The Cost of Curtailing Wind Turbines for Secondary Frequency Regulation Capacity

Stephen Rose and Jay Apt

S. Rose is with the Department of Engineering & Public Policy, Carnegie Mellon University, Pittsburgh PA 15213, USA (email: srose@cmu.edu)

J. Apt is with the Tepper School of Business and the Department of Engineering & Public Policy, Carnegie Mellon University, Pittsburgh PA 15213, USA (email: apt@cmu.edu)

List of Symbols

$R_{max}(k)$	Max available upregulation capacity in k^{th} interval (MW-h)
$R(k \mid \Delta)$	Upregulation capacity in k^{th} interval, given curtailment of Δ (MW-h)
Δ	Curtailment of wind farm power output below the possible power output (MW)
LOL	Lower operating limit: minimum curtailed power (p.u.)
$P_{poss}(t)$	Instantaneous possible power output of wind farm (MW)
$P_{curt}(k \mid \Delta)$	Instantaneous curtailed power output of wind farm, given curtailment of Δ (MW)
$P_{\mu}(\mathrm{k})$	Mean wind farm power in the k^{th} interval (MW)
$P_{COV}(\mathrm{k})$	Coefficient of variation of wind farm power in the k^{th} interval (MW)
$E_{loss}(k \mid \Delta)$	Energy production lost to curtailment in the k^{th} interval (MWh)
t	Time (sec)
T	Duration of frequency-regulation dispatch intervals (sec)
k	Index of dispatch intervals
$AC(k \mid \Delta)$	Average (opportunity) cost of up-regulation capacity [MWh/MW-h]
$MC(k \mid \Delta)$	Marginal (opportunity) cost of up-regulation capacity [MWh/MW-h]

1. Introduction

When large numbers of wind turbines are connected to the electrical grid, the rapid variability of wind power on short time scales can cause the grid frequency to deviate significantly from its nominal value [1, 2]. In the many regions where pumped hydroelectric storage is not available, the grid frequency is typically regulated by adjusting the power output of fast-ramping thermal generators, including gas turbines. Gas turbines can be expensive to operate and produce power less efficiently and with higher levels of NO_x emissions when quickly varying their power output [3]. Other technologies are capable of compensating for the short-term variability of wind power, but they are either currently too expensive to be practical or cannot be scaled large enough to meet the rapidly-increasing penetration of wind power on the electrical power grid. Batteries and flywheels are technically well-suited to rapidly generating or absorbing power but commercially available units are too expensive at the moment on the scale needed for the penetrations of wind power expected in the next 10 – 20 years, although some promising systems are in the development phase. Hydroelectric power, and especially pumped hydro storage, is technically well-suited to rapidly changing power output and is relatively inexpensive. However, little new pumped hydro storage is being developed in the United States.

Denmark, Ireland, Great Britain, and Germany now include requirements in their national grid codes that wind farms be able to increase or decrease their power output to aid in regulating the grid frequency [4-7]. The active power output of a wind farm can be decreased by reducing the aerodynamic efficiency of the wind turbines or completely shutting down some turbines, but it is impossible to increase the output of a wind farm beyond the power level provided by the current wind velocity. Operating a wind farm at less than the currently available wind capacity ("curtailing") creates a reserve of power that allows power output to be increased on demand. Prior research has demonstrated the technical feasibility of curtailing the power output of a wind farm to regulate the grid frequency [5, 8-10]. Those authors consider primary frequency regulation, in which a generator responds to frequency deviations in a few seconds. In this paper we consider *secondary* frequency regulation, in which the generator responds over tens of seconds or minutes to a dispatch signal from the grid operator. Other research has demonstrated the feasibility of curtailing a wind farm to reduce the variation in power output or limit power ramp rates [11-13]. Many modern wind farms already have most or all the equipment necessary to curtail for frequency regulation and many others can be retrofitted inexpensively. However, the revenue a wind farm foregoes due to curtailment may be greater than the cost of procuring the same frequency regulation from traditional sources such as gas turbines or hydroelectric power plants, or using an energy storage technology such as a battery to store a reserve of energy [9, 10]. Kirby et al. analyzed historical energy and regulation prices in Texas to show that wind would infrequently be able to supply regulation at competitive prices and uses a simple wind farm simulation to demonstrate the technical feasibility [14].

We calculate the cost and quantity available of frequency upregulation capacity from a curtailed wind farm as a function of wind conditions. Upregulation capacity is the ability of a generator to increase its active power output to increase the grid frequency. We also compare the cost of upregulation from a curtailed wind farm to market prices for upregulation capacity (typically set by gas turbines), though grid operation rules do not generally allow wind power to bid into the market for upregulation capacity (or other ancillary services). This comparison gives a first-order estimate of how often grid operators may call on wind turbines for secondary frequency regulation.

The method of curtailing a wind farm we consider in this research does not reduce the variability of the wind farm power output. We control the wind farm to maintain a constant difference Δ between the possible power P_{poss} and curtailed power P_{curt} outputs, which provides only a reserve of power that the grid operator can use to regulate the supply of power to match demand. There is an alternative curtailment scheme that would reduce the variability of wind power by curtailing the wind farm output to a low fixed amount, but we do not consider it here because the cost would be extremely high and there is still active debate about how much cost wind power variability imposes on a grid operator [15, 16].

2. Model

We calculate the amount of upregulation capacity a curtailed wind farm can produce and its cost by simulating the operation of a 100-MW wind farm. The wind farm produces upregulation capacity by curtailing its power output a fixed amount below the power possible in given wind conditions; this creates a reserve of power than can be dispatched on demand. The wind farm we model consists of twenty 5-MW pitch-regulated turbines that receive power setpoint commends from a closed-loop wind farm controller that regulates the aggregate power output of the wind farm. The turbines are driven by wind speed data that is a hybrid of measured low-frequency wind speed data and simulated high-frequency turbulence. We analyze the curtailed power output of the wind farm, relative to its uncurtailed output, to calculate how much upregulation capacity it can produce and at what cost.

