NOTES ON THE ECONOMICS OF ENERGY STORAGE

Geoffrey Heal

Notes on the Economics of Energy Storage
Geoffrey Heal
NBER Working Paper No. 22752
October 2016
JEL No. Q4,Q53

ABSTRACT

The increasing importance of intermittent renewable energy sources suggests a growing importance for energy storage as a way of smoothing the variable output. In this paper I investigate factors affecting the amount of energy storage needed, including the degree of intermittency and the correlations between wind and solar power outputs at different locations.

Geoffrey Heal
Graduate School of Business
516 Uris Hall
Columbia University
New York, NY 10027-6902
and NBER
gmh1@columbia.edu

Notes on the Economics of Energy Storage

Geoffrey Heal

October 13, 2016

Columbia Business School and National Bureau of Economic Research

Abstract

The increasing importance of intermittent renewable energy sources suggests a growing importance for energy storage as a way of smoothing the variable output. In this paper I investigate factors affecting the amount of energy storage needed, including the degree of intermittency and the correlations between wind and solar power outputs at different locations.

1   Introduction

The increasing importance of intermittent renewable energy sources such as wind and solar power has stimulated interest in energy storage. Storing energy from times of high output to times of low, is an obvious way of responding to intermittency. So far the rapid growth of renewable power has been managed without big investments in storage: in the US grids generally compensate for the intermittency of renewable power by backing it up with gas turbines, which are relatively inexpensive to construct and can vary their outputs rapidly to compensate for fluctuation in wind or solar output. In northern Europe, hydro power from Norway, which has a surplus of such power, is widely used as a balancing item to match supply and demand when wind output falls. The US also uses hydro, but pumped hydro, as a means of storage. However, if countries meet their greenhouse gas reduction targets under the 2015 Paris UNFCCC COP, they will need to phase out the use of fossil fuels in most electric power generation, meaning that gas will no longer be acceptable as a back-up for renewables, and as renewable penetration moves from 15-25% of grids to over half of all generating capacity, it will be necessary to make specific provisions for coping with the stochasticity of renewable power output.

There are several ways in which this might be approached. One, as noted, is to store energy from times of abundant production to those of low production. Another is to invest in technologies that are carbon-free and whose outputs can very quickly so as to fill in for solar and wind. Hydropower and geothermal power fit in this category, although in the US the scope for radical increases in

the outputs from these power sources seems to be limited. A third possibility is to install so much wind and solar capacity that there is never a shortage: even when the wind is low, there will be sufficient power. This of course implies a huge excess of power at some times. A variant on this is to draw on a geographically diverse set of wind and solar power plants in the hope that the covariance between their outputs is low or negative and pooling their outputs over a large area effectively smooths the output: this is the "diversified portfolio" approach. A fourth option is demand-side management, contracting with energy consumers to reduce their use of power at times of shortage. Currently the first and last of these options - storage and demand management - are emerging as the main carbon-free strategies for coping with stochastic power output in the US. The US currently has 24 pumped hydro storage systems, of which ten can generate more than one gigawatt of power,[1] and one compressed air energy storage system, and utilities are beginning to install grid-scale battery storage.

In an earlier paper (Heal 2016), I investigated how the US could meet its stated ambition of reducing greenhouse gas emissions by 80% by 2050. In these calculations, it emerged that the amount of energy storage needed, and its cost, were critical parameters. I estimate that it might be necessary to spend trillions of dollars on storage. I also noted there that we have no clear model to determine the amount of storage needed, and the point of this paper is to set out the basic issues with which such a model has to deal. I do this in the context of an economy where all energy is generated by intermittent renewable sources.

In reality there are two parts of the intermittency we experience from renewables: solar has a natural daily cycle, producing power in the day and not at night, according to a fully predictable schedule: there is nothing stochastic about sunrise and sunset. In addition intermittent cloud cover can cause its output to vary stochastically even during the day. And the output of wind is generally stochastic: there is a tendency for output to be greater at night than in the day, but it is never predictable over more than short periods. This implies that there are two parts to the energy storage problem: one is shifting energy produced during the day to the night, and the other is smoothing out the stochastic output of wind or solar power. The first is mainly about storing solar energy from day to night, and the second is largely about smoothing intermittent wind output. Here I am going to focus first on the predictable aspect of intermittency from solar power, its diurnal cycle. Throughout the paper I take the time path of demand for energy as given, although as mentioned above there is some scope for modifying this in the light of the availability of power through demand-side management. Increasing sophistication of power grids should increase this scope.

