible in Figure 1. The relation is considerably stronger if the three ERCOT outliers just mentioned are removed (ρ = -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 CF in both time-series and cross-section dimensions.

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 (ρ = -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 probabilities of generation than the other facilities and much higher values of CF^+ than all but one other facility.

[19] 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.

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

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.[21] 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 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 capacity factor in the region not also have the highest value-adjusted capacity factor.[22] Taking the logarithm of the identity defining *VCF* and decomposing the variance of log(*VCF*) for wind plants as above reveals that variations in log(*CF*) and log(ψ) 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.

Table 2 shows substantial regional differences in the average capacity factors for facilities in each ISO. In the sample of wind facilities, the SPP average is nearly twice the average in ISONE and NYISO. The average CF of CAISO solar facilities was more than twice the average for ISONE and PJM. Differences in value factors are less substantial, though the

[21] The National Renewable Energy Laboratory has compiled and makes available a great deal of location-specific information on wind and solar resources: http://www.nrel.gov/renewable_resources/

[22] Lewis (2010) used actual spot prices plus wind speed data to compare alternative sites in Michigan and concluded that the best sites were generally those with the best wind speed.

high ψ for wind facilities in ISONE does stand out. The final row in Table 2 shows the fraction of the total sample variance for each measure accounted for by differences between ISO averages, as opposed to within-ISO differences.[23] Almost all the variance in CF among plants in our small solar sample reflects the difference between performance in CAISO and elsewhere, and over two-thirds of the variance in ψ reflects regional differences. Regional effects are less important on this measure for wind facilities, but more than two-thirds of the variance in CF reflects differences between rather than within ISOs.[24]

Together, Table 1 and 2 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, most historical wind data were collected at ground level or slightly above, well below the hub height of wind turbines.[25] It is thus possible that poor data led to poor siting decisions for some wind plants. Second, wholesale electricity prices tend to be higher in ISONE and NYISO than in the Midwest or California,[26] 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, some states' renewable portfolio standard (RPS) programs (notably, Texas's) require the facilities to be located in-state; all except Colorado require them to be located relatively nearby. The rationale is generally economic development, but the result is clearly sub-optimal generator locations from the point of view of the nation as a whole.

Table 3 decomposes the early/late changes in value factors per equation (6). While only 13 of the 25 wind-plant ψs 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

[23] In a standard F-test of the null hypothesis of no differences among ISO means, the 5% critical values correspond to fractions of 0.47 for wind facilities and 0.95 for solar facilities. The 1% critical value for wind corresponds to a between-ISO share of 0.57.
[24] 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, Table 2 shows that 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.
[25] Gunturu and Schlosser (2011) discuss the available wind energy data. Note that this (potential) problem does not arise for solar facilities.
[26] The FERC provides information on average prices at selected points within the ISOs: http://www.ferc.gov/market-oversight/mkt-electric/overview.asp.

facilities. Table 3 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. In short, even though we have at most two years of data, a general movement of prices against wind and solar generators is visible. As Hirth (2013) has argued, this is exactly what would be expected to occur as penetration of these non-dispatchable facilities with positively correlated outputs increases.[27]

3.3 Some Possible Correlates of Value Factors

Table 4 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 Table 4 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.

In the U.S., system peak loads generally occur in the summer, so it is interesting to examine the ratio of generation in summer hours to generation in other hours. 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: one in ISONE and all three in CAISO.[28] 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.

[27] See also Mills and Wiser (2013). Correlation of generators' outputs within ISOs is discussed in Section 3.
[28] The CAISO result is consistent with Fripp and Wiser (2009), who find that winds in coastal passes in California are stronger in the summer that in other seasons.

Wind or solar generators reduce the need for other capacity to the extent that they produce during periods when system-level demand is high. Lacking system-level demand data, we can investigate 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.[29] Nonetheless, whatever the cause of high spot prices, they signal a high social value of incremental generation. To examine the contributions these wind and solar facilities made to system output in peak-price periods, I found a price level for each facility such that it was exceeded for about 100 hours in 2011 and considered generation during those hours.[30]

Table 4 gives the ratio of average generation during these peak-price hours to average generation at other times. The contrast between wind and solar plants is dramatic. Wind facilities generated on average 28% less during peak 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 inter-temporal variability, presumably in part because, as noted above, high prices can be produced by supply-side outages as well as high levels of system demand.[31]

3.4 Negative Spot Prices

All generators outside ISONE faced negative spot prices for at least 18 hours during 2011.[32] These prices often reflected transmission congestion, and it is well known that Texas has for some time lacked adequate transmission capacity between its wind generators in the west and its load centers in the east.[33] Thus it is not surprising that one of the western ERCOT units had the highest incidence in this sample of negative prices in 2011: 1542 hours, or just under 18% of the

[29] A good deal of graphical information on ISO-level daily generation and average spot price is provided by the FERC: http://www.ferc.gov/market-oversight/mkt-electric/overview.asp. See also Figure 7.1 in MITEI (2011).
[30] The range was 95 to 104 hours.
[31] If the peak-price periods in this sample were mainly peak load periods, one could imagine using peak-price-period capacity factors to assign capacity credits in some market designs. But these capacity factors varied between 0.08 (ERCOT West) and 0.46 (ERCOT Coast) for wind plants and between 0.12 (ISONE) and 0.30 (CAISO) for solar facilities in 2011. These large differences suggest the danger in relying on averages instead of site-specific information in such contexts.

year. The other two western ERCOT units faced negative prices in 857 and 529 hours during 2011. But other 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 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 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. 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, as Tropical Storm Irene was moving north through New England, causing outages and reducing air conditioning demand. Because of ISONE's unusual market structure during this period, the analysis in the remainder of this section excludes the facilities in that ISO.