We create 60 days of wind speed data sampled at 5 Hz for three locations in central North America listed in Table 1. We use the method developed by Rose and Apt to create the wind speed data as a hybrid of measured 10-minute data and simulated high-frequency turbulence [17]. This method simulates high-frequency turbulence to “interpolate” between empirical data while maintaining the statistical properties of the empirical data. In this research, the empirical data is 10-minute mean and standard deviation of the wind speed measured at 50-meter height at three locations listed in Table 1. We randomly draw 60 days of data from each year and each location; 15 days from each season where possible. The high-frequency turbulence is simulated using Veers’ method with the Kaimal spectrum recommended by the IEC wind turbine design standard [18, 19]. We create wind-speed time series for each turbine location in the wind farm. The wind speed time series for each turbine location have identical mean wind speeds, but the turbulences are related to each other by the lateral coherence relation proposed by Sørensen et al. [20]. We use only the lateral coherence relation because all the turbines are arranged in a straight line perpendicular to the wind.

We use the hybrid wind speed data to dynamically simulate the power output of each turbine in a 20-turbine wind farm. Dynamically simulating individual turbines is an improvement on steady-state turbine models that relate power output to wind speed with a simple power curve, or aggregate farm models that lump the entire farm into a single equivalent turbine [20-22]. Each turbine is a pitch-regulated 5-MW turbine with a 126-meter rotor designed by the National Renewable Energy Laboratory [23]. The turbine rotor is large enough that it has a smoothing effect on the wind turbulence at higher frequencies; we model this smoothing effect with the wind turbine admittance function $F_{wt}(f)$ proposed by Sørensen et al. [20]. The turbines in the wind farm are spaced 5 rotor diameters apart (630 m) in a line perpendicular to the wind direction; the wind direction is constant.

The wind farm simulations are run using the SimWindFarm toolbox (version 0.8) for Matlab, developed as part of the Aeolus project [24]. The active power output of the wind farm is curtailed by a closed-loop controller that reduces the power setpoint for each turbine so the aggregate actual power output is a fixed number of megawatts below the possible power output; the Danish grid code refers to this control scheme as “Delta production constraint”, where Δ is the fixed difference between the possible and actual power outputs [25]. Individual turbines are not curtailed equally—each turbine is curtailed proportional to its available power [26]. We do not allow the wind farm controller to command any wind turbine to curtail below its lower operating limit (LOL) of 20% of its rated power, though a turbine’s power output may go below that limit when the wind speed is low.

We simulate the power output of the wind farm for levels of curtailment $0 \leq \Delta \leq 30$ MW with identical wind inputs; the power output is sampled at 1 Hz. To limit the number of simulations, we increase curtailment in one-megawatt steps up to 10 MW, two-megawatt steps up to 20 MW, and five-megawatt steps up to 30 MW. We repeat the simulations of all curtailment levels with 60 days of hybrid wind speed data described above. To generalize the results of these experiments, we repeat the simulations with five wind speed data sets of 60 days each, listed in Table 1: three sets of wind speed measurements from a wind farm site in west Texas in 2007, 2008, and 2009, a set of measurements from a site in the northern Great Plains of the U.S. in 2008, and a set of measurements from a site in Ontario, Canada in 2008.

Table 1: Empirical wind speed data used to simulate wind farm power

Location	**Period**	**Mean wind speed**	**Wind turbulence intensity (TI)**	**Wind farm capacity factor**
West Texas	29 Mar. – 8 Dec. 2007	6.8 m/s	13%	29%
	21 Dec. 2007 – 14 Dec. 2008	7.2 m/s	13%	33%
	25 Dec. 2008 – 1 Sept. 2009	7.5 m/s	13%	35%
Northern Great Plains (U.S.)	2 Jan. – 17 Dec. 2008	8.0 m/s	10%	41%
Ontario (Canada)	29 Dec. 2007 – 20 Dec. 2008	6.7 m/s	12%	28%

3. Analysis

We calculate the amount of upregulation capacity that a curtailed wind farm can provide by retrospectively analyzing the wind farm simulations described in Section 2. The upregulation capacity $R(k \mid \Delta)$ that a wind farm can supply for given curtailment Δ in the k^{th} dispatch interval of length T is the smallest difference between the uncurtailed ("possible") power at time $P_{poss}(t)$ and curtailed power at time $P_{curt}(t)$:

$$R(k \mid \Delta) = \min\left(P_{poss}(t) - P_{curt}(t \mid \Delta)\right) \quad (1)$$

for t in the k^{th} interval: $kT \le t \le (k+1)T$.

An example of data for this calculation are shown in Fig. 1 for a dispatch interval of T = 300 seconds and a curtailment of Δ = 5 MW. During most of the dispatch interval, the curtailed power $P_{curt}(t \mid \Delta = 5)$ (green line) closely tracks the desired curtailment $P_{ref}(t) - \Delta$ (dashed line), as shown at point A. However, $P_{curt}(t)$ (green line) is sometimes not curtailed as much as desired, as shown at point B. In this case illustrated by point B, the power output of several of the turbines in the wind farm reached the lower operating limit (LOL = 0.2 p.u.) and could not curtail further. The amount of available upregulation capacity $R(\text{k} \mid \Delta = 5)$ for the interval, shown as the shaded area, is limited by the smallest difference between $P_{poss}(t)$ and $P_{curt}(t \mid \Delta = 5)$ in that interval. Upregulation capacity R is always less than or equal to the curtailment Δ. We measure upregulation capacity in megawatts hours (MW-h), which is different from the unit of energy "megawatt hours (MWh)".

We analyze dispatch intervals T of 60 minutes and 15 minutes. Most power systems with markets for frequency regulation use a dispatch interval of 60 minutes, which means market participants bid an amount of regulation capacity they can sustain for the full 60 minutes. We test a dispatch interval of 15 minutes to determine whether a curtailed wind farm can better compete over shorter intervals.