Dealing with the intermittency of renewable energy is not the only role for storage. Amongst the problems facing any electric grid are the massive diurnal and seasonal fluctuations in demand, caused for example by daytime peaks in hot regions as demand for air conditioning soars, accompanied by far lower

---

[1]See FERC, Pumped Storage Projects, http://www.ferc.gov/industries/hydropower/gen-info/licensing/pump-storage.asp

demands at night and in the winter. As a consequence much generating capacity is used only a fraction of the time, raising the effective cost of power from these sources. Inexpensive storage could allow generators to produce a more or less constant output, with surplus being sent to storage at times of low demand and shortfalls met from storage at times of high demand. This could be a valuable role even without the intermittency of renewables - provided storage costs are low enough.

2   Literature

The economic literature on energy storage is limited, not surprisingly given the recent emergence of the topic. I mentioned intermittency and the need for storage in a survey paper Heal (2010). Schmallensee (2014) has written about the management of intermittent power sources, and indicates clearly the seriousness of the problem that they pose. He studies the hourly output of wind and solar plants in seven ISOs covering around 2/3 of US electricity production for the year 2011, and finds a significant number of incidents of ISO-wide zero output from wind plants. He notes that "a large fraction of the hours of ISO-wide zero generation hours from wind facilities in this sample in 2011 were part of spells lasting at least three hours, although the importance of such spells was substantially lower at the ISO level that at the plant level." The last sentence here illustrates the importance of geographical diversification. Schmallensee also comments that "the three northeastern ISOs - NYISO, ISONE, and PJM - stand out not only in terms of the incidence of hours with zero output 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." These data show clearly that intermittency is a real problem, and can cause difficulties in providing reliable power to consumers. Bradbury (2013) does not talk about storage, but discusses in detail the smoothing effect of spatial diversification, and I use his results extensively below. Chang et al. (2010) do model storage explicitly in the context of renewable penetration in two ISOs, and their results suggest two conclusions - one that high levels of storage may be needed (much greater than assumed in Heal (2016)), and the second that there is a tradeoff between installing storage capacity and installing more renewable generation capacity. Denholm and Margolis (2016) also model storage explicitly, asking how much storage would be needed if California were to produce 50% of its power from solar PV generators. They emphasize that there are alternatives to storage - more flexible generation capacity (such as dispatchable nuclear, or CCGTs), power trade with neighboring regions, demand response and schedulable loads from electric vehicles. They suggest that with an optimal mix of all these alternatives, the amount of storage needed is small, less that one day's energy consumption. Gowrisankaran et al. (2015) also model the integration of solar power into a diverse grid. They make a point similar to that of Schmalensee, quoting Der Spiegel of January 16th 2012 to the effect that "For one week now the 1.1 million solar power systems in Germany have generated almost no electricity ..... As is so often the

case in winter, all solar panels more or less stopped generating electricity at the same time. To avert power shortages, Germany currently has to import large amounts of electricity." They compute the cost of intermittency for the grid in Arizona, note that spatial diversity reduces fluctuations in total solar output, and comment on the role of storage in stabilizing power output.