Table 5 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.[34] 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 5 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.

---

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

[33] See, e.g., Weiss et al (2013, pp. 22-23).

[34] 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. At any rate, it appears likely that wind generation during negative-price periods would have been somewhat greater in the absence of curtailments. 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.

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 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 5 shows that the solar plants in our PJM and CAISO samples faced negative spot prices roughly as often on average sa wind plants in 2011. But the average hides a large inter-ISO difference: 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. But 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.

A major reason for this difference is that, as just noted, negative prices are more likely to occur at night, when solar facilities don'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 5 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.[35]

Electricity generation in the face of negative prices is an unintended consequence of governments' tax and subsidy policies. The most important federal policy in support of wind generation is the production tax credit, which provides a reduction in corporate tax liability of $0.023 per kwh generated, regardless of when that output is produced. Thus as long as the

[35] 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.

production tax credit plus the spot price is positive, wind generation, which has essentially zero marginal cost, is definitely profitable at the margin.[36] And as wind generation increases in importance, one can expect the incidence of negative spot prices to increase.[37]

The major federal policy in support of solar generation, however, is an investment tax credit, which is proportional to initial investment cost and does not depend on output. So why do solar facilities generally not shut down when the sun is shining but spot price of their output is negative? The answer is most likely state RPS programs. Solar units are everywhere classified as renewable, and 17 of the 30 RPS programs explicitly favor or require solar generation. Thus solar generators often receive significant RPS-based compensation over and above the spot price. All U.S. RPS regimes treat all kilowatt-hours as equivalent, regardless of when they are generated, and one would expect RPS-based compensation to follow suit. Thus state RPS systems provide solar facilities an incentive to generate even when the value of incremental electricity is negative, and they provide wind facilities an additional incentive to do so.

4. Output Variability

At low levels of penetration, wind and solar power simply reduce the net demand that must be met by dispatchable generators during some hours and, typically make that residual demand somewhat more variable. At higher levels of penetration, significant changes in system operations and the composition of the dispatchable generation fleet may be required, and both storage facilities and measures to make demand more responsive to system conditions will likely become more attractive.

Quantitative analyses of the latter changes and their costs have necessarily involved modeling an actual or hypothetical system in detail and examining how the presence of

[36] One might think that the occurrence of negative prices establishes that most facilities in this sample lack significant market power, since a monopoly generator could simply reduce output until the market price became positive. Unfortunately, it is trivial to show that if a monopoly generator faces a linear demand curve for its output, a negative price will maximize its profit if the per-kWh subsidy exceeds the zero-output price.

[37] 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.

additional wind and/or solar capacity would alter optimal operating and investment decisions.[38] Such an analysis is clearly beyond the scope of this study. There is nonetheless value in understanding how the outputs of wind and solar plants vary over time and how effectively geographic averaging (with adequate transmission, which is assumed here) across plants that resemble those in our sample can reduce system-level variability in various regions.

This section considers two dimensions of variability. Section 4.1 presents four standard-deviation-based measures of hour-to-hour and day-to-day changes in output. These measures should shed light on the implications of increased wind or solar penetration for the need for flexibility in the rest of the system (including the demand side).[39] Section 4.2 develops a measure of the effectiveness of geographic averaging for reducing output variability and computes it for the plants in each ISO using the measures developed in Section 4.1. Because there is generally some reduction in variability from geographic averaging, it is not appropriate to model the impact of adding incremental wind or solar capacity by simply scaling up current wind or solar output.[40]

Section 4.3 considers the incidence of low or zero output, which should shed light on the requirements for backup generation. Individual facilities are first considered in isolation, and the relation between low or zero output and the plant and ISO levels is then investigated.

4.1 Changes in Output

Perhaps the most 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:[41]

(7a) $$Vh \equiv \sigma\left(Q_h - Q_{h-1}\right) \big/ \overline{Q}_{2011},$$

[38] See the references cited in note 8, above.
[39] I make no claim that these measures are in any sense optimal, but they do seem easy to interpret. For alternative measures see, e.g., Ela and O'Malley (2012).
[40] Gunturu and Schlosser (2011) provide a useful overview of the literature on geographic averaging of wind generation.
[41] Using the standard deviation as a measure of variability would most obviously be appropriate if output changes were normally distributed. In Section A.3 of the Data Appendix it is argued that normality is a reasonable approximation.