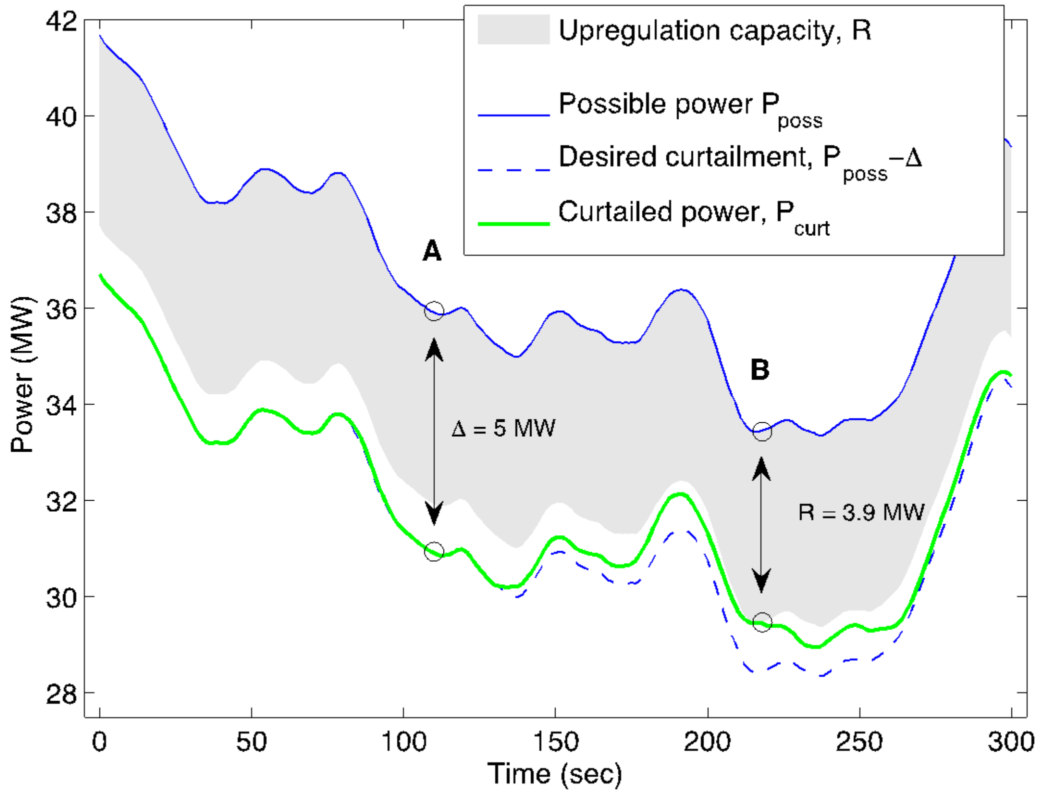

Fig. 1 Method for calculating the available regulation capacity. In this example, the 100 MW-capacity wind farm is curtailed by Δ = 5 MW (point A), but the available regulation capacity $R(k \mid \Delta)$ is 3.9 MW (point B), which is the largest curtailment that can be maintained through the entire dispatch interval T, i.e. the smallest difference between the uncurtailed and curtailed power outputs of the wind farm

The upregulation capacity available in a given dispatch interval, calculated in (1) cannot be known in advance without perfect forecasting of future wind conditions. Perfect foresight is an unrealistic assumption, but it sets an upper bound on the amount of regulation capacity available and the opportunity cost, in terms of energy production lost, to produce it. The results in Section 4.1 relax this assumption slightly—there we assume perfect forecasting of only the mean and standard deviation of wind speed in a given dispatch interval, rather than perfect forecasting of the wind speed at every moment.

We calculate the cost of upregulation capacity from a curtailed wind farm by retrospectively analyzing the wind farm simulations described in Section 2. The average cost of upregulation capacity $AC(k \mid \Delta)$ in the k^{th} dispatch interval with curtailment Δ is the energy generation lost in that interval due to curtailment $E_{loss}(k \mid \Delta)$ divided by the quantity of upregulation capacity provided during that interval $R(k \mid \Delta)$:

$$AC(k \mid \Delta) = \frac{E_{loss}(k \mid \Delta)}{R(k \mid \Delta)} \qquad (2)$$

where E_{loss} is measured in MWh and AC is measured in MW-h/MWh. We calculate the marginal cost of upregulation capacity $MC(k \mid \Delta)$ as the additional upregulation capacity divided by the additional energy loss resulting from curtailing one more step. The steps are not always one megawatt; to limit the number of simulations, we increase curtailment in one-megawatt steps up to 10 MW, two-megawatt steps up to 20 MW, and five-megawatt steps up to 30 MW. The average power was often not high enough to simulate curtailment up to 30 MW-- only approximately 20% of the 1-hour intervals we examined produced enough power to curtail by 30 MW.

4. Results

We present three results: estimates of the maximum regulation capacity available in given conditions, the cost of regulation capacity in terms of unproduced energy, and the cost premium for curtailment-derived regulation capacity compared to the market price.

4.1 Maximum Available Regulation Capacity

The maximum upregulation capacity available from curtailing a wind farm in the k^{th} interval $R_{max}(k)$ is a nearly-perfect linear function of the minimum power wind farm power sampled at 1 Hz in that interval based on our definition of upregulation capacity R in (1). However, the minimum power in a interval is difficult to predict, so instead we assume it can be estimated from the mean $P_\mu(k)$ and standard deviation $P_\sigma(k)$ of wind farm power and we model $R_{max}(k)$ with a linear function of $P_\mu(k)$ and $P_\sigma(k)$ given in (3).

$$\hat{R}_{max}(k \mid \tau) = a_0 + a_1 P_\mu(k) + a_2 P_\sigma(k) \quad (3)$$

The model in (3) is fit to simulated wind power data using quantile regression, which determines the plane below which of τ% of the R_{max} values lie for a given quantile $0 \le \tau \le 1$ [27]. A separate model is fitted for each quantile τ, so the values of parameters a_0, a_1, and a_2 shown in Table 2 are functions of τ. For example, we estimate a 95% probability that $R_{max}(k) < -18.4 + 0.926P_\mu(k) - 0.940P_\sigma(k)$ for a 15-minute dispatch interval. We do not include an interaction term $P_\mu * P_\sigma$ in the model given in (3) because the coefficient is not statistically significant for lower quantiles (0.05 and 0.25) for 15-minute intervals and most of the fitted quantiles (0.05, 0.5, 0.75, and 0.95) for 60-minute intervals. We use quantile regression to fit the model instead of ordinary least squares (OLS) because our data are not well-suited to OLS regression-- the residuals have longer tails than a normal distribution and are heteroskedastic, which result in poor estimates of the confidence intervals of the fitted parameters. For comparison, Table 2 lists parameters for two models fitted with OLS regression: the first model identical to (3) and the second model adds an interaction term $P_\mu * P_\sigma$. The quantile regression model is fitted using the "quantreg" package (version 4.97) [28] in the R statistical software and the OLS regression is fitted using the "lm" function of R (version 3.0.0) [29].

