

Excess Variance in Decentralized Renewable Energy Investment

Richard L. Sweeney and Joseph Wilske*

July 15, 2025

Abstract

The system value of variable renewable energy (VRE) investments depends not only on expected power production, but also on the covariance of production with other intermittent resources. Deregulated electricity markets do not provide sufficient incentives for renewable developers to fully internalize their impact on system variance. We empirically investigate the extent of this inefficiency in wind power investments in the United States. Using high-frequency, spatially granular estimates of wind production potential, we show that alternative investment programs which reallocate existing investment to locations with less correlated wind resources could substantially reduce system variability without sacrificing total output.

Keywords: Renewable Energy, Electricity Market Design

*Economics Department, Boston College. Financial support from NBER and the Sloan Foundation is gratefully acknowledged. We thank Matt Kotchen and Catherine Wolfram for detailed comments.

1 Introduction

Electricity markets must balance at every point in space in time. Rather than relying on an unbounded spot energy market and a responsive demand system to achieve this goal, wholesale power markets in the United States impose price caps, and employ a combination of ancillary services and out of market actions by the system operator to ensure reliability. Although essential for decarbonization, variable renewable energy (VRE) generators pose unique challenges to this system for two reasons ([Joskow, 2019](#)). First, VRE resources, such as wind and solar, are inherently intermittent. Second, VRE plants have zero marginal cost, meaning that, when available, they will be dispatched first. Compared to traditional generation, VRE thus increases the variability of supply. This, in turn, raises the operational challenges and costs of balancing the grid.

Although individual VRE plants are inherently intermittent, aggregate supply variability and uncertainty could be substantially smaller if renewable developers select sites where resource availability is relatively uncorrelated. However, developers do not face incentives to internalize the system costs of their entry. Increased VRE uncertainty requires the grid operator to procure a larger amount of backup generation than it would for an equivalent amount of non-VRE generation. But ancillary service fees are typically socialized across all demand, and not directly tied to the impact of specific generators. The combination of uncertainty of availability and certainty of dispatch conditional on being available also shifts the non-VRE generation mix towards more flexible resources, which are generally more expensive. Although capacity markets are supposed to provide a signal for the system value of entry, compensation schemes for VRE have historically been based on coarse measures of average availability by technology class. For these reasons, it is likely that potential gains from diversifying the location of VRE investments are not fully realized in practice.

In this paper, we quantify excess variability in wind production in the United States. Wind capacity has expanded considerably over the past decade, and now represents 12% of total generating capacity. At the same time, the grid has increasingly struggled to utilize this increased capacity. In 2023, 4.6% of wind production in the seven largest electricity markets (ISO's) was curtailed ([U.S. Department of Energy, 2023](#)), highlighting that the system value of wind capacity was considerably lower than the sum of its parts, due to the high correlation in wind speeds across locations.

To conduct our analysis, we construct high frequency (20-minute) potential wind power estimates at 2 square kilometer grid cells covering the entire continental United States. We then group these cells into twelve regions to measure the impact of entry at each location on aggregate reliability. Using these data, we document three facts. First, we show that there is

considerable heterogeneity in the correlation of production across potential wind farm sites within each region. Importantly, this heterogeneity persists even if we condition on a site's expected production, suggesting that aggregate variance can be reduced without sacrificing aggregate output. Second, using regression, we demonstrate that while private site-specific characteristics (mean output and variance) strongly predict actual wind investment decisions, the covariance of wind production across sites does not. Third we show that, as a result, considerable regional variability persists. The coefficient of variation (CV) for a typical wind site in the United States is around 1. Despite the fact that there are now over 70,000 wind turbines installed across the country, we find that the aggregate coefficient of variation for wind production ranges from 0.4 to 0.6 across the seven largest electricity markets in the United States.

Having documented that there appears to be scope for variance reduction in every region, we adopt a portfolio optimization approach to quantify the potential gains from a coordinated investment program (Wolak, 2016). Holding total regional investment and annual production fixed, we search for an optimal turbine location portfolio which minimizes regional variance. Initially, we constrain the optimizer to reallocate capacity across locations with existing wind turbines. This intensive margin reallocation reduces regional system coefficients of variation by 10 to 20 percent, with the largest gains happening in the Western Interconnection. If we allow for reallocation from existing sites to locations without current wind investment, but which are near existing transmission lines, the potential reductions in regional variability are more than twice as large. In a robustness section, we show that these reductions persist if we restrict output to match observed output by time of day and season, and if we instead constrain the optimization to match observed revenue, rather than output.

We then explore the potential mechanism behind our measured excess variance. As the underlying driver of covariance across plants comes from spatial proximity, we find that counterfactual investment portfolios succeed in reducing system variability by reducing the spatial density of investment. Consistent with the fact that the source of inefficiency is driven not simply by VRE variability, but also by the fact that deregulated markets reward primarily private benefits, we find that states more reliant on traditional regulated utility investment are associated with lower levels of counterfactual reallocation. We also look for effects of state level renewable portfolio standards (RPS), but find no evidence that states with higher renewable requirements see larger reductions in wind investment in our coordinated investment programs.

Existing Literature

Our analysis is motivated by several papers discussing the conceptual problems that VRE pose for electricity markets, some of which contain empirical evidence supporting the validity of these concerns (Joskow, 2019). Hogan (2010) provides a thorough discussion of the potential problems posed by VRE intermittency, but notes that “the individual variability of solar or wind facilities is less of a problem when there is sufficient regional diversification of sources where winds speeds are not correlated.” We show that, despite wind capacity now making up over 12% of total generating capacity, substantial correlation across locations persists. Wolak (2016) quantifies the scope for variance reduction from a portfolio reoptimization of capacity across existing solar and wind sites in California as of 2012. We adopt the optimization framework from that paper, and extend that analysis to include all wind farms installed as of 2024. We also consider extensive margin reallocation, shifting capacity to sites that have not yet experienced wind investment. Weber and Woerman (2024) estimate the impact of hourly wind intermittency and generation uncertainty on prices and marginal operating costs in Texas. In contrast, we consider the aggregate potential impact of wind variability over longer time horizons, inclusive of portfolio effects.