Where $\sigma(\bullet)$ denotes the standard deviation of the quantity in parentheses, Q_h is output in some hour h, and $\overline{Q}_{2011}$ is average hourly generation in 2011.[42] It is also of interest to consider day-to-day variability by focusing on differences between output in the same hours in adjacent days:

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

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 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:[43]

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

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

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

[42] 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.

[43] While hourly data may be adequate to capture most variability in wind power, the output of photovoltaic facilities can change from minute to minute as clouds pass over, and these changes can pose serious operational challenges; see Pérez-Arriaga (2010).

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.

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 the corresponding coefficients of variation were much larger for solar plants than for wind plants, which themselves showed substantial variation in cross-section and between early and late periods.[44] 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 desirable, all else equal, but I am unaware of collections of historical data on wind energy variability that could be used to help select 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 different sites to remain constant.

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

[44] All the Early/Late comparisons in Table 6 involve 8760-hour periods, but since computation of all these quantities involves losing observations at the beginning of the sample (*Vh* and *Vd*) or at both the beginning and the end of the sample (*Vht* and *Vdt*), these periods overlap somewhat in all cases.

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.

Table 7 gives the average values of our output variability measures for the wind and solar plants in each ISO. Since each of the numbers in Table 7 reflects at most five facilities, one should not read too much into differences among them. Still, this Table does demonstrate 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. 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. The last line of Table 7, like the last line of Table 2, gives the percentage of the total sample variance of each measure associated with differences between ISO means, as opposed to differences within the ISOs. It seems clear that ISO-level differences are important in all cases and account for the bulk of the total sample variance except for *Vh* and *Vht* for wind plants, where most of the variation is within ISOs.[45]

4.2 Geographic Averaging

The goal of this section is to use the sample of plants in each ISO to investigate the potential of geographic averaging to reduce the variability of system-level total output, using the four measures introduced above. In order to eliminate the effects of differences in the size distributions in our within-ISO samples, we first re-scale each facility’s output so that all in each ISO have the same average size and then compare the variability of the sum of these re-scaled

[45] But see note 23, above: the usual F-test would reject the null hypothesis of no differences among ISOs only for *Vd* and *Vdt* for wind plants.

outputs to what it would have been under two polar case assumptions about the relations between the individual 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.[46] The average correlations for each ISO are shown in the first column of Table 8, and the ISOs are listed in decreasing order of those averages. The individual correlations differed more than the averages shown might suggest: 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 will depend on the number of plants involved and their relative scales, as well as the properties of the stochastic process that generates their outputs. As noted above, to focus on the last of these it is necessary first to adjust for differences in facility scale. 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 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:

$$(8) \qquad Q_h^T = Q_h^1 + \sum_{i=2}^{N} (A^1 / A^i) Q_h^i, \qquad \textit{for all h.}$$

One can then compute each of the four variability measures defined by equations (7) for each ISO's 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 Q^T, Z^T, will be the function of the Z^i given by equation (8). That is,

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

(9) $$Z_h^T = Z_h^1 + \sum_{i=2}^{N} (A^1 / A^i) Z_h^i, \qquad \textit{for all h.}$$

If the Z^i were uncorrelated, the variance of Z^T would be the sum of the variances of the terms on the right of (9). It follows that in this first polar case the aggregate variability measure based on Z^T would be given by

(10a) $$\left.\frac{\sigma\left(Z^T\right)}{NA^1}\right|_{\rho=0} = \frac{1}{NA^1}\left[\sigma^2\left(Z^1\right) + \sum_{i=2}^{N}\left(A^1 / A^i\right)^2 \sigma^2(Z^i)\right]^{1/2} = \frac{1}{\sqrt{N}}\left\{\frac{1}{N}\sum_{i=1}^{N}\left[\frac{\sigma\left(Z^i\right)}{A^i}\right]^2\right\}^{1/2}.$$

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 (10a) is the root-mean-square of these plant-level statistics, divided by the square root of N. 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:

(10b) $$\left.\frac{\sigma\left(Z^T\right)}{NA^1}\right|_{\rho=1} = \frac{1}{N}\sum_{i=1}^{N}\frac{\sigma\left(Z^i\right)}{A^i}.$$

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

(10c) $$R\left(Z^T\right) = \left[\left.\frac{\sigma\left(Z^T\right)}{NA^1}\right|_{\rho=1} - \frac{\sigma\left(Z^T\right)}{NA^1}\right] \Bigg/ \left[\left.\frac{\sigma\left(Z^T\right)}{NA^1}\right|_{\rho=1} - \left.\frac{\sigma\left(Z^T\right)}{NA^1}\right|_{\rho=0}\right].$$

With positive correlations among plants outputs, one expects *R* generally to be between zero and one, with higher values indicating greater gain from geographic averaging. With plants of roughly equal size, when *R*=1 aggregate variability using the measures presented here should fall roughly as the square root of the number of plants, while when *R*=0 aggregate variability should remain roughly constant when new plants come on line.