Table 2: Regression coefficients for max regulation capacity R_{max} in (3) fitted using quantile regression. For comparison, coefficients fitted with OLS regression for the same model and a model with a statistically-significant interaction term ($P_\mu * P_\sigma$.) are shown in the two right-most columns labeled "OLS". Standard errors calculated with bootstrapping methods are shown in parenthesis. All coefficients are significant at a 99% level.

Statistic	Quantile					OLS	
	0.05	0.25	0.5	0.75	0.95		
15-minute intervals (n = 1964)							
Constant (a_0)	-17.8	-18.2	-18.3	-18.3	-18.4	-17.9	-18.9
	(0.22)	(0.16)	(0.13)	(0.12)	(0.15)	(0.14)	(0.21)
P_μ coeff. (a_1)	0.870	0.893	0.904	0.912	0.926	0.890	0.912
	(0.0094)	(0.0064)	(0.0055)	(0.0039)	(0.0055)	(0.0042)	(0.0055)
P_σ coeff. (a_2)	-1.62	-1.39	-1.21	-1.07	-0.940	-1.22	-0.936
	(0.054)	(0.037)	(0.032)	(0.017)	(0.033)	(0.016)	(0.0470)
P_μ* P_σ coeff.							0.00583
							(0.00092)
60-minute intervals (n = 422)							
Constant (a_0)	-14.5	-15.9	-16.1	-16.2	-16.9	-15.0	-19.0
	(1.4)	(0.62)	(0.51)	(0.53)	(0.70)	(0.47)	(0.83)
P_μ coeff. (a_1)	0.677	0.747	0.774	0.791	0.835	0.743	0.823
	(0.059)	(0.023)	(0.016)	(0.024)	(0.02)	(0.014)	(0.019)
P_σ coeff. (a_2)	-1.23	-1.13	-1.03	-0.966	-0.882	-1.03	-0.477
	(0.15)	(0.09471)	(0.04554)	(0.08668)	(0.04690)	(0.042)	(0.10)
P_μ* P_σ coeff.							-0.0100
							(0.0017)

The a_1 coefficient in Table 2 shows that R_{max} increases 0.68 - 0.83 MW when the mean power in a 60-minute dispatch interval increases by 1 MW and by 0.87 - 0.93 MW for a 15-minute dispatch interval. The a_2 coefficient shows that the expected value of R_{max} decreases as the standard deviation of the wind farm power increases. For example, if the standard deviation increases by 1 MW, R_{max} decreases by 1.23 - 0.88 MW for a 60-minute interval and 1.62 - 0.94 MW for a 15-minute interval.

These results assume perfect forecasting of the mean and standard deviation of wind farm power for a given interval. This is important because upregulation capacity must typically be bid into the market hours or a full day ahead. If future wind conditions cannot be forecast with perfect accuracy, the wind farm must bid less upregulation capacity into the market than it could theoretically produce or risk being unable to meet its commitment. However, we show in Section 4.2 (below) that a wind farm should bid less than the R_{max} because the costs rise steeply as the upregulation capacity approaches its maximum.

The model parameters listed in Table 2 are fit to measurements derived from the power output of the 100-MW wind farm described in Section 2 simulated with wind speed data from west Texas in 2008 described in Table 1. The dependent variable, maximum upregulation capacity in the k^{th} interval $R_{max}(k)$, is calculated as the maximum of (1) in that interval over the range of simulated curtailments $\Delta \in [0\ 30]$ MW. The independent variables P_μ and P_σ are the mean and standard deviation of wind farm power sampled at 1 Hz in the given interval. We exclude intervals with minimum power less the LOL (20 MW) because the $R(k)$ is zero for those intervals and exclude intervals with minimum power greater the sum of the LOL and the largest simulated curtailment (20+30 = 50 MW) because $R(k)$ is constant for those intervals. For the 60 days of west Texas 2008 wind speed data, these exclusions yield 1964 15-minute intervals or 422 60-minute intervals. We present a table of summary statistics for these “training” data in the online supporting information.

We validate the fitted models in Table 2 with two tests described by Vaz et al.: a correct classification test and a rank correlation test [30]. For both tests, we validate our model, which is fitted to west Texas 2008 data, against three of the data sets described in Table 1: west Texas 2007, northern Great Plains 2008, and Ontario 2008. Tables of summary statistics for these “validation” data are presented in the online supporting information.

The correct classification test computes the fraction of observed $R_{\max}$ values that are less than the $\hat{R}_{\max}$ values predicted by the model in (3) for a given quantile τ, a quantity that Vaz et al. call the “correct classification statistic” (CCS) [30]. For example, a successful model of the 75th percentile ($\tau = 0.75$) would predict $\hat{R}_{\max}$ values that are greater than approximately 75% of the $R_{\max}$ values in a validation data set. We consider a model for a given quantile τ successful if the 95% confidence interval around the CCS contains τ. The 95% confidence intervals (CI) of the CCS are calculated from 1000 bootstrap resamples of the validation data sets (with replacement). The calculated CCS and their 95% CI for all quantiles, validation data sets, and interval lengths are given in the online supporting information.

The models in Table 2 are successful for the west Texas 2007 data: the desired quantiles τ are within the CCS values for all quantiles and time intervals except 5th percentile for 15-min intervals. The models have mixed success with the Ontario 2008 data: the desired quantiles are within the CI for most quantiles in 60-minute intervals but none of the higher quantiles in 15-minute intervals. The models have poor success with the northern Great Plains 2008 data: the desired quantiles are consistently higher than the CI for the CCS values, which means our models predict lower maximum available regulation capacity than is observed. For example, only 84.5 – 88.0% of the observed $R_{\max}$ values are less than the predicted 95th percentile value in the northern Great Plains 2008 data for 15-minute intervals.

The second validation test calculates Spearman’s rank correlation coefficient r_s, which estimates the correlation between the *rank* of the predicted $\hat{R}_{\max}$ values and the observed $R_{\max}$ values. The coefficient r_s is calculated by separately putting the $\hat{R}_{\max}$ and $R_{\max}$ values in rank order and calculating Peasrson’s correlation coefficient between the ranks. An r_s value of 1 indicates that the two variables are related by a monotonically-increasing function. We estimate the 95% confidence interval (CI) for the r_s values from 1000 bootstrap resamples of the validation data sets (with replacement). Spearman’s rank correlation coefficient is better-suited to quantile regression than Pearson’s correlation coefficient because Spearman’s does not assume a linear relationship between the two variables. We consider a model to be more successful according to this rank correlation test if value of r_s for a given validation data set is closer to 1 and statistically significant [30].