3 Time-Shifting Solar Output

I tackle the easiest problem first, that of shifting solar energy produced in the day to night-time. I work with a heavily stylized model and aim to get some idea of orders of magnitude. I assume there are two phases of demand for electric power, a base load of b megawatts, which runs from 8.00pm to 8.00am, and a daytime load of $d > b$ megawatts that covers the remaining 12 hours. Solar panels begin to generate power at 8.00am and cease generation at 8.00pm, peaking at 2.00pm: their output rises linearly to its peak and then falls linearly from this. The maximum output of the solar installation is denoted S megawatts: this is its nameplate capacity, assumed reached at 2pm. The total megawatt hours produced by the solar installation over a 24 hour period is $6S$. The total number of megawatt hours consumed over the same period is $12(b + d)$. For total production to equal total consumption over a 24 hour period we need these two amounts to be equal so that $S = 2(b + d)$. Maximum solar output has to be twice the sum of baseload and daytime loads. So if baseload is 30 gigawatts and the daytime load is 65 gigawatts we need a total of 190 gigawatts of solar power.[2]

Figure 1 shows this configuration: solar output is in red and load in black. How much storage capacity is needed? I assume that any solar output in excess of consumption between 8.00am and 8.00pm is stored to meet the later shortfall, in which case the storage capacity S needed is the area of the red triangle above the daytime load d. This area is

$$S = 3\left\{4(d + b) + \frac{d^2}{d + b} - 4d\right\} \tag{3.1}$$

Total daily consumption of power C is $12(d + b)$, and

$$\frac{S}{C} = 1 + \frac{d}{d + b}\left\{\frac{1}{4}\frac{d}{d + b} - 1\right\} \leq 1 \tag{3.2}$$

and for $b = 1$, $d = 2$ this is 0.44 – storage needed is 44% of daily consumption. For the Texas-like numbers used above, we have 38%: for plausible parameter values this formula gives us storage need of the order of one third to one half of a day's energy consumption for the time-shifting aspect of energy storage, taking predictable solar output from the daytime to the night. Hence if we were

---

[2]These figures correspond roughly to the numbers for Texas for 2015: maximum power use was 68,945 gigawatts and minimum 27,504. See ERCOT Hourly Load Data Archives, http://www.ercot.com/gridinfo/load/load_hist/
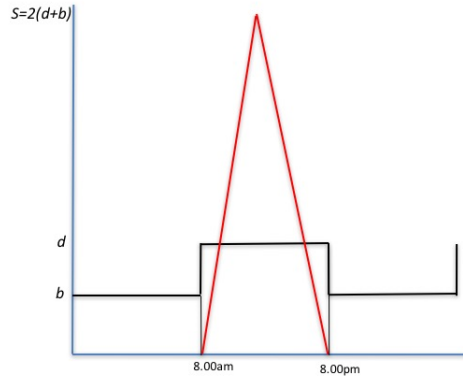
Figure 3.1:

to rely entirely on solar power, we would need storage capacity in the range of 30% to 50% of total daily consumption in order to store enough power to meet demand all around the clock.[3]

4 Smoothing Production

Dealing with the intrinsic stochasticity of renewables, wind especially, is more complex. There are two approaches to smoothing stochastic production. One is to use multiple intermittent power sources in the hope that they are stochastically independent, or indeed have negative covariances, so that diversity of energy sources reduces the risk of shortage. The second is to develop infrastructure that can store energy so that energy output can be held in reserve for use when there is a drop in production. We model the diversification approach first. In practice both will certainly be used, and indeed are being used already.

4.1 Multiple Stochastic Sources

We compare the alternatives of producing a given level of output from one large wind turbine vs many smaller ones with the same mean output in total, and investigate how the covariances between the turbines affect the variability of total output. Consider I identical wind turbines with outputs w_i each having a mean output μ and variance σ^2. The mean of the sum of their outputs is $I\mu$ and the variance of the sum is $VS = I\sigma^2 + 2\sum\sum_{i<j} Cov\left(w_i w_j\right)$ where $Cov\left(w_i w_j\right)$ is the covariance between the outputs of turbines i and j. The alternative is

---

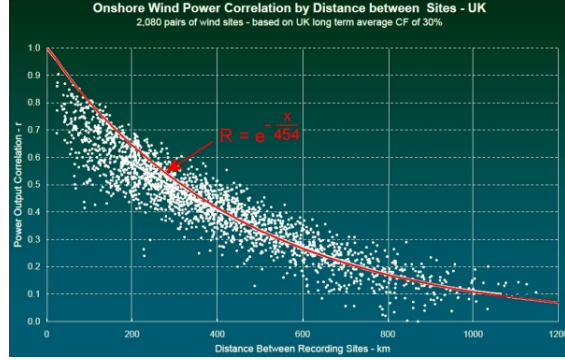[3]This is abstracting from any seasonal fluctuations in solar output, for example between winter and summer.