This paper also builds on the “excessive entry” literature in industrial organization. Mankiw and Whinston (1986) show that markets characterized by fixed entry costs and imperfect competition tend toward excessive entry due to a business-stealing effect. Amir et al. (2014) generalize this result, and show that it holds in the absence of convex operating costs. With large up-front costs and zero marginal operating costs, VRE contain the basic ingredients identified by this literature for entry to be socially excessive (ignoring other positive externalities, such as emissions displacement). However, we note that electricity markets often institutionalize an extreme form of business stealing, whereby the production of existing generation is displaced by new entrants not through pecuniary margins, but through out-of-market actions (curtailment). Using historical hourly generation data in California, Novan and Wang (2024) find that that only 91% of the output supplied by new solar capacity goes towards increasing the state’s renewable supply, with the remaining potential production being lost to curtailment.

2 Background


The central challenge of an electrical power system is maintaining instantaneous balance between electricity generation and consumption. Due to the physical characteristics of electricity, the power dispatched must precisely match power consumed at every moment across

the entire network. This requires that, if a wind turbine currently producing suddenly becomes unavailable, there must be sufficient slack capacity at another nearby resource capable of ramping up production immediately. Overproduction is also problematic. This means that if a wind resource experiences a sudden increase in output, there must be sufficient scope for currently generating resources to ramp down.

Conceptually, a fully market-based electric power system could use prices to manage this problem. In the short-run, an unbounded continuum of interconnected local spot prices should efficiently coordinate availability and dispatch among the existing generation mix. In practice, prices are locationally and temporally coarsened, and capped at price levels deemed to be exorbitant. In the long run, capacity markets could compensate generators for the reliability value of their presence during hours when they are not actually called upon to produce. However, to date, these markets have failed to provide sufficiently targeted investment signals in practice. Although the exact rules vary across markets, compensation for renewable energy was typically based on average availability rather than marginal system value, failing to reflect the spatial externalities created by correlated renewable generation. As a result, wind plant developers make entry decisions largely based on expected revenue.

Given the inadequacy of price signals in the market, grid operators often resort of non-pecuniary methods of ensuring adequate supply and safety. One important tool is the ability to curtail renewable energy production, despite the plant being willing to supply power at the market clearing price. The exact rules for curtailment again vary across markets, but they are generally implemented in a manner which does not rely on some well-defined production “right” or ex ante ordering of plant priority. Figure 1 shows the rate of wind curtailment in each of the seven largest ISOs in the United States. The average rate of curtailment was 4.6%, with much higher levels observed in Texas and the Southwest Power Pool in recent years. These curtailment rates are suggestive of potentially large gains from reallocating investment towards less correlated locations in these markets.

Figure 1: Wind Curtailment and Penetration by ISO



Sources: ERCOT, MISO, CAISO, NYISO, PJM, ISO-NE, SPP

Source: LBNL Land-Based Wind Market Report: 2024 Edition

3 Data

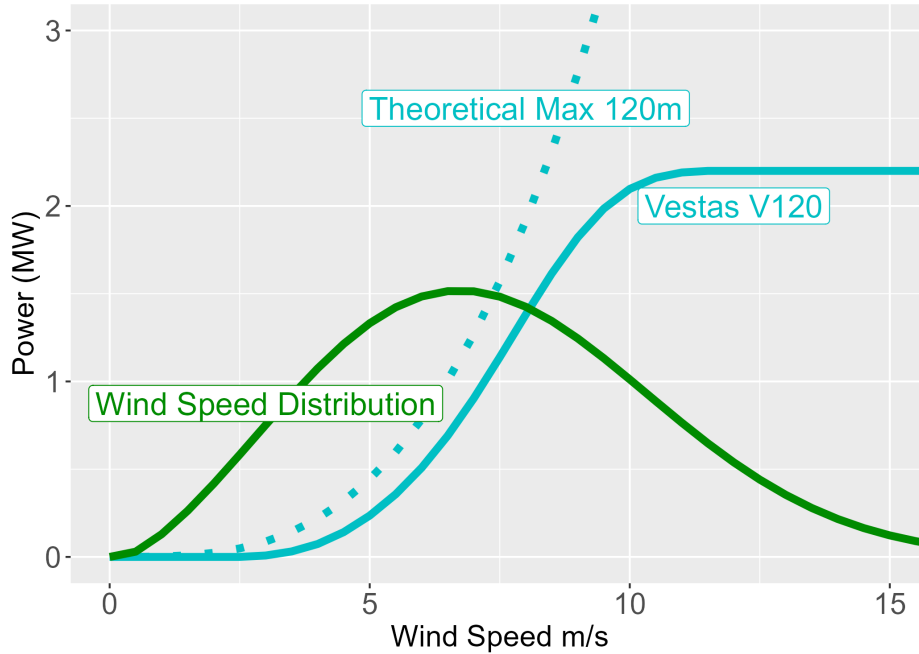
To conduct our analysis, we construct a database of engineering-based 20-minute frequency wind production estimates.¹ The National Renewable Energy Laboratory produced high frequency reanalysis estimates of wind speeds over two square-kilometer grids covering the continental United States. We use the data from 2019, at a hub height of 80 meters.

We use a wind turbine “power curve” to convert high frequency wind speed data into estimates of hourly production. Power curves are manufacturer provided production functions mapping wind speed to electric power production. This conversion is important, as the relationship between wind speed and wind power is nonlinear. We compute output measures using a common turbine for all sites - Vestas’ V120 2.2 MW turbine. This is one of the most common turbines installed in the United States in the past decade.² Figure 2 shows the power curve for this turbine, along with the wind distribution at a typical US site.

¹For additional discussion of this approach to engineering based production estimates, and a discussion of their ability to predict actual production, see [Aldy et al. \(2023\)](#).

²Power curves vary across turbines, driven primarily by the length of the rotor blades ([Covert and Sweeney, 2022](#)). Although we observe the exact turbine installed at every existing wind farm, we do not observe what turbine would be installed on locations that have not yet experienced entry. For this reason, we use a common power curve for all locations.

Figure 2: Example Wind Distribution and Power Curve



Notes: “Wind Speed Distribution” is a probability density of wind speeds at a typical location with installed wind turbines in the United States. The “Theoretical Max” function plots the theoretical maximum output that could be captured from a 120-meter wind turbine at each speed. The “Vestas V120” power curve displays the amount of electricity produced by the Vestas’ 120-meter turbine at each speed.