Table 8 presents these ratios and their averages for the variability measures in Table 7, where *Rh* corresponds to *Vh* and so on.[47] Given the high correlations among the plant-level variability measures noted above, one might expect high correlations among the ratios shown for wind plants in Table 8, and correlations across ISOs in fact range from 0.62 (*Rht* and *Rdt*) and 0.99 (*Rd* and *Rdt*). Looking across ratios for wind plants, ISONE would seem generally to have less to gain from geographic averaging than the other ISOs, particularly the four appearing below it in the Table. Looking at solar plants, PJM stands out as seeming to have less to gain from geographic averaging 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 difference too seriously, of course.

One might also expect that higher average output correlations, shown in column 1, would generally lead to lower gains from averaging and thus lower values of the ratios shown in columns 2 – 5. While all but *Rht* are in fact negatively correlated with the average correlations in column 1, these correlations are substantial only for *Rd* (-0.66) and *Rdt* -(0.68). It appears that one cannot reliability infer potential gains from geographic averaging from correlations among plants' outputs, particularly at the hourly timescale.

Finally, except for the 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, the assumption of statistical independence seems a good approximation for *Vht* except for ISONE and for *Vh* except for SPP, ISONE and CAISO.

---

[47] 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 8 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 8 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.

4.3 No or Low Generation

One measure of the need to provide backup capacity for wind or solar facilities 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. I deal first with wind and then with solar.

The within-ISO averages for hours with zero generation from the wind plants in our sample in 2011 are shown in the first column in Table 9. Every plant had at least 100 hours with zero output, and all but two (one in CAISO and one on the Gulf coast in ERCOT) had at least 300. The average across all units was 948 hours, just under 11% of the year. It would be interesting to know how many of these zero-output hours were the result of ISO curtailment orders. We cannot observe this directly but the available evidence suggests that most zero-output hours occurred either because there was simply not enough wind. 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 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.[48]

The fourth column in Table 9 shows that on average across facilities within each ISO, most hours with zero 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).

[48] 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 and so would find production profitable at even lower 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.

The rest of Table 9 is concerned with episodes of zero output at all sample plants in each ISO.[49] The third column shows that in 2011 only ERCOT had no hours in which none of the wind plants in the sample were generating, and only it and CAISO had fewer than 40 such hours. 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 in such hours. 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. Except for ISONE, the overwhelming majority of ISO-wide zero generation hours 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.[50]

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.[51]

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. The figures in the last two columns in Table 9 shed light on this matter. 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

[49] See Gunturu and Schlosser (2011) for one study of regional calm periods with little or no wind generation possible and a discussion of others.
[50] In 2011 wind plants in ISONE produced only 7% more per hour at night than at other hours, a smaller difference than in any other ISO, and well below the 15% average difference shown in Table 4.
[51] 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 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 9.

corresponding plants' probabilities. The last column in Table 9 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 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.[52]

Thus the last two columns of Table 9 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. Moreover, some regional low-wind events, a large fraction in some regions, will last more than an hour or two. Based on 30 years of wind energy data, Gunturu and Schlosser (2011) find that geographic averaging among wind plants should be less valuable in the center of the country (SPP and MISO) than near the coasts (NYISO, ISONE, CAISO, and PJM), with ERCOT in between these two groups. Except for ERCOT, the statistics in the last column of Table 9 seem broadly consistent with this finding.

Let us now turn to solar plants. They 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 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

[52] 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.

same hour in the two adjacent days, defining "low generation" as an output less than half the mean of the two comparison hours. The results of this inquiry are summarized in Table 10.

The first column indicates that incidents of low generation relative to adjacent hours are not common even at the plant level – the plant-level average in 2011 was 133 hours, 1.5% of the hours in the year. The second column shows that such events are even less common at the ISO level. Still, the statistics in the third and fourth columns, which are computed as in Table 9, indicate that plant-level hour-to-hour low-generation events, which one can think of as reflecting hours that are much cloudier than adjacent hours, are not statistically independent.

The right-hand half of Table 10 shows that hours with low generation relative to the same hour in adjacent days were much more common at both plant and ISO levels, particularly in PJM. Independence is again strongly rejected, and probability ratios are substantial. 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.

At least in this small sample of solar facilities, low hourly generation relative to adjacent hours would not appear to be a concern at the ISO level – not because plant-level episodes of low generation are statistically independent, since they do not seem to be, but simply because they are so rare. Low generation relative to the same hour in adjacent days, likely reflecting cloudy days rather than cloudy hours, is a more serious matter. The last two columns in Table 10 suggest that the scope for geographic averaging to reduce this measure of variability is limited. On the other hand, cloudy days may be easier to forecast than cloudy hours and, again, we have a very small sample.

5. Concluding Remarks

This study has used a unique dataset to produce a large number of results on the performance of individual wind and solar generators across the US – some new results and some that confirm the prior literature. In interpreting these results, it should be kept in mind that, as noted in the introduction, these data are incomplete on important dimensions. And they can shed no direct light on performance of wind and solar generators in regions that still lack organized wholesale electricity markets. But they do have some clear implications for the design of policies to support use of wind and solar energy.