Figure 4.1:

a single large turbine I times larger than the smaller turbines with output w, mean output $I\mu$ and variance of output $I^2\sigma^2$.

For each of these we look at the coefficient of variation (standard deviation over mean), which for the single large turbine is σ/μ. For the sum of the smaller turbines it is $\sqrt{I\sigma^2 + 2\sum_{i<j}\sum Cov\left(w_i w_j\right)}/I\mu$. If these turbines are independent (covariances zero) then this reduces to $\sigma/\mu\sqrt{I}$. In the case of independence the ratio of the coefficients of variation of the many turbines to the one is $1/\sqrt{I}$. So multiple independent sources producing a given average total output lead to less variation of output than a single source with the same average output and the same distribution. If the covariances of the small turbines are negative then this reduces the variance of the total output even more, and it is important to note that the coefficient of variation may be reduced even if the covariances are positive. This is the statistical theory behind the idea that diversifying intermittent sources of electric power can reduce the risk of power shortage.

What in fact are the covariances between wind turbine outputs? Over short to medium distance at least, they are positive (Schmallensee (2014)). Figure 2 shows correlation as a function of distance for the UK (from PelaflowConsulting (2016, http://www.wind-power-program.com/aboutus.htm)). This shows covariances as positive up to about 1000 km.

While there seems to be no comparable data set for the US, there are reasons to think that correlations would fade to zero faster than in the UK as the topography is more varied: the UK generally is affected by a single weather system whereas the US may be subject to several different ones. However Schmallensee's work (Schmallensee (2014)) shows that up to the level of an ISO, there are positive correlations between wind sites in the US. It seems possible that for the UK once the distances between wind sites are of the order of 500km or more, the covariances are low enough that the coefficient of variation of power output is reduced by geographic diversification.

Bradbury (2013) carries out an interesting exercise for the US. Using data on the location of wind farms in the continental US in 2012, and using UK data
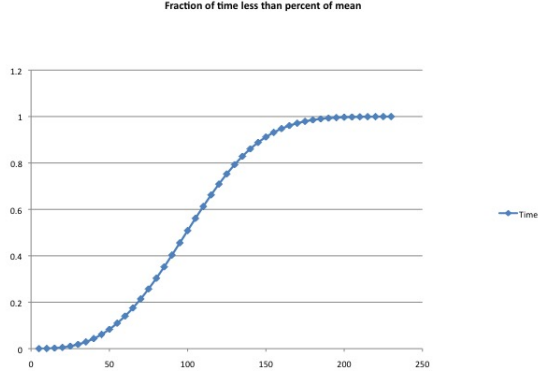
Figure 4.2:

on the decay of correlation with distance, he computes the distribution of the total output of all US wind farms. He models this as a Weibull distribution, with cumulative distribution

$$C\left(u\right) = 1 - e^{\left[\frac{u}{U_m}\Gamma\left(1+\frac{1}{k}\right)\right]^k}$$

where $C\left(u\right)$ is the proportion of the time the wind speed is between 0 and u, U_m is the mean wind speed, k is a shape parameter taken to be 2, and $\Gamma\left(1 + \frac{1}{k}\right)$ is the gamma function defined by $\Gamma\left(z\right) = \int_0^\infty e^{z-1}e^t dt$. Bradbury shows that the aggregate output has a lower coefficient of variation than the outputs of individual farms. From his estimated distribution we can compute figure 4.2.

The horizontal axis here shows the total output of all US wind farms as a percentage of their mean output, from zero to 250%. The vertical axis shows the fraction of time the output is less than the percent of the mean shown horizontally. So we see that output is less than 100% of the mean 50% of the time: it is less than 50% of the mean about 8% of the time, and less than 5% of the mean about 0.1% of the time. For 95% of the time, output exceeds 42% of the mean. So we can be 95% sure that aggregate US wind output will not fall below 42% of the mean output.