According the rank correlation test, the models in Table 2 fit the validation data well. The r_s values range from 0.925 to 0.986 and all are statistically significant. Models for 15-min intervals fit better (have higher r_s values) than models for 60-min intervals. The r_s values for all data sets and quantiles of 15-minute periods are similar to each other, and the r_s values for all data sets and quantiles of 60-minute periods are similar to each other. A table of the r_s values and their 95% CI for all validation data sets, quantiles, and time intervals is given in the online supporting information.

4.2 Cost of Curtailing for Regulation Capacity

The average and marginal costs of curtailing a wind farm for frequency regulation as a function of regulation capacity are calculated with (2) and the results are shown in Fig. 2 for 15-minute and 60-minute dispatch intervals. We calculate the average cost (AC) as the opportunity cost, i.e. the amount of energy production (in MWh) lost to produce 1 MW-h of steady regulation capacity. The marginal cost (MC) is the opportunity cost of additional regulation capacity R produced by curtailing the wind farm one additional megawatt. Fig. 2 plots cost curves for three representative ranges of R_{max}: 5 MW, 10 MW, and 15 MW. The circles denote the median cost and the error bars show the 5th and 95th percentile costs.

We find several trends in the cost of regulation capacity from a curtailed wind farm. First, the cost of regulation capacity is lower in intervals with larger maximum available regulation capacity R_{max}. Second, the cost is high for quantities of regulation capacity near zero or near R_{max}. Third, costs are slightly lower for shorter dispatch intervals, e.g. 15-minute intervals vs. 60-minute intervals. These trends are consistent for wind power simulated with wind data from the three sites listed in Table 1.

The cost of regulation capacity is lower in intervals with larger R_{max}, as shown in Fig. 2. The cost curves representing intervals with smaller R_{max} (e.g. 5 MW) have higher costs for all levels of regulation capacity then the curves representing intervals with larger R_{max} (e.g. 15 MW). For example, minimum median AC for 60-minute dispatch intervals in Fig. 2c is 1.51 MWh/MW-h when R_{max} = 5 MW, 1.11 MWh/MW-h when R_{max} = 10 MW, and 1.06 MWh/MW-h when R_{max} = 15 MW.

Similarly, the marginal cost MC decreases as R_{max} increases. For example, the minimum median MC for a 60-minute dispatch interval in Fig. 2d is 1.31 MWh/MW-h when R_{max} = 5 MW, 1.03 MWh/MW-h when R_{max} = 10 MW, and 1.00 MWh/MW-h when R_{max} = 15 MW.

The cost of regulation capacity is high for quantities of regulation capacity R near zero or near R_{max}. The high cost of regulation capacity near zero ($R \leq$ 1 MW) can be seen on the left side of the AC curves in Fig. 2a and 2c. For example, Fig. 2c shows the median AC of 1 MW of regulation capacity is 1.39 MWh/MW-h but the median AC of 2 MW is 1.18 MWh/MW-h in a 60-minute dispatch interval with R_{max} = 10 MW.

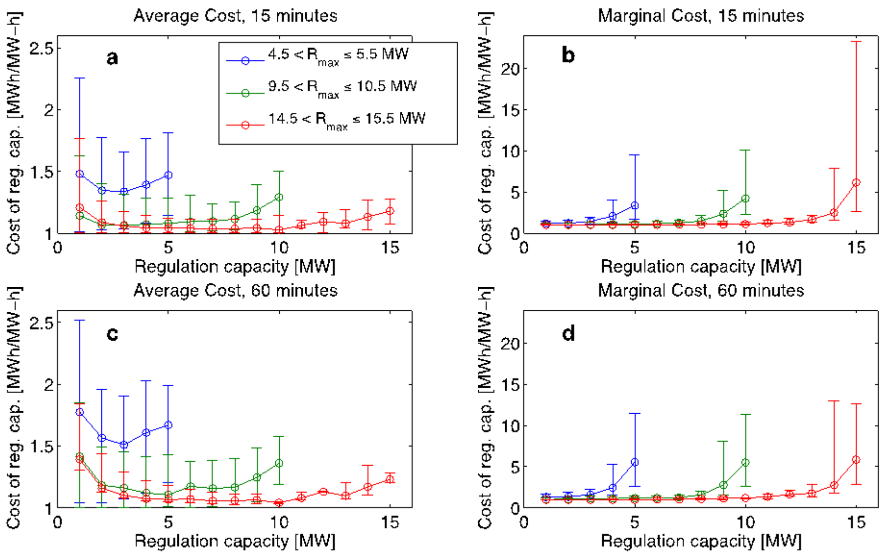

Fig. 2 The average cost (a, c) and marginal costs (b, d) of regulation capacity for 15-minute (a, b) and 60-minute (c, d) dispatch intervals. Each line represents cost data for dispatch intervals with maximum available regulation capacity in a certain interval described in the figure key. Circles denote the median and error bars show the 5th and 95th percentile costs. These results are calculated for a 100-MW wind farm simulated with wind speed data from west Texas in 2008

The opportunity cost when R is near zero can be reduced by curtailing a few turbines more deeply instead of curtailing all turbines in the wind farm equally. For example, if all turbines are curtailed equally, the median AC of R = 1 MW is 1.48 MWh/MW-h for a 15-min dispatch interval with R_{max} = 5 MW. If only half the turbines in the wind farm are curtailed, the median AC is 1.12 MWh/MW-h and if a quarter of the turbines are curtailed, the median AC is 1.00 MWh/MW-h. However, concentrating the curtailment on a few turbines reduces the maximum available regulation capacity; for example, concentrating the curtailment on 25% of the turbines reduces R_{max} by a factor of four.

The high cost of regulation capacity near the maximum can be seen on the right side of the AC curves in Fig. 2s and 2c and the MC curves in Fig. 2b and 2d. As R approaches R_{max}, the average and marginal costs increase sharply. For example, Fig. 2d shows the marginal cost of increasing R from 8 to 9 MW is 1.58 MWh/MW-h and the marginal cost of increasing R from 9 to 10 MW is 2.80 MWh/MW-h in a 60-min interval with R_{max} = 10 MW. The cost increases sharply as R approaches R_{max} because individual wind turbines reach the lower operating limit (LOL) of their power output. When turbines reach their LOL, even for a short time, that limits the available regulation capacity R for the entire interval, as described in equation (1). However, energy loss E_{loss} is cumulative over the entire interval, so it is not affected much when turbines briefly reach their LOL. Thus the cost of regulation capacity, calculated with equation (2) increases sharply when the denominator R decreases (because turbines reach their LOL) but the numerator E_{loss} changes very little.