Perhaps the most striking result regarding the *value* of generation from wind and solar energy was that in this sample in 2011, a kilowatt-hour produced by a solar facility was on average worth about 32% more than a kilowatt-hour from a wind plant. The average value of wind and solar output relative to baseload output may have declined during the sample period as more of both non-dispatchable technologies were deployed. Wind capacity factors were generally, but not always, above solar capacity factors, and in both cases regional differences accounted for most of the dramatic within-technology differences in capacity factors.

True to stereotype, wind plants tended to generate more at night than in the day and less in the summer that in other seasons, but there were exceptions to both generalizations, and the average differences were not dramatic. Solar plants all generated more on average during peak-price periods than at other times, while wind plants generally produced less. Outside ISONE, which alone barred negative bids, all of the plants in the sample faced negative spot prices during 2011, and 12 faced negative prices for more than 500 hours. Both wind and (during the day) solar plants generally produced positive outputs during hours with negative prices – they paid the market to take their electricity.

Two dimensions of the *variability* of output from wind and solar facilities were studied: hour-to-hour and day-to-day variation in generation and the incidence of low or zero output. Different measures of variability are generally highly correlated across plants and show substantial variation. On average, short-term trends or ramping events accounted for over a third of the hour-to-hour changes in wind plant output. Both wind and solar plants showed more day-to-day than hour-to-hour changes, and for wind plants differences among ISOs were considerably more important for day-to-day changes than for hour-to-hour changes. On the other hand the potential gain from geographic averaging seems greater for hour-to-hour variability, at least for wind plants.

Wind plants in this sample averaged 948 hours a year with zero generation, mainly in episodes of three or more hours. While hours with zero generation in all sample plants in an ISO were considerably rarer, such hours occurred much more frequently than if plant-level zeros were statistically independent. If generation was positive in the adjacent hours, solar plants rarely produced less than half the mean output of those hours, and instances of zero generation in such cases were very rare indeed.

All but one of the tables discussed above provide information on the cross-section *variation* of various dimensions of individual generating plant or ISO average performance, and half the tables provide information on variation between early and late periods. Figures 1 and 2 depict the cross-section variation in value and capacity factors, and Table 2 demonstrates that regional differences drive the striking variation in value factors. Variation on some other dimensions of performance is also substantial among plants in 2011, between early and late periods, and, in some cases, between ISOs. Most wind plants generate less in the summer than other seasons, for instance, but not those in CAISO. Most wind value factors are less than those of constant-output baseload plants, but not those of the two coastal plants in ERCOT. For other technologies, historical averages may be good predictors of the performance of new facilities, but that is clearly not true for wind or solar generation. Site selection is important, regions differ, and performance varies over time.

Two of the patterns noted above that seem particularly robust have clear implications for public policy. First, when spot prices are negative and they can generate, wind and solar plants generally do so. One can debate whether it is desirable to subsidize renewable generation at all, since doing so is clearly more costly than taxing emissions of carbon dioxide as a means of slowing climate change and is a similarly inefficient approach to reducing local air pollution. But there can be no doubt that encouraging renewable generation when its marginal value to the electric grid is negative raises costs to society, but that is what both the federal production tax credit (for wind) and state RPS programs (for both wind and solar) do – along with most of the feed-in-tariff schemes in widespread use outside the US.[53] In regions with organized wholesale markets and nodal pricing, it would be more efficient to pay output subsidies only when the spot price is positive or even to make them proportional to the spot price. In regions that have not yet adopted this modern design, however, there is no obvious way to provide incentives for wind or solar generators to reduce their output when it has negative social value.

[53] See Schmalensee (2012) and the references there cited. It should also be noted that the efficiency of the production and investment tax credits are further reduced because firms without substantial taxable income must engage in tax equity financing in thin markets with high transactions costs (Bipartisan Policy Center, 2011). But, of course, tax expenditures don't look like spending at first glance, and their costs are generally well hidden.

A second robust pattern is the huge regional differences in facility performance – most clearly the capacity factor differences presented in Table 2 and depicted in Figures 1 and 2. One important reason why 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 not clear that these limits can ever in fact have much impact on any state's employment. But it is clear 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. If there is a national interest in subsidizing the generation of electricity from wind and solar power, a national RPS program or feed-in-tariff would give a much higher return per dollar spent than a collection of state plans that restrict generator siting.[54]

References

Baker, Erin, Meredith Fowlie, 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 athttp://www.ucei.berkeley.edu/PDF/csemwp176.pdf.

__________ (2012). "The Private and Public Economics of Renewable Electricity Generation," *Journal of Economic Perspectives*, 26:1 (Winter), pp. 67-92.

[54] A feed-in tariff would make the cost of the subsidy visible, while RPS programs and tax expenditures hide those costs.

Coughlin, Katie, and Joseph H. Eto (2010). "Analysis of Wind Power and Load Data at Multiple Time Scales," Lawrence Berkeley National Laboratory, Report LBNL-4147E (December).