This observation gives us a way of dealing with intermittency without storage: suppose we want to be 95% (or 99%) certain of meeting demand: then (in the 95% case, and given the above distribution of aggregate wind output for the US) we need to install sufficient capacity that demand is 42% of mean output. Note that this is assuming that electric power generators anywhere in the continental US are perfect substitutes, as we are taking our numbers from the distribution of aggregate output: this implies that we have a "perfect" grid in the sense that surplus power at any point can be moved to meet any source of demand. This clearly requires the installation of capacity in excess of what
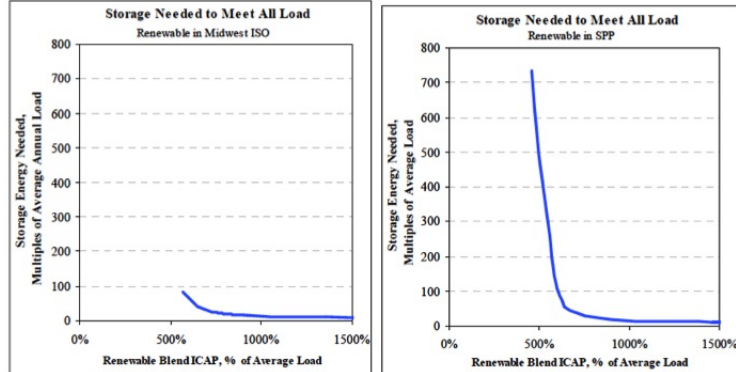
7

Figure 4.3:

will normally be used: in effect we are substituting spare production capacity for storage capacity. To give a sense of the scale of capacity required, suppose we need to meet a 1 gigawatt demand with 95% certainty using the distribution in figure 4.2. We would need capacity with a mean output of 2.381 gigawatts (1 gig is 42% of this), and assuming capacity factor of 40% this means a nameplate capacity of 5.95 gigs - almost 6 times the demand we want to be 95% sure of meeting. When the wind is blowing it could lead to extensive curtailment of wind power (on which see Denholm and Margolis (2016) in the case of California), but this could still be less expensive than building storage capacity.

Another very interesting paper is Chang, Spees and Weiss (Chang et al. (2010)). They take load data from the Midwest ISO (MISO) and the Southern Power Pool (SPP), assume that this load is met from wind and solar capacity just sufficient to meet demand in total over the year (the integral of output equals the integral of consumption), and take time paths of output generated by simulating hourly wind and solar generation profiles for the two regions from NREL. They then simulate these two systems with various assumptions about storage capacity. What emerges, as suggested above, is that storage capacity and generating capacity are substitutes for each other: installing more generating capacity lowers the need for storage.

Their results are illustrated by the graphs in figure 4.3. The horizontal axis shows the amount of capacity installed as a % of average load, and the vertical the amount of storage needed to meet demand, measured in hours of average load. Above we spoke of installing 600% of average load if we want to be 95% certain of meeting this without storage: Chang et al suggest that even with 600% of average load installed, the MISO would need of the order of 70 hours

of storage capacity, and the SPP about 50 hours. (These solutions would lead to extensive curtailment of renewables under certain conditions.) It's important to recognize that Chang et al are meeting demand in each ISO from wind and solar generated only in that ISO, so they are not pooling risks across the entire country, as is in effect done by Bradbury and by the calculations based on his work. They are also pooling the two types of storage distinguished above, time shifting and stochastic smoothing. Their graphics illustrate very nicely the capacity-storage tradeoff, to which I will return below.