The wind data cover over 4.7 million grid points, most of which are not suitable for wind development. At most locations, the wind speeds are too low to make investment economically viable. At other locations, surface use patterns or restrictions preclude installation. To restrict attention to potentially relevant points of entry, we adopt the sample of locations identified the National Renewable Energy Laboratory’s (NREL) Techno-Economic WIND Toolkit. This sample is based on a combination of wind resource suitability criteria, and legal and practical restrictions on wind development. These restrictions included the full exclusion of airfields, urban areas, wetlands, and water bodies, and partial exclusion of forested regions that were not situated on ridgetops.³ NREL also added back all remaining sites with observed wind turbines as of 2015. The final sample includes 126,693 onshore grid cells across the contiguous United States.

We combine this data on potential wind turbine locations with detailed information on the universe of wind turbines currently installed in the US from the USGS. This database

³Complete detail on the Techno-Economic site selection is given in [Hodge \(2016\)](#).

contains the exact location of each wind turbine and the date of initial operation. It also contains, in nearly all cases, the exact turbine model installed and rated capacity. Any locations with observed entry that fell outside of the NREL sample were added back to our sample population.

Table 1 presents counts of potential wind farm locations, divided up geographically into NERC subregions.⁴ We excluded regions from the sample with less than 1000 MW of installed capacity, which effectively removes the south-eastern United States. The top map in figure 3 shows the regions in our sample. The bottom map displays the full set of potential wind locations identified by NREL. The points with observed entry are displayed in black. In our analysis below, we consider an additional restriction of the NREL sample to include only grid cells within 2 miles of a transmission line. Cell counts for this sample are displayed in the “Near Grid” column. The final column presents total installed regional capacity as of 2024.

Table 1: Sample Observations

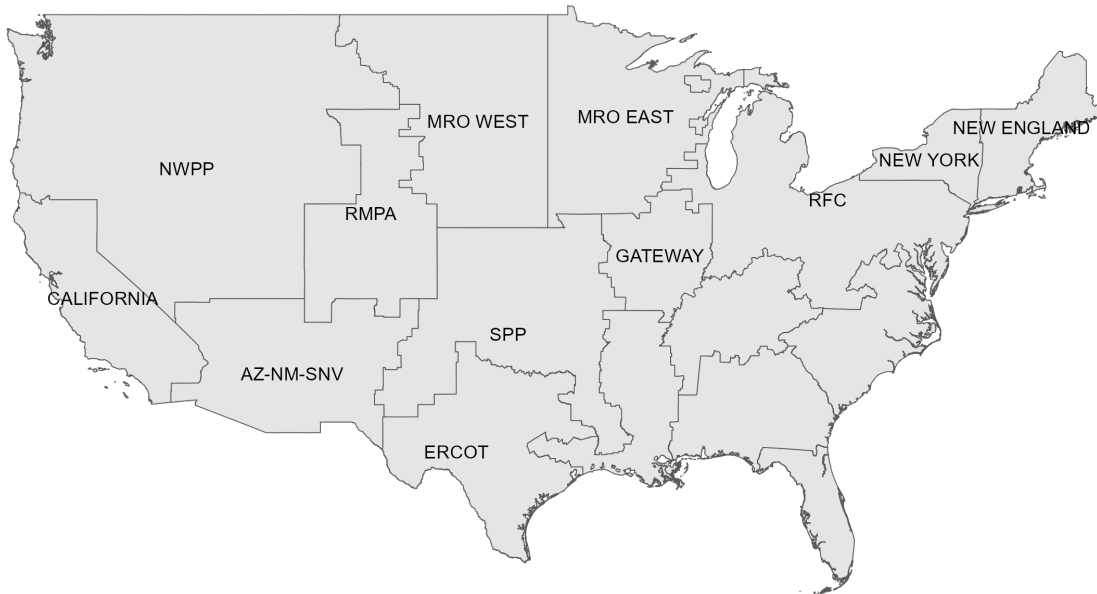
NERC	Subregion	N Sites	N Near Grid	N Entry	Capacity (GW)
WECC	AZ-NM-SNV	7781	1551	399	3.7
WECC	CALIFORNIA	3692	1578	256	5.6
TRE	ERCOT	7568	5069	3633	33.7
NPCC	NEW ENGLAND	2855	652	218	1.6
WECC	NWPP	19181	4062	1097	12.3
WECC	RMPA	10007	2950	962	7.3
SPP	SPP	17680	9119	4187	34.5
NPCC	NEW YORK	2960	1235	323	2.8
RFC	RFC	14235	7226	1696	12.8
SERC	GATEWAY	4558	2273	868	6.7
MRO	MRO WEST	14510	5593	1748	12.3
MRO	MRO EAST	14510	7957	2425	17.4

Notes: Observations are the number of grid cells in each region. The full sample includes all sites within a NERC subregion. The “Near Grid” sample includes all sites within 2 miles of a transmission line. The “N Entry” sample includes all sites with at least one wind turbine installed as of 2024.

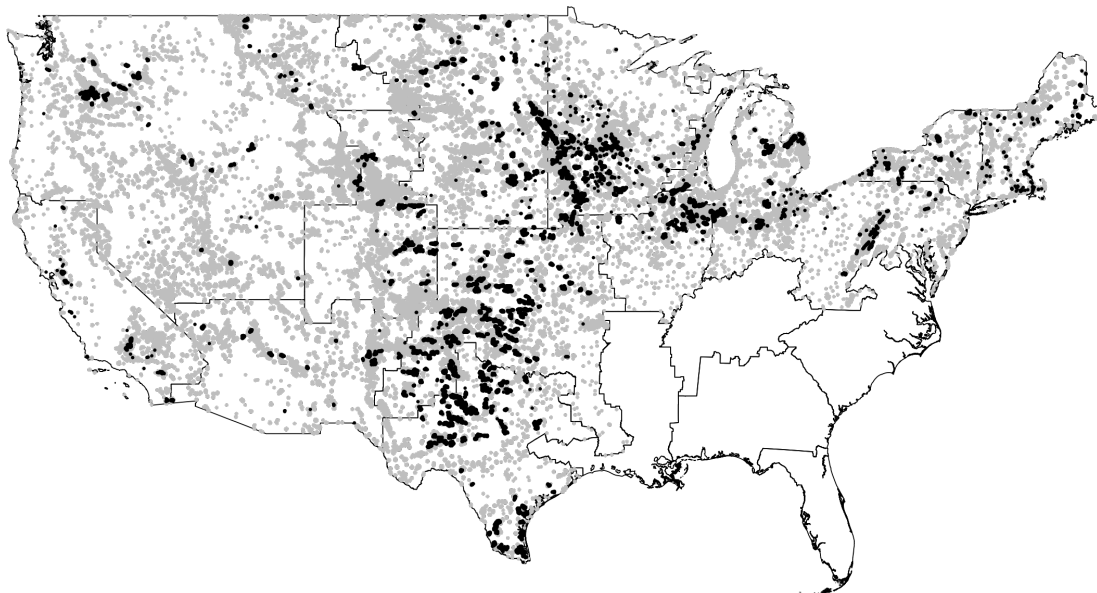
⁴The last two regions in the table, MRO EAST and MRO WEST, are not NERC designated subregions. The geographic area of MRO is quite large, and contained the largest number of potential points of wind entry. We split the region into an east and west to facilitate the computations below.