Ela, Erik, and Mark O'Malley (2012). "Studying the Variability and Uncertainty Impacts of Variable Generation at Multiple Timescales," *IEE Transactions on Power Systems*, 27:3 (August), 1324-33.

Fripp, Matthias, and Ryan W. Wiser (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 (May), 477-485.

Green, Richard, and Nicholas Vasilakos (2012). "Storing Wind for a Rainy Day: What Kind of Electricity Does Denmark Export?" *The Energy Journal,* 33:3, pp.

Gunturu, Udaya Bashar, 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).

Hirth, Lion (2013). "The Market Value of Variable Renewables," *Energy Economics,* 38 (July), pp. 218-236.

Hogan, William W. (1992). "Contract Networks for Electric Power Transmission," *Journal of Regulatory Economics*, 4:2, pp. 211-242.

Holttinen, 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 (March), 179-92.

Hsu, Michael (1997). "An Introduction to the Pricing of Electric Power Transmission," *Utilities Policy*, 6:3 (September), 257-270.

Huntowski, Frank, Aaron Patterson, and Michael Schnitzer (2012). "Negative Electricity Prices and the Production Tax Credit," The NorthBridge Group (September), report available at http://www.nbgroup.com/publications/Negative_Electricity_Prices_and_the_Production_Tax_Credit.pdf.

Joskow, Paul L. (2011). "Comparing the Costs of Intermittent and Dispatchable Electricity Generating Technologies," *American Economic Review*, 100:3 (May), pp. 238-241.

__________ and Richard Schmalensee (1987). "The Performance of Coal-Burning Electric Generating Units in the United States," *Journal of Applied Econometrics*, 2:2 (April), pp. 85-109.

Kopsakangas-Savolainen, Maria, and Rauli Svento (2013). "Economic Value of Intermittent Power Generation," Finnish Environment Institute, Working Paper (February), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2226363.

Lamont, Alan D. (2008). "Assessing the Long-Term System Value of Intermittent Electric Generation Technologies," *Energy Economics*, 30, pp. 1208-1231.

Lewis, Geoffrey McD (2010). "Estimating the Value of Wind Energy Using Electricity Locational Market Price," *Energy Policy*, 38:7, pp. 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.

Pérez-Arriaga, Ignacio J. (2010), "Managing Large-Scale Penetration of Intermittent Renewables," MIT Energy Initiative, Symposium Framework Paper, available at http://mitei.mit.edu/system/files/intermittent-renewables-whitepapers.pdf.

Schmalensee, Richard (2012). "Evaluating Policies to Increase the Generation of Electricity from Renewable Energy," *Review of Environmental Economics and Policy*, 6 (Winter), pp. 45-64.

Ueckerdt, Falko, Lion Hirth, Gunnar Luderer, and Ottmar Edenhofer (2013). "System LCOE: What are the Costs of Variable Renewables?" Potsdam Institute for Climate Impact Research, Working Paper (January), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2200572.

U.S. Energy Information Agency (2011). *Direct Federal Financial Interventions and Subsidies in Energy in Fiscal Year 2010*, Washington, DC (July), available at http://www.eia.gov/analysis/requests/subsidy/pdf/subsidy.pdf.

__________ (2012a). *2011 Wind Technologies Market Report*, Washington, D.C. (August), available at http://www1.eere.energy.gov/wind/pdfs/2011_wind_technologies_market_report.pdf.

__________ (20112b). Annual Energy Review 2011, Washington, DC (September), available at http://www.eia.gov/totalenergy/data/annual/index.cfm.

Weiss, Jurgen, Heidi Bishop, Peter Fox-Penner, and Ira Shavel (2013). "Partnering Natural Gas and Renewables in ERCOT," The Brattle Group (June), report available at http://www.texascleanenergy.org/Brattle%20report%20on%20renewable-gas%20FINAL%2011%20June%202013.pdf.

Data 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.[55]

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

[55] Recall that 2012 was a leap year.

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.

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.

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:

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

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.[56]

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. In an attempt to provide a more intuitive measure of tail heaviness, Table A.1 presents ISO-specific averages of *ckurt*, defined as follows:

[56] 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.

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

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.[57]

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 ISO's 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*

[57] 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.

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.

A.4. Output Change Distributions

In a study of changes in aggregate wind output in ERCOT, PJM, and NYISO using high frequency data (output measured at 4 to 6 second intervals), Coughlin and Eto (2010) found evidence of substantial departures from normality in the direction of heavy tails for ERCOT and PJM at time intervals ranging from a few seconds to just under an hour. To see if output changes in our plant-level data also departed substantially from normality, I computed the standard measures of skewness and kurtosis as well as *rskew* and *ckurt* for hour-to-hour output changes ("hourly changes") and (as in equation (7b)) changes from the same hour in the preceding day ("daily changes") using output data for 2011.