4.2   Storing Stochastic Power Output

Next we address directly the issue of using storage capacity to smooth the output of a stochastic energy source such as wind. Suppose the stochastic source produces an output R_t at time t, with demand being deterministic and given by C_t. Energy can be stored and the stock available at time t is S_t. Clearly the stock evolves according to the equation

$$\frac{dS_t}{dt} = R_t - C_t \tag{4.1}$$

where R_t follows a stochastic process. In principle we can integrate this equation using Ito's lemma and stochastic calculus. I will simplify matters by replacing the differential equation (4.1) by a difference equation

$$S_{t+1} = S_t + R_t - C_t \tag{4.2}$$

I further simplify this by assuming that $C_t = C \; \forall t$ and that R_t is generated by a very simple stochastic process: each period the energy output is $R > 0$ with probability p and zero with probability $1 - p$. Expected output is thus pR and variance $R^2 p (1 - p)$. Each successful generation episode adds the amount $R_t - C$ to the energy stock, and each unsuccessful one adds $-C$. After t periods (trials) the probability that there have been k successes and $t - k$ failures in generation is given by the binomial distribution

$$f\left(k : p, t\right) = \left(\begin{array}{c} t \\ k \end{array}\right) p^k \left(1 - p\right)^{t-k} = \frac{t!}{k! \left(t - k\right)!} p^k \left(1 - p\right)^{t-k} \tag{4.3}$$

If there have been k successful generation episodes in t periods then the stock is $kR - tC + S_0$ and so

$$Pr\left\{S_{t+1} = kR - tC + S_0\right\} = \frac{t!}{k! \left(t - k\right)!} p^k \left(1 - p\right)^{t-k} \tag{4.4}$$

It follows that

$$Pr\left\{S_{t+1} \geq 0\right\} = \sum_{x=0}^{x=\infty} \frac{t!}{\left(\frac{tC - S_0}{R} + x\right)! \left(t - \frac{(tC + S_0)}{R} - x\right)!} p^{\frac{tC + S_0}{R} + x} \left(1 - p\right)^{t - \frac{(tC - + S_0)}{R} - x} \tag{4.5}$$

9

$$= \sum_{k=\frac{tC-S_0}{R}}^{\infty} \frac{t!}{k!\,(t-k)!} p^k \, (1-p)^{t-k}$$

We are interested in understanding the factors that influence this probability, and in particular those amenable to influence by policy choices. These are S_0, the initial stock of energy, and R, the level of generation capacity installed. The probability of a non-negative stock increases with both of these factors, confirming analytically the finding of Chang et al that more installed capacity lowers the need for storage.

Storage capacity has not been explicitly modeled here: it has been taken to be unlimited. Its effect would be to put an upper limit on the stock of energy, and also a lower limit as stored energy cannot be negative, and so change the dynamics of stock movement over time from that given by equation (4.2). Assume that the maximum amount of energy that can be stored is $D > 0$: in this case the stock dynamics are

$$\begin{aligned} S_{t+1} &= D \; if \; S_t + R_t - C \geq D \\ S_{t+1} &= 0 \; if \; S_t + R_t - C \leq 0 \\ S_{t+1} &= S_t + R_t - C \; otherwise \end{aligned} \qquad (4.6)$$

This changes the overall dynamics. Consider a situation where $S_t = D$: then if the next four energy outcomes are $R, R, 0, 0$ this lead to a drop in stocks of $2C$, whereas if the next four are $0, 0, R, R$ then there is a drop of $2C$ followed by an increase of $2(R - C)$, unless this leads to a stock in excess of D. So now the order in which outcomes happen, matters for the movement of the stock and the difference equations need to keep track not only of the total number of generations episodes but also of their sequence.

5   Conclusions

This paper has explored some basic issues in the economics of energy storage. There are two different functions that storage has to perform: one is to shift solar power produced in the daytime to the night, assuming that there are not sufficient other sources of power available at night. The second is to smooth out fluctuations in the output of renewable energy.

It might be possible to design a system in which either or both of these functions are redundant. There could be sufficient carbon-free generating sources to meet demand out of daytime hours - nuclear, geothermal, hydro and others - in which case no time-shifting of solar power would be needed.

These same sources could also be used to smooth the output of a stochastic renewable source. But it could also be possible to avoid the need for storage to smooth output by spatial diversification of renewable energy sources, so lowering the correlation between different sources, and by building large amounts of capacity. And in both cases the need for storage could be further reduced by demand management, giving consumers incentives to reduce demand on the

grid at times of shortage and to shift loads to times of surplus.[4] Consumers have many ways of storing energy at their disposal - they can store hot or cold water in tanks, make ice when power is available and use it to cool air when it is not, and store energy in batteries, in particular in the batteries of electric vehicles. Denholm and Margolis (2016) consider this last option in detail.