The results in Fig. 2 exclude data where the actual regulation capacity is very close to the maximum available regulation capacity ($R_{max} - R < 0.1$ MW) because the costs approach infinity. Excluding those points significantly reduces the 95th percentile cost but changes the median MC very little. In addition to excluding those points, we also exclude dispatch intervals when the mean wind farm power is outside the linear range ($P_{\mu}(k) < 21$ MW or $P_{\mu}(k) > 50$ MW) and dispatch intervals when the maximum available upregulation capacity is near zero ($R_{max}(k) < 0.5$ MW).

The cost of regulation capacity decreases for shorter dispatch intervals, though the difference becomes smaller as R_{max} increases. For example, the results in Fig. 2a and C show that the median AC of 10 MW of regulation capacity is 1.03 MWh/MW-h for a 15-minute dispatch interval and 1.04 for a 60-minute dispatch interval when R_{max} = 15 MW.

The results in Fig. 2 are calculated for wind speed data from west Texas in 2008, but results calculated for the other locations and other intervals listed in Table 1 show the same trends. Results calculated for the Great Plains wind data show lower average and marginal costs than the other sites; we believe the costs are lower because the a wind farm at the Great Plains site has a significantly higher capacity factor than the other two sites—41%, as compared to 33% for the west Texas site in 2008 and 28% for the Ontario site.

4.3 Cost-Effectiveness of Curtailing for Frequency Regulation

Curtailing a wind farm can very rarely provide frequency regulation for less than the market price of regulation. We compare the minimum AC in each 1-hour dispatch interval to the market price of upregulation (MCPCU = "Market-Clearing Price of Capacity – Up") in the corresponding dispatch interval in the ERCOT (Texas) market [31]. The cost of upregulation capacity from a curtailed wind farm in a given interval is the minimum AC in that interval multiplied by the opportunity cost of a megawatt of wind power production. We plot the results in Fig. 3 as a cumulative distribution (CDF) of premiums that must be paid for regulation capacity from a curtailed wind farm, above the market price. The results are sensitive to the market price of upregulation and to the opportunity cost of curtailment.

Negative cost premiums in Fig. 3 correspond to intervals when wind farm curtailment produces upregulation capacity for less the market price. The cost of curtailment-derived upregulation capacity is less than the market price in approximately 1% of the 1440 1-hour dispatch intervals studied in 2008, and approximately 0% of the dispatch intervals in 2007 and 2009. If the grid operator is willing to pay a premium up to $50/MW-h, the wind farm can provide regulation capacity in 8% of the 1-hour dispatch intervals in 2007, 15% in 2008, and 5% in 2009. These results assume an opportunity cost for curtailing the wind farm of $62/MWh, made up of a wholesale energy price of $40/MWh, the federal Production Tax Credit (PTC) of $21/MWh, and a Renewable Energy Certificate (REC) price of $1/ MWh estimated by Wiser and Bollinger [32]. To examine the sensitivity of the results to the opportunity cost, we plot in Fig. 4 the percentage of 1-hour dispatch intervals when upregulation capacity from a curtailed wind farm costs less than the market price against the opportunity cost of curtailment. These results show that a wind farm with a lower opportunity cost will be competitive in the upregulation capacity market more often than a wind farm with higher opportunity costs.

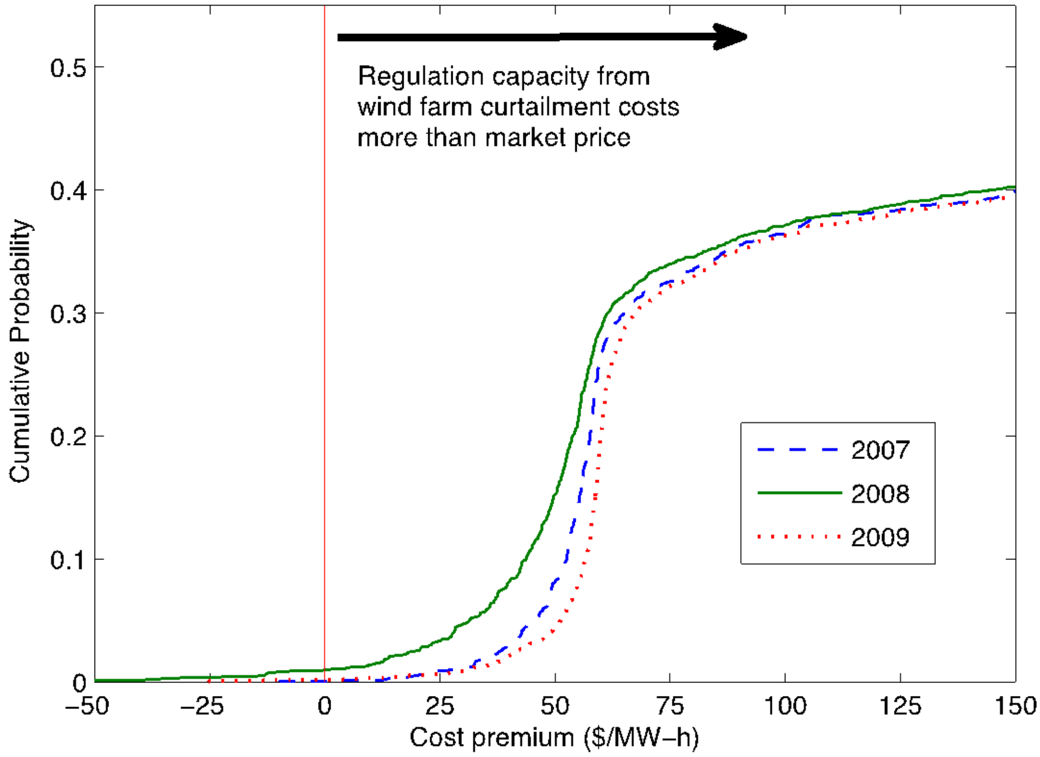

Fig. 3 A CDF of the cost premium for regulation capacity from a curtailed wind farm, as compared to ERCOT market prices for upregulation. The curtailed wind farm can provide upregulation at less than the market price in fewer than 1% of the 1440 1-hour intervals studied. We assume an opportunity cost for curtailing the wind farm of $62/MWh.