Figure 3: Region Definition and Sample Grids and Entry

(a) Subregions



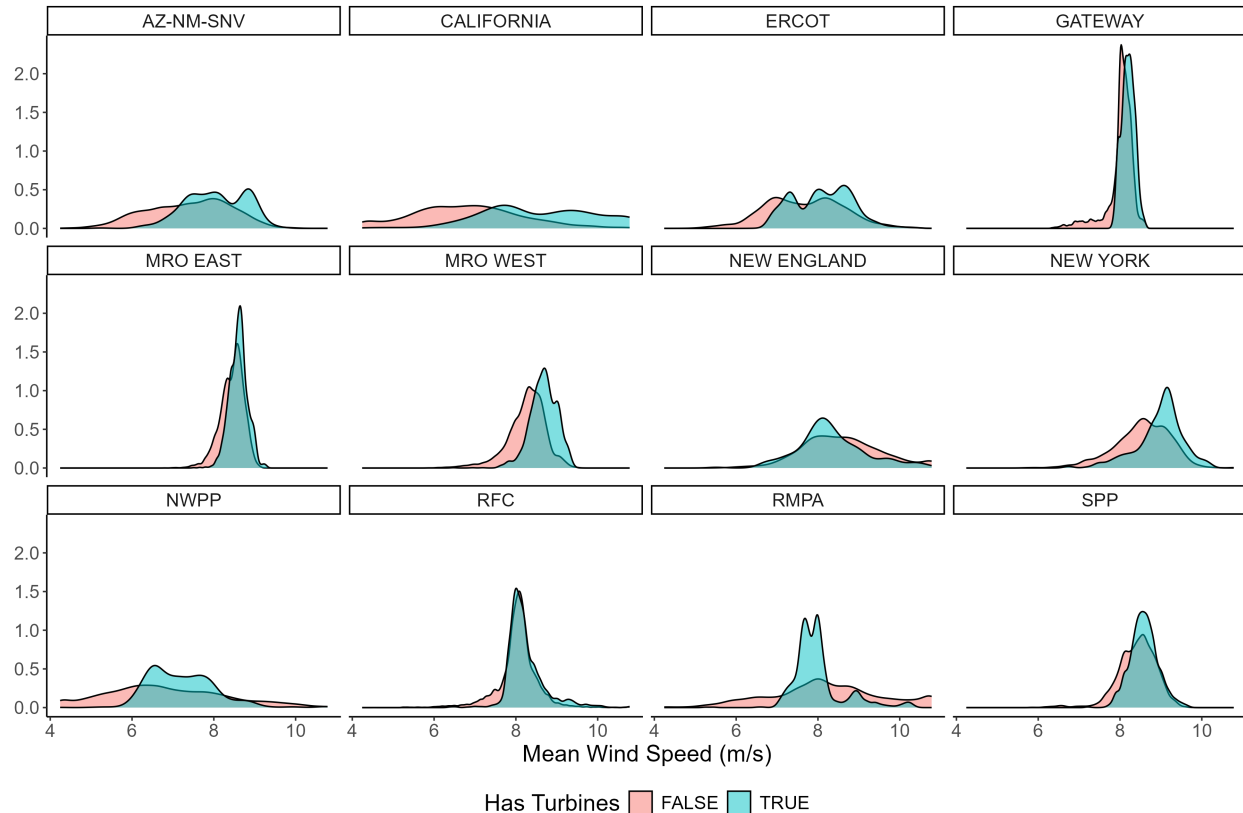
(b) Sample Grids



Notes: The top map shows the NERC subregions used in this analysis. The bottom map shows the full set of potential wind locations identified by NREL. The points with observed entry are displayed in black.

To provide a sense of the dispersion in wind quality within each region, figure 4 displays the distribution of grid mean wind speeds, separated out by whether the grid contains any existing wind capacity. As expected, sites with observed entry have higher mean wind speeds. However, there is a substantial amount of overlap between the distribution of sites with and without entry.