The bottom line is that the question we are focussing on - how much energy storage would be needed if an economy such as that of the US were to move to much heavier dependence on renewable energy, as is implied by the goal of substantial reduction of greenhouse gas emissions - is probably not well-posed. As de Sisternes et al. (2016) note, "In general, while energy storage appears essential to enable decarbonization strategies dependent on very high shares of wind and solar energy, storage is not a requisite if a diverse mix of flexible, low-carbon power sources is employed, including flexible nuclear power." There are some routes to low GHG emissions that travel via extensive use of storage, and others that make little if any use of these technologies and use other ways of managing intermittent power supplies. Which is best seems to be a matter of costs. If storage costs continue to fall, storage will feature prominently in the ultimate solution: otherwise we will work with a range of alternatives.

In my earlier paper Heal (2016) I assumed that the ability to store about two days of energy consumption would be needed in a world with 66% renewable energy, divided equally between wind and solar PV. Where does this discussion leave that assumption? It is convenient to discuss solar and wind separately. One third of total consumption would come from solar PV, and several studies (for example Denholm and Margolis (2016)) suggest that this level of solar penetration could be accommodated without storage: that it could all be used during the day. If we were to seek to shift some of this to other times of day, we would probably want to shift less than one half of solar output, which is less than one sixth of daily consumption.

The output from the one third of capacity that is wind could be smoothed by spatial diversification, as analyzed in section 4.1 above, but not completely, and at the cost of improving the connectivity of the grid (in my earlier paper I included the cost of an increase in miles of high voltage lines by 25%). We could adopt the approach of overbuilding wind capacity to reduce the probability of a power shortage to some acceptable level, combining this with demand management to cope with the low-probability eventualities. The illustrative calculations in section 4.1 suggest that in this case we might need to install six time the demand we want to meet, which would mean doubling the amount of wind capacity in my earlier paper. This would increase cost by $0.78 trillion. As I noted in my earlier paper, a battery large enough to store one day's output from a wind turbine would cost twice as much as the turbine itself. It follows the extra capacity, with possible curtailment in the event of strong and persistent winds, is probably less costly than storage. In the earlier paper I allowed between $2.2 and $5.1 trillion for storage capacity (table 2), and for less than

---

[4]Denholm and Margolis (2016) review the literature on the potential for shifting demand by demand management.

the lower limit here we could both build extra wind capacity for smoothing and construct some storage. The bottom line is that the storage figures I used in Heal (2016) are probably too high: less than one day of storage capacity might be adequate.

References

Bradbury, L.J.S., "The probability density distribution for the power output from arrays of wind turbines and the intermittent nature of wind power," *Journal of Wind Engineering and Industrial Aerodynamics*, October 2013, *123*, 121–129.

Chang, Judy, Kathleen Spees, and Jurgen Weiss, "Renewables and Storage - Does Size Matter?," Technical Report, Brattle Consulting Group March 2010.

de Sisternes, Fernando, Jesse Jenkins, and Audun Botterud, "The value of energy storage in decarbonizing the electricity sector," *Applied Energy*, 2016, *175*, 368–379.

Denholm, Paul and Robert Margolis, "Energy Storage Requirements for Achieving 50Technical Report NREL/TP-6A20-66595, National Renewable Energy Laboratory August 2016.

Gowrisankaran, Gautam, Stanley Reynolds, and Mario Samano, "Intermittency and the Value of Renewable Energy," Technical Report, University of Arizona May 2015.

Heal, Geoffrey, "The Economics of Renewable Energy," *Review of Environmental Economics and Policy*, December 2010, *4* (1), 139–154.

_ , "What would it take to reduce US greenhouse gas emissions 80 percent by 2050?," Technical Report 22525, NBER August 2016.

PelaflowConsulting, "About the WindPower project," 2016, http://www.wind-power-program.com/aboutus.htm.

Schmallensee, Richard, "The Performance of US Wind and Solar Generating Plants," August 2014.