In practice, a wind farm operator may bid less than the full opportunity cost for upregulation capacity because he or she would receive payments for extra energy produced when that upregulation capacity is dispatched, sometimes called upregulation *energy*. Most other players in the upregulation capacity market already bid prices based on expectations of the amount of upregulation energy that will be dispatched.

The results in Fig. 3 and Fig. 4 are calculated for wind speed data from west Texas in 2008, but results calculated for the other locations and other intervals listed in Table 1 show the same trends. We compare regulation capacity costs calculated with wind speed data from the other sites listed in Table 1 to regulation capacity market prices from Texas, so the results do not account for any correlation between regulation market prices and wind conditions. Regulation capacity costs calculated with Great Plains wind data are approximately half of costs calculated with wind data from the other two sites, though they are still only competitive with the market price in 3.5% of the intervals studied. We believe the costs are lower because the a wind farm at the Great Plains site has a significantly higher capacity factor than the other two sites—41%, as compared to 33% for the west Texas site in 2008 and 28% for the Ontario site.

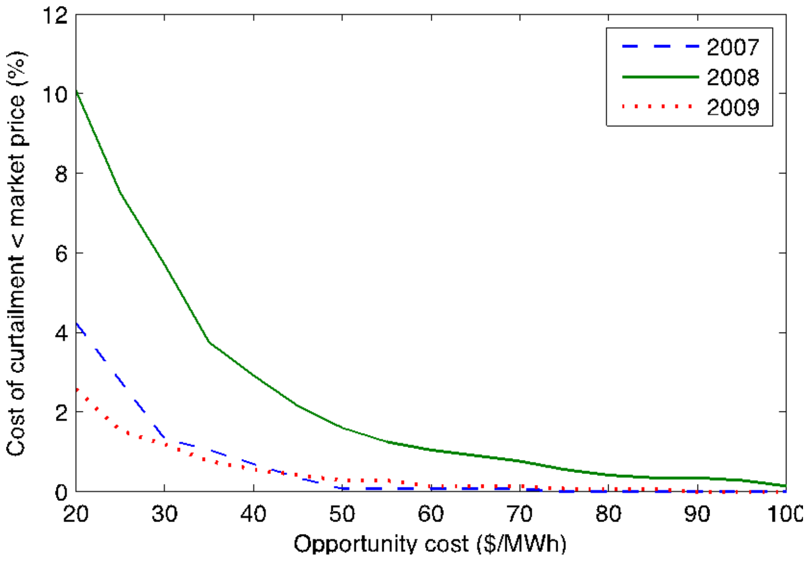

Fig. 4 Cost-effectiveness of wind farm curtailment for upregulation as a function of opportunity cost of curtailment. As the opportunity cost of curtailment increases, upregulation from curtailment is cheaper than the market price for curtailment in fewer 1-hour dispatch intervals

The results in Fig. 3 and Fig. 4 exclude dispatch intervals when the mean wind farm power is outside the linear range ($P_{\mu}(k)$ < 21 MW or $P_{\mu}(k)$ > 50 MW) and dispatch intervals when the maximum available upregulation capacity is near zero ($R_{max}(k)$ < 0.5 MW). Those results also exclude points where the actual regulation capacity is very close to the maximum available regulation capacity ($R_{max} - R$ < 0.1 MW) because the costs approach infinity.

These results are best-case scenarios based on the assumption of perfect forecasting. The results are likely to be similar even with imperfect forecasting because Fig. 2 shows that the cost of upregulation capacity does not diverge much from the minimum when the wind farm is not curtailed by the optimum amount. However, the cost of upregulation capacity from a curtailed wind farm is rarely competitive with Texas market prices.

5. Conclusions

A curtailed wind farm can rarely provide frequency upregulation at a cost lower than the present U.S. regulation market price, even if wind conditions can be forecast with perfect accuracy. We find that even with the high prices for upregulation capacity seen in the Texas (ERCOT) market in 2008, a curtailed wind farm with average opportunity costs could produce upregulation capacity at a cost less than the market price only 1% of the time. In other years with lower market prices for upregulation capacity, a curtailed wind farm would almost never be competitive. Curtailing a wind farm for frequency upregulation may be worthwhile in electrical grids where fast-ramping conventional generators are very expensive such as Hawaii, where diesel generators produce most of the regulation.

Several factors put wind farms at a disadvantage in a competitive market for upregulation capacity. First, the structure of some government subsidies for wind energy increase the opportunity cost of unproduced energy. When wind turbines are subsidized based on energy production, a curtailed wind farm loses both the revenue and the subsidy for unproduced energy. Second, thermal generators have lower opportunity costs for unproduced energy because their lost revenue is partially offset by fuel savings. Wind turbines have no significant variable costs, so they receive no offsetting savings for curtailment.

However, it is reasonable for grid operators to require that wind farms install the capability to curtail for frequency regulation. Curtailing a wind farm has a high operating cost, the opportunity cost of unproduced energy, but a very low capital cost. For grid operators, requiring delta curtailment capability from wind farms creates an emergency source of frequency regulation that may become more useful as wind power penetration increases. Wind farm owners already accept the requirement to be able to curtail for frequency regulation as a cost of connecting to the grid in some places.

If it is necessary to curtail a wind farm to provide upregulation capacity, there are several ways to minimize the cost. First, wind farms should be curtailed to approximately half of the maximum available upregulation capacity, as shown in Fig. 2. If a wind farm is required to provide a small quantity of regulation capacity, it is better to deeply curtail a few of the turbines than evenly spread the curtailment over all the turbines. Some grid codes, such as those for E.On (Germany) and EirGrid (Ireland) require wind farms to curtail by a few percent of their rated power to create reserve power for frequency regulation [5, 7]. Curtailing only a few turbines in each wind farm to meet these requirements would increase the amount of reserve power for a given curtailment and decrease the variability of the size of the reserve. Second, wind farms with low opportunity costs should be curtailed first. Wind farms that sell power on low-price long-term contracts or farms that no longer receive production subsidies are the most likely candidates.