Figure 4: Distribution of Site Mean Wind Speeds by Region and Entry Status



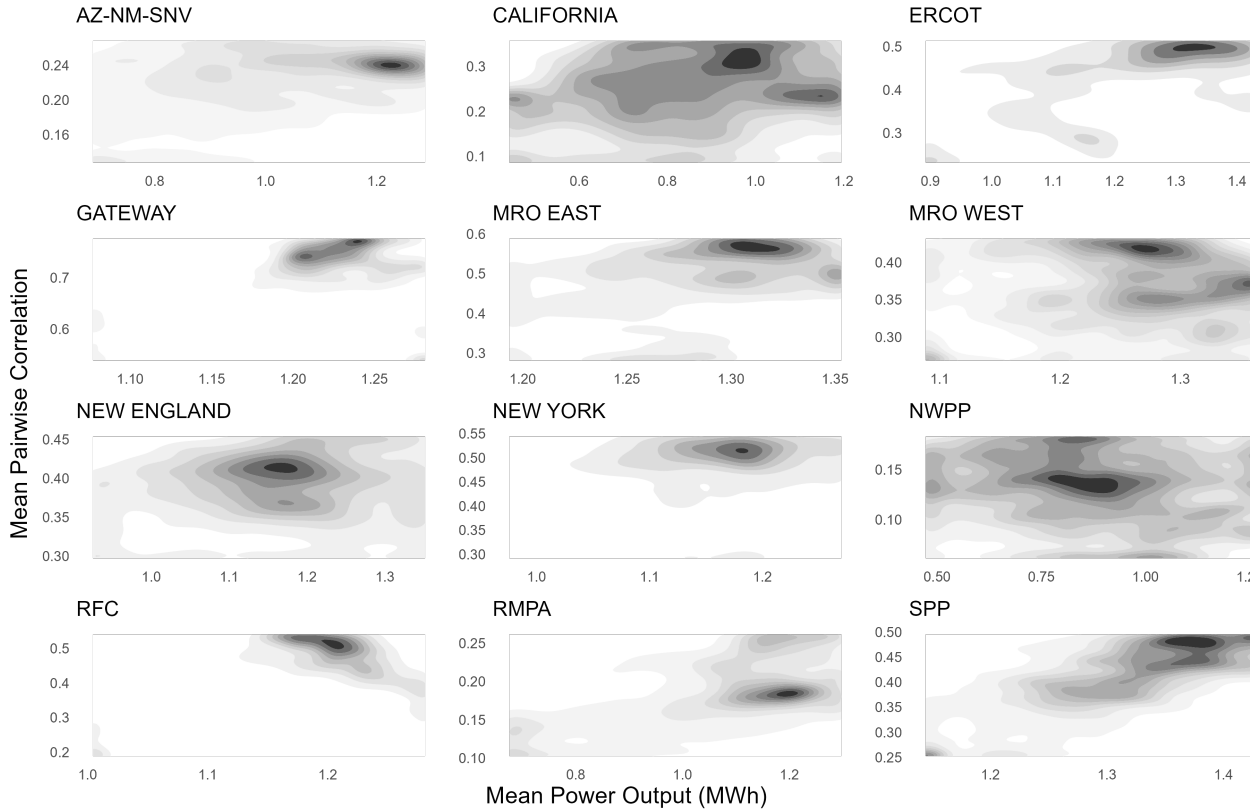
Notes: The sample is restricted to onshore locations identified as feasible for wind development by [Hodge \(2016\)](#). “Has Turbines” reflects grids with at least one observed wind turbine as of 2024.

4 Preliminary analysis

Our primary interest in this paper is understanding the extent to which aggregate (regional) wind production variance could be reduced by reallocating capacity across sites. A necessary condition for this to be true is that there exists meaningful differences in the covariability of production across sites. To explore this, we compute the pairwise correlation in predicted power output across all site pairs in the same region. We then take the average of these pairwise correlations for each site.

Figure 5 plots the mean pairwise density in the average pairwise correlation for each site against that site’s expected power output, by region. Taking California as an example, we see that average pairwise correlation is more than three times larger at some sites than others (span of the y-axis). Looking at the x-axis, the difference in expected output across sites is nearly as large. Based on the shading, it’s clear that these two measures are correlated, suggesting a tradeoff between expected output and system stability, on average. However, there is also considerable vertical dispersion at each expected output level, suggesting that substantial differences in aggregate variability would obtain across sites, holding the change in expected output fixed. Similar patterns hold for all regions, although the scale of the differences and correlation vary.

Figure 5: Density of Mean Wind Speeds and Pairwise Correlations



Notes: The x-axis is the expected output at each site, in MWh. The y-axis is the average pairwise correlation in output across all other sites in the same region.

Figure 5 shows considerable differences in covariance across sites within region. Next we ask whether these differences actually influence entry decisions. To analyze this, we split the sample into grids that had at least one turbine as of 2019, and those that did not. For the

latter group, we compute the marginal increase in aggregate variance that would be added to each region if the site were to enter.⁵ We then estimate the following linear probability model.

$$\text{Entry}_i = \beta_1 \text{Power}_i + \beta_2 \text{Variance}_i + \beta_3 \Delta \% \text{ System Variance} + \beta_4 \text{Transmission Dist.} + \epsilon_i \quad (1)$$

where Power and Variance are the expected output and variance in output at site i . $\Delta \% \text{ System Variance}$ is the marginal increase in system variance from adding site i , divided by the aggregate variance of production from sites that entered prior to 2019. Table 2 provides summary statistics for these variables for the sample of grids which had not experienced entry as of 2019.

Table 2: Entry Regression Sample Summary Statistics

Variable	Mean	Standard Deviation
Power Mean (MWh)	1.13	0.22
Site Variance	0.76	0.06
Marginal System Variance (%)	0.24	0.33
Transmission Distance	7.27	9.13
Entry	0.05	0.22

Notes: Sample restricted to grid cells that had not entered as of 2019. The marginal system variance is the percent increase in regional variance from adding each grid. Transmission distance is the distance to the power grid, in miles. Entry is a binary indicator for whether wind investment is observed by the end of 2024.

Table 3 present the regression results from estimating equation 1. As the variables have considerably different scales, we first convert them to z-scores, so the coefficients can be interpreted as the impact of a one standard deviation increase on the probability of entry post 2019. Column 1 includes region fixed effects. As expected, the productivity of a site has a strong positive influence on entry, while site variability has a strong negative influence, as does distance to the transmission grid. Our primary coefficient of interest is the percent change in system variance. If wind developers were internalizing the impact of their entry on system reliability, system variance should *negatively* influence entry. However, here we estimate a coefficient that is actually positive and significant. Column 2 replaces region fixed effects with state fixed effects. In this model, the impact of power, own variance and

⁵The difference in regional aggregate variance with and without site i is the variance at site i plus two times covariance between i and all existing sites. This latter term is the marginal covariance for site i .

Table 3: Entry Regressions

	(1)	(2)	(3)	(4)
Power	0.043 (0.011)	0.045 (0.009)	0.091 (0.025)	0.101 (0.022)
Variance	-0.025 (0.007)	-0.026 (0.009)	-0.050 (0.016)	-0.046 (0.018)
Δ % System Variance	0.041 (0.011)	0.015 (0.008)	0.043 (0.013)	0.012 (0.011)
Transmission Distance	-0.009 (0.005)	-0.010 (0.006)	-0.034 (0.010)	-0.035 (0.010)
Observations	107,140	107,140	116,103	116,103
Entry Year	2019	2019	2009	2009
Region FE	X		X	
State FE		X		X

Notes: Sample restricted to grid cells that had not entered as of the listed Entry Year. The dependent variable is an indicator for entry by 2024. Power and Variance refer to the expected output and variance in output at each site i . All variables are converted to z-scores. Standard errors clustered at the state level.

transmission distance are similar. The impact of system variance is much smaller, however we can still statistically reject the hypothesis that system variance reduces the probability of entry.