For hourly changes, 22 of the 25 distributions were right-skewed, but the largest value of *rskew* was 1.13, and all the others were less than 1.11. At the ISO level, the highest average was 1.06 (CAISO), well below all the corresponding statistics in Table A.1. For daily changes, 15 of the 25 distributions were *left*-skewed, but, again, departures from symmetry seem small. Values of *rskew* ranged from 0.93 to 1.10, and ISO-level averages ranged only from 0.98 to 1.04.

Departures from normality in the direction of heavy tails seem only a little bit more important, particularly at the hourly time-scale. All hourly output change distributions were leptokurtic according to the standard measure, and *ckurt* ranged from 1.32 to 1.79. ISO-level averages ranged from 1.41 to 1.60 (MISO). While 20 of the 25 daily change distributions were also leptokurtic, the largest value of *ckurt* was 1.28, and the largest ISO-level average was 1.25 (CAISO).

In short, at both time-scales, but particularly at the daily scale, normality seems a fairly good approximation for wind output change distributions – certainly a better approximation than for price distributions.

Table 1. Summary Measures of Generator Performance

	In 2011		Coefficient of Variation	
	Mean	Range	In 2011	Early/Late*
Wind Generators				
CF	31.2	18.8 - 42.9	0.22	0.13
CF^{+}	34.9	22.4 - 46.8	0.20	0.11
Pr(Q>0)	0.89	0.79 - 0.99	0.06	0.03
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
Solar Generators				
CF	14.1	6.91 - 25.0	0.52	0.18
CF^{+}	34.6	17.6 - 49.8	0.39	0.28
Pr(Q>0)	0.42	0.19 - 0.52	0.31	0.04
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

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

Table 2. Average Capacity and Value Factors, 2011

	Wind Plant Average		Solar Plant Average	
Region	CF	ψ	CF	ψ
ISONE	0.23	0.99	0.11	1.12
NYISO	0.24	0.92		
PJM	0.31	0.86	0.09	1.21
MISO	0.33	0.85		
SPP	0.42	0.87		
ERCOT	0.35	0.86		
CAISO	0.27	0.84	0.24	1.13
between % of σ^2	68.8	10.9	93.1	67.5

Table 3. 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
$\Delta\psi_q$	0/25	0.002	2/9	0.050

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

Table 4. Some Possible Correlates of Value Factors: Output Ratios

	In 2011		Coefficient of Variation		Correlation
	Mean	Range	In 2011	Early/Late#	with ψ, 2011
Wind					
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
Solar					
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.
#Computed as the standard deviation of early-period to late-period changes, divided by the 2011 mean of the corresponding variable.

Table 5. Negative Prices and Value Factors*

	In 2011		Coefficient of Variation		Correlation
	Mean	Range	In 2011	Early/Late[#]	with ψ, 2011
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					
Wind:					
P < 0	0.04	0.002 - 0.18	1.08	0.73	-0.72
Q > 0 when P < 0	0.92	0.60 - 1.00	0.11	0.12	-0.06
Solar:					
P< 0	0.04	0.004 - 0.08	0.99	0.43	0.95
Q > 0 when 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 when 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.

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

Table 6. Meaures of Plant-Level Ouput Variability				
	In 2011		Coefficient of Variation	
Measure	Mean	Range	In 2011	Early/Late*
Wind Plants				
Vh	0.32	0.16 - 0.48	0.22	0.11
Vd	1.12	0.79 - 1.52	0.17	0.12
Vht	0.21	0.09 - 0.33	0.27	0.13
Vdt	0.94	0.60 - 1.30	0.18	0.13
Solar plants				
Vht	0.68	0.25 - 1.64	0.68	0.25
Vdt	1.61	0.43 - 3.96	0.71	0.23

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

Table 7. ISO Averages of Plant-Level Output Variability Measures, 2011						
	Wind Facilities				Solar Facilities	
Region	*Vh*	*Vd*	*Vht*	*Vdt*	*Vht*	*Vdt*
ISONE	0.38	1.35	0.26	1.14	1.15	2.85
NYISO	0.33	1.36	0.22	1.34		
PJM	0.32	1.17	0.21	0.98	0.59	1.40
MISO	0.34	1.13	0.21	0.97		
SPP	0.31	1.05	0.20	0.89		
ERCOT	0.31	0.93	0.19	0.77		
CAISO	0.22	0.97	0.13	0.78	0.30	0.59
between % of σ^2	37.3	70.6	42.5	68.9	65.5	74.6

Table 8. Potential Gains from Geographic Averaging, All Data

	Wind Facilities					Solar Facilities	
Region	mean ρ*	*Rh*	*Rd*	*Rht*	*Rdt*	*Rht*	*Rdt*
NYISO	0.61	0.88	0.39	0.97	0.41		
SPP	0.59	0.83	0.39	0.96	0.40		
ISONE	0.42	0.61	0.17	0.67	0.18	0.86	0.68
CAISO	0.41	0.82	0.72	0.92	0.77	0.63	0.80
PJM	0.35	0.96	0.79	0.99	0.81	0.47	0.28
MISO	0.27	0.91	0.77	0.94	0.80		
ERCOT	0.20	0.92	0.74	0.91	0.80		
*Average***	*0.41*	*0.85*	*0.57*	*0.91*	*0.60*	*0.66*	*0.59*

*Average of pairwise correlations of hourly outputs of wind plants in the ISO.