This work was supported in part by the EPA STAR fellowship program, a grant from the Alfred P. Sloan Foundation and EPRI to the Carnegie Mellon Electricity Industry Center, and by the Doris Duke Charitable Foundation, the Department of Energy National Energy Technology Laboratory, the R.K. Mellon Foundation and the Heinz Endowments for support of the RenewElec program at Carnegie Mellon University. This research was also supported in part by the Climate and Energy Decision Making (CEDM) center. This Center has been created through a cooperative agreement between the National Science Foundation (SES-0949710) and Carnegie Mellon University.

6. References

1. Brooks, D.: Evaluation of the Effectiveness of Automatic Generation Control (AGC) Alterations for Improved Control with Significant Wind Generation. Electric Power Research Institute, Palo Alto (2009).

2. Evaluation of the Impacts of Wind Generation on HELCO AGC and System Performance -- Phase 2. Electric Power Research Institute, Palo Alto (2009).

3. Katzenstein, W., Apt, J.: Air Emissions Due To Wind And Solar Power. Environ Sci Technol. 43, 253–258 (2009).

4. Tsili, M., Patsiouras, C., Papathanassiou, S.: Grid code requirements for large wind farms: A review of technical regulations and available wind turbine technologies. Proceedings of EWEC. (2008).

5. Ramtharan, G., Jenkins, N., Ekanayake, J.: Frequency support from doubly fed induction generator wind turbines. IET Renew. Power Gener. 1, 3 (2007).

6. Singh, B., Singh, S.: Wind power interconnection into the power system: a review of grid code requirements. The Electricity Journal. 22, 54–63 (2009).

7. EirGrid: EirGrid Grid Code: WFPS1, Controllable Wind Farm Power Station Grid Code Provisions, version 3.3. (2009).

8. de Almeida, R., Lopes, J.: Participation of Doubly Fed Induction Wind Generators in System Frequency Regulation. IEEE T Power Syst. 22, 944–950 (2007).

9. Holdsworth, L., Ekanayake, J., Jenkins, N.: Power System Frequency Response from Fixed Speed and Doubly Fed Induction Generator-based Wind Turbines. Wind Energy. 7, 21–35 (2004).

10. Jecu, C., Teninge, A., Roye, D., Bacha, S., Belhomme, R., Bousseau, P.: Contribution to frequency

control through wind turbine inertial energy storage. Presented at the 2008 European Wind Energy Conference & Exhibition, Brussels (2008).

11. Lubosny, Z., Bialek, J.: Supervisory control of a wind farm. IEEE T Power Syst. 22, 985–994 (2007).
12. Rawn, B., Lehn, P., Maggiore, M.: Control methodology to mitigate the grid impact of wind turbines. IEEE T Energy Conver. 22, 431–438 (2007).
13. Vigueras-Rodríguez, A., Sørensen, P., Cutululis, N.A., Viedma, A., Gómez-Lázaro, E., Martin, S.: Application of ramp limitation regulations for smoothing the power fluctuations from offshore wind farms. Presented at the 2009 European Wind Energy Conference and Exhibition, Brussels (2009).
14. Kirby, B., Milligan, M., Ela, E.: Providing Minute-to-Minute Regulation from Wind Plants. Presented at the 9th Annual International Workshop on Large-Scale Integration of Wind Power into Power Systems, Quebec October 19 (2010).
15. Parsons, B., Milligan, M., Zavadil, B., Brooks, D.: Grid impacts of wind power: a summary of recent studies in the United States. Wind Energy. 7, 87–108 (2004).
16. Ela, E., Milligan, M., Parsons, B., Lew, D., Corbus, D.: The evolution of wind power integration studies: Past, present, and future. Presented at the IEEE Power & Energy Society General Meeting, 2009, Calgary July (2009).
17. Rose, S., Apt, J.: Generating wind time series as a hybrid of measured and simulated data. Wind Energy. 15, 699–715 (2012).
18. Veers, P.S.: Three-Dimensional Wind Simulation. SAND88-0152. 40 (1988).
19. IEC 61400-1: Wind Turbine Design Requirements (ed. 3.0). International Electrotechnical Commission. 22–31 (2005).
20. Sørensen, P., Cutululis, N.A., Vigueras-Rodríguez, A., Madsen, H., Pinson, P., Jensen, L.E., Hjerrild, J., Donovan, M.H.: Modelling of power fluctuations from large offshore wind farms. Wind Energy. 11, 29–43 (2008).
21. De Tommasi, L., Gibescu, M., Brand, A.J.: A dynamic aggregate model for the simulation of short term power fluctuations. Procedia Computer Science. 1, 269–278 (2010).
22. Sørensen, P., Cutululis, N.A., Vigueras-Rodríguez, A., Jensen, L.E., Hjerrild, J., Donovan, M.H., Madsen, H.: Power Fluctuations From Large Wind Farms. IEEE T Power Syst. 22, 958–965 (2007).
23. Jonkman, J., Butterfield, S., Musial, W., Scott, G.: Definition of a 5-MW reference wind turbine for offshore system development. National Renewable Energy Laboratory, Golden, CO (2009).
24. Grunnet, J.D., Soltani, M., Knudsen, T., Kragelund, M., Bak, T.: Aeolus Toolbox for Dynamics Wind Farm Model, Simulation and Control. Presented at the 2010 European Wind Energy Conference & Exhibition, Warsaw April 22 (2010).
25. Elkraft, Eltra: Wind Turbines Connected to Grids with Voltage above 100kV. (2004).
26. Hansen, A.D., Sørensen, P., Iov, F., Blaabjerg, F.: Centralised power control of wind farm with doubly fed induction generators. Renewable Energy. 31, 935–951 (2006).
27. Koenker, R., Bassett, G.: Regression Quantiles. Econometrica. 46, 33–50 (1978).
28. Koenker, R.: quantreg: Quantile Regression, http://CRAN.R-project.org/package=quantreg.
29. R Core Team: R: A Language and Environment for Statistical Computing, http://www.R-project.org/.
30. Vaz, S., Martin, C.S., Eastwood, P.D., Ernande, B., Carpentier, A., Meaden, G.J., Coppin, F.: Modelling species distributions using regression quantiles. J Appl Ecol. 45, 204–217 (2007).
31. ERCOT: Day-Ahead Ancillary Services Market Clearing Prices for Capacity Archives, http://www.ercot.com/mktinfo/prices/mcpc.
32. Wiser, R., Bolinger, M.: 2010 Wind Technologies Market Report. (2011).