Columns 3 and 4 present the same model, looking at entry post 2009. The results are similar, although the coefficient on system variance is smaller and not statistically significant. Across these two models, we fail to find any evidence that covariance with existing capacity deters entry or investment in the US wind industry, consistent with the discussion above explaining that these plants face little incentive to internalize system variability costs.

5 Portfolio reoptimization

To explore the potential gains from a centralized investment program, which would internalize the covariance in production across locations, we follow the portfolio choice approach of [Wolak \(2016\)](#). For each region, we search for an optimal turbine location portfolio which minimizes the variance for a given level of expected aggregate production.

Let M be a set of potential wind entry points in a region, indexed by i .⁶ Let C be a vector of potential capacity factors at each site, which is the expected output, in MWh, per

⁶Regional subscripts are suppressed for ease of exposition.

unit of capacity installed. If W is a vector of capacity allocated to each site under some allocation, then the expected level of aggregate production from the allocation is $W'C$. Let Ω be a covariance matrix of capacity factors for all sites in M . Then the variance from an allocation is given by $W'\Omega W$.

With this notation, we can write the optimal portfolio as the solution to the following quadratic program,

$$\begin{aligned} \min_{\{w_i\}} \quad & W'\Omega W \quad \text{s.t.} \\ & \widetilde{W}'C \leq W'C \\ & 0 \leq w_i \leq \bar{w} \\ & \sum_{i \in M} w_i \leq \sum_{i \in M} \tilde{w}_i \end{aligned}$$

\widetilde{W} is a vector of *observed* installed capacities. The first constraint therefore imposes that any counterfactual allocation generate at least as much aggregate output, in expectation, as the observed real world allocation. The second constraint imposes non-negative capacity allocations, and sets a maximum capacity per site of \bar{w} .⁷ Finally, the third constraint imposes that the counterfactual allocation not include more total investment than what is actually observed in the data in each region.⁸ It is possible, however, that some variance minimizing allocation achieves the observed level of output at a *lower* level of aggregate regional investment. In this case, focussing on variance alone would potentially understate the gains from portfolio reoptimization.

For each region, we solve this quadratic program over an increasingly large set of potential sites M .⁹ In the first exercise, which we label “Intensive Margin” reallocation, we restrict M to the set of locations containing existing wind turbines. This is a natural starting point, as these locations are demonstrably suitable for wind investment and sufficiently close to transmission infrastructure. As the total level of capacity is fixed, the algorithm simply attempts to reallocate capacity away from sites with high covariance towards sites with lower covariance.

We then consider two extensive margin reallocations of capacity away from sites with existing investment towards sites which are not yet developed. Although NREL specifically

⁷ \bar{w} is set to 22 MW (10 Vestas V120 turbines), which is approximately the 99%th percentile of capacity per site in the sample with observed investment.

⁸We exclude capacity installed prior to 2000. These turbines are considerably smaller and less productive than our reference turbine. As of 2024, they are also likely not operating often.

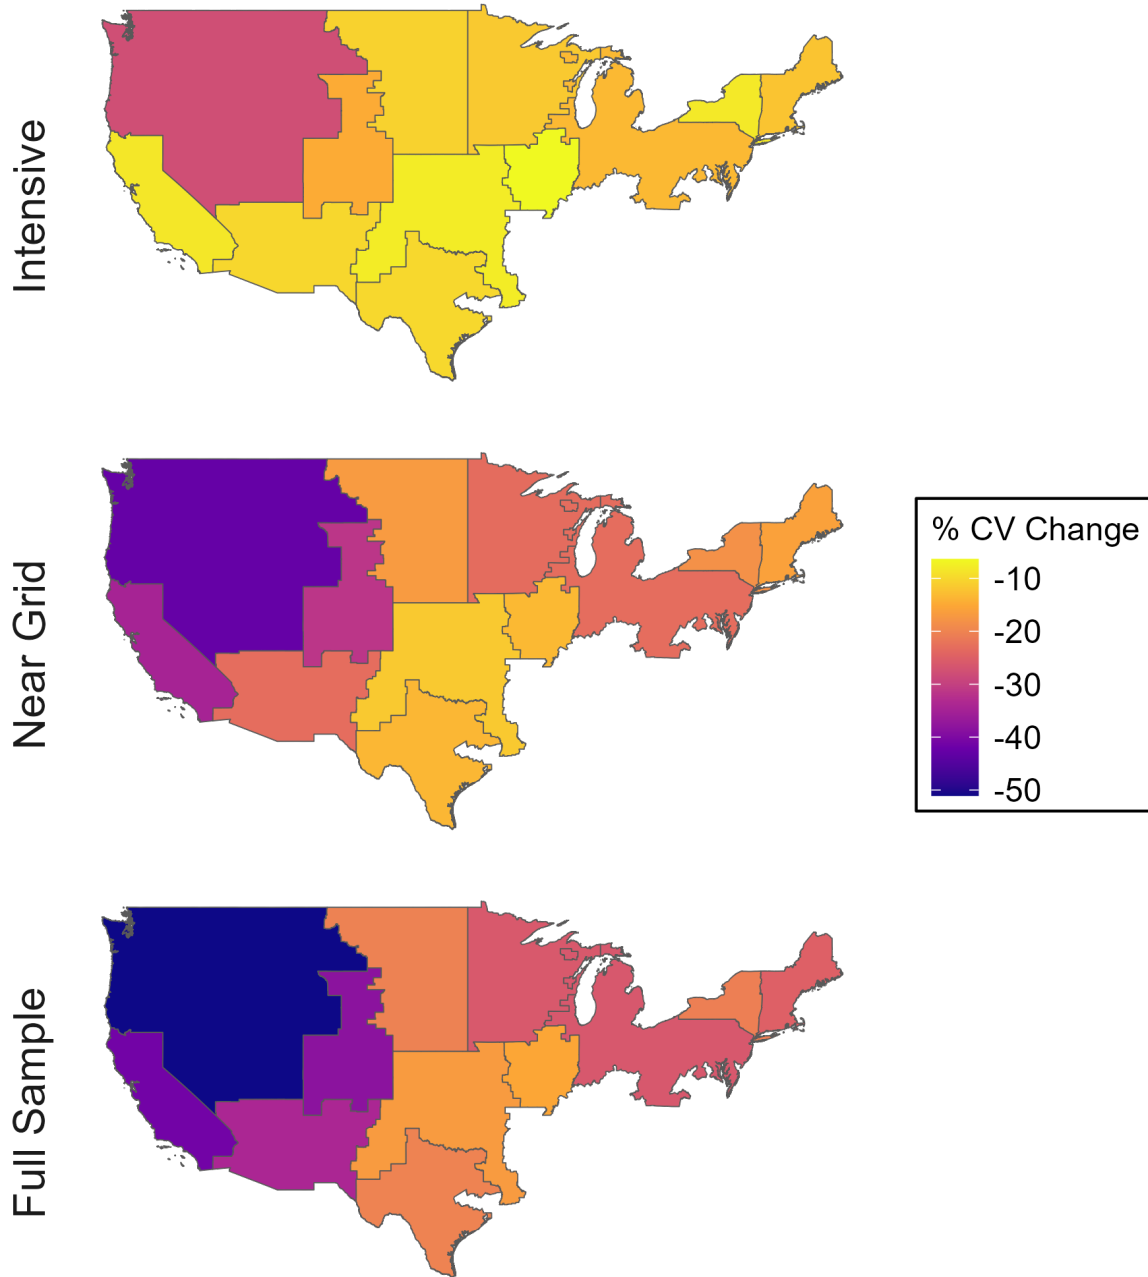
⁹We use the OSQP Julia package, <https://github.com/osqp/OSQP.jl>, which is able to solve large quadratic programming problems with many constraints.