**Averages of correlations are across all plants in the sample; other averages are across ISOs.

Table 9. Hours with No Wind Generation, 2011

	Hours with Zero Generation, 2011			% Q=0 Hours in 3+ Hour Spells			Are Plant Zeros Independent**	
		All Plants						
ISO (# plants)	Plant Average	Hours	% in Day	Plant Average	All Plants	Longest Q=0 Spell, All Plants**	Z-Statistic	Probability Ratio
NYISO (3)	1136	178	80	87	76	19	56.6	11.3
ISONE (3)	1424	173	76	90	71	15	27.9	4.0
PJM (3)	1267	106	89	86	68	13	19.9	4.2
SPP (3)	616	74	92	80	47	6	51.0	19.3
CAISO (3)	744	16	100	79	21	4	13.4	11.3
MISO (5)	1096	11	100	85	45	4	29.8	42.4
ERCOT (5)	532	0	–	80	*	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 were computed using all available data, not just 2011 data.

Table 10. Hours with Low Solar Generation, 2011

	Relative to Adjacent Hours				Relative to Same Hours in Adjacent Days			
	Hours in 2011		Plants Independent?#		Hours in 2011		Plants Independent?#	
ISO	Plant Average	All Plants	Z-Statistic	Probability ratio	Plant Average	All Plants	Z-Statistic	Probability Ratio
ISONE	97	1	*	145	530	82	82	74
PJM	188	14	60	186	870	473	199	54
CAISO	114	0	**	41	349	18	55	56

*One low hour was observed in the entire sample: November 13, 2011, 11:00am - 12:00pm. Under independence the probability of one or more low hours would be 0.0069.

**One low hour was observed in the entire sample: June 27, 2012, 6:00am - 7:00am. Under independence the probability of one or more low hours would be 0.024.

#Statistics in the columns below were computed using all available data, not just 2011 data.

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

	Average Across Nodes Within Each ISO							Percentage of Hours with Price Range	
					Correlation of $P^i(h)$ with				
ISO (# of Nodes)	Price, $/MWh	Coeff. of Variation	rskew*	ckurt*	$P^i(h\text{-}1)$	$P^i(h\text{-}24)$	$P^j(h)$	< $1	> $10
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 1. Value Factor (*ψ*) v. Capacity Factor (*CF*): Wind, 2011

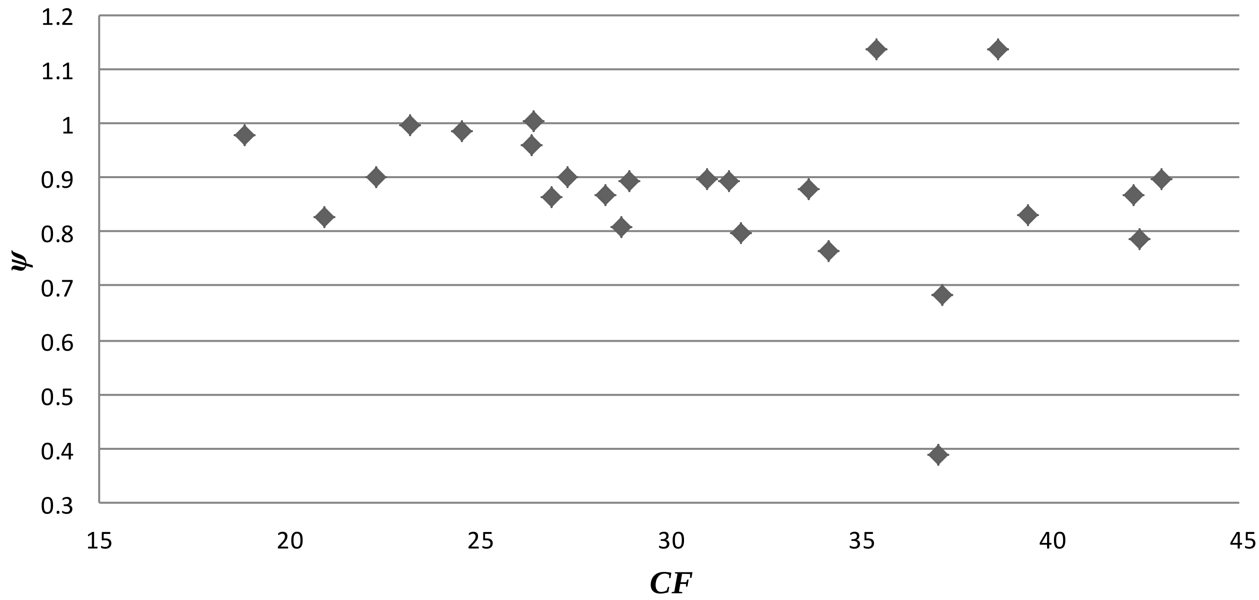

Figure 2. Value Factor (*ψ*) v. Capacity Factor (*CF*): Solar, 2011