selected sites suitable for wind development based on surface characteristics, they did not account for proximity to the power grid. To account for this, we first restrict the set of potential undeveloped points to locations within two miles of a transmission line.¹⁰ We label this exercise “Near Grid” reallocation. The second extensive margin reallocation allows for reallocation across all sites in the NREL Techno-Economic dataset. We label this exercise “Full Sample” reallocation.

Figure 6 contains maps summarizing the regional level reallocation results. The shading in each region is determined by the percent reduction in the coefficient of variation (CV) under the reallocation, relative to the baseline CV. Starting at the top, there is a considerable amount of variance reduction obtainable simply by reallocating capacity across existing sites. This result matches the findings of Wolak (2016) for the state of California. Moving to the “Near Grid” scenario, the reduction in CV is roughly twice as large, on average. Allowing for reallocation across all sites produces a further reduction in CV, but the average reduction is not much larger than the “Near Grid” scenario.

¹⁰Transmission line locations were obtained from an archived version of the now restricted Homeland Infrastructure Foundation Level Database (HIFLD). According to the map metadata, this shapefile was last updated in 2023.

Figure 6: Reductions Regional Variance from Reallocation



Notes: For each investment program, we compute the aggregate coefficient of variation (CV) of output across all sites in the region. The shading indicates the percent reduction in CV from the observed allocation to the reallocation. The “Intensive Margin” reallocation restricts investment to existing sites. The “Near Grid” reallocation allows for investment at undeveloped sites within 2 miles of a transmission line. The “Full Sample” reallocation allows for investment at all undeveloped sites.

5.1 Additional constraints

In the preceding re-optimization exercises, we minimized aggregate variance under the constraint that aggregate output, over the course of the year, was at least as large as the observed output. However, electricity is more valuable at some times than others. It might be the case that existing wind production is highly correlated because developers place more value on output at certain times.

To explore whether this force could be important in driving wind developers towards highly correlated sites, we consider two additional constraints on the optimization problem. In our first exercise, we break hours of the day into a “peak” period, between 8 a.m. and 8 p.m., and an “off-peak” period, from 8 p.m. to 8 a.m. We break months of the year into a “winter” season, from November to April, and a “summer” season from May to October. The cross of these two categories divides the year into four equal time periods. We then amend the output constraint above to restrict all counterfactual allocations to obtain at least as much output during each time period. Table 4 presents the results. For context, the third column presents the observed coefficient of variation (CV) for each region, computed using the observed investment locations. The fourth column presents the percentage CV reduction from the preceding section, where annual output is constrained to match in any counterfactual investment allocation. The fifth column, labeled “Output-Time” requires that aggregate match by peak and off-peak for winter and summer months. Although the percentage reductions are smaller than the output case, the CV is much closer to the output constrained case than it is to the observed CV. The two largest deviations from the Output case occur in the WECC NERC region, suggesting that if timing concerns are driving investment anywhere, it is in these regions.

One limitation of the previous exercise is that we do not know which times of day or year are most valuable. This could also differ across regions. To more directly proxy for system value, we bring in hourly wholesale price information. We obtained hourly day-ahead nodal prices for 18 representative nodes from SNL Energy.¹¹ We then match each grid cell to the closest price point, and compute expected revenue as the product of expected output and the day-ahead price.¹² We then add an additional constraint to the optimization program that expected revenue in any counterfactual allocation must match observed aggregate annual revenue. The resulting percentage reductions in regional CV are presented in the “Revenue” column of table 4. In most cases, the results are identical to the changes in the Output

¹¹We use day ahead prices rather than real time prices, as the latter will be more endogenously determined by wind realizations at observed investment locations.

¹²Although not every grid point in our sample lies within an ISO, we still assign an ISO day-ahead price to every location. As the electric power grid is interconnected, and regions trade power, these prices still provide an informative measure of system value.

Table 4: Coefficient of Variation Reductions - Near Grid Sample

NERC	Subregion	Observed CV	Counterfactual CV Change (\%)		
			Output	Output-Time	Revenue
MRO	MRO EAST	0.51	-23	-16	-24
MRO	MRO WEST	0.42	-17	-16	-17
NPCC	NEW ENGLAND	0.50	-16	-16	-15
NPCC	NEW YORK	0.55	-18	-17	-13
RFC	RFC	0.48	-23	-22	-22
SERC	GATEWAY	0.61	-13	-13	-14
SPP	SPP	0.43	-12	-11	-9
TRE	ERCOT	0.45	-14	-12	-13
WECC	AZ-NM-SNV	0.44	-23	-22	-22
WECC	CALIFORNIA	0.49	-35	-26	-36
WECC	NWPP	0.45	-43	-38	-40
WECC	RMPA	0.46	-32	-29	-31

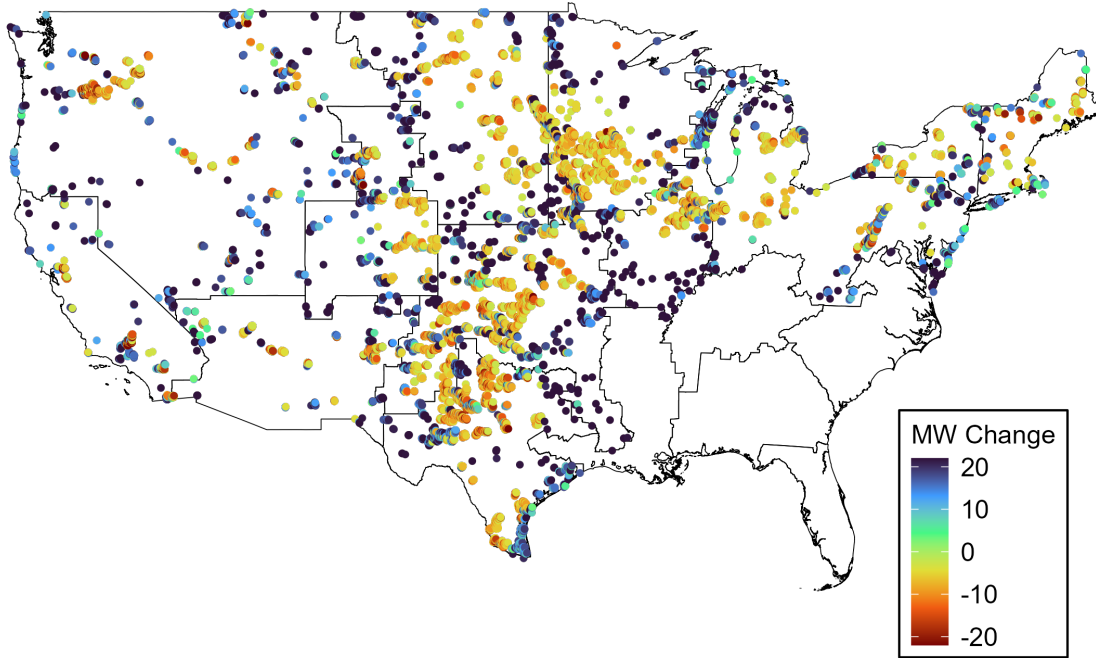
Notes: Observed CV is the aggregate coefficient of variation for each region computed using observed investment locations. The final three columns report the percent reductions in CV, relative to the observed CV. The Output column restricts the reallocated aggregate output match observed output. The Output-Time column requires that aggregate match by peak and off-peak for winter and summer months. The Revenue column requires that counterfactual “revenue” (predicted output times day ahead nodal prices) match. All counterfactual allocations restrict entry to grid cells within 2 miles of a transmission line.

column, suggesting that the output constrained reallocation was already achieving at least as much aggregate revenue over the course of the year as the observed allocation.

5.2 Discussion of mechanisms

To explore the potential drivers of misallocation, we first provide a visual comparison of the differences in allocations across space. Figure 7 maps the changes in capacity under the “Near Grid” reallocation. The shading indicates the change in capacity at each site, relative to the observed capacity. Negative values (dark red) indicate points in space where capacity was removed, and positive values (dark blue) indicate points where the portfolio reoptimization increased capacity. Although the sample is restricted to points close to transmission infrastructure, the blue points (capacity increases) are visually much more disperse than the red dots (capacity decreases), which tend to be quite concentrated. Since wind patterns are spatially correlated, this map is consistent with the real-world decentralized investment program concentrating too much investment in a few high output locations.

Figure 7: Capacity Changes Under Near Grid Reallocation



Notes: Plot of site MW under the Near Grid optimized portfolio minus observed MW. Sites with an absolute difference of less than one turbine are excluded from the plot for clarity.

Next we consider differences in misallocation across states. There are two relevant state-level policies. The first is electricity deregulation. As described in [Joskow \(2019\)](#), variable renewable energy poses challenges to electricity markets because renewable energy developers base entry decisions on solely on a project's expected revenue in wholesale markets. Their variability increases system costs, but these costs are not directly internalized by the developer. In these markets, the entry decision is driven by expected revenue, which does not account for the impact on system variability. However, roughly half of the United States electric power system still operates under a regulated utility investment and dispatch mode. For these power systems, the operational costs of supply variability are similar, but projects are procured by a regulated utility tasked with minimizing average costs. In these markets, it is possible that the regulated utility considers the impact on system variability when selecting investment locations.

Another potentially important driver of misallocation across states is renewable policy. As of 2023, 28 states and Washington, DC had passed some form of renewable portfolio standard (RPS). These policies require that a certain share of consumption in each state come from qualifying renewable sources. Using panel variation, previous research found that these policies cause an increase in wind energy generation within a state's borders ([Feldman](#)

and Levinson, 2023; Hollingsworth and Rudik, 2019). If RPS policies bind more in states that already have a lot of wind, or where resources are relatively more correlated with regional wind output, then these policies could increase system wind variability. If, on the other hand, RPS policies direct investment towards locations where wind speeds are lower or less correlated with other regional resource, then RPS’s could actually reduce system variability.

We construct measures of these policies for each state. As discussed in Borenstein and Bushnell (2015), although electricity regulation is typically described as a binary state-level choice, the real situation is more continuous. We use data from the 2023 EIA 860 survey to characterize each existing wind plant as being regulated or not. We then compute the share of deregulated wind capacity in each state. For renewables policy, we use a continuous measure of stringency from Feldman and Levinson (2023).¹³ For each state, we compute a continuous measure of the state RPS share in 2019 as the ratio of required renewable generation in 2019 and the total state consumption of electricity in that year.

We then project changes in state level capacity under portfolio reallocation on these two measures. To deal with the fact that some states have many more potential sites than others, we compute each state’s utilization rate as the ratio of installed capacity to the maximum total capacity if every potential site in the sample were fully developed. We then estimate the following model,

$$\Delta\text{Utilization}_s = \alpha + \beta\text{RPS Share}_s + \gamma\text{Deregulated Share}_s + \epsilon_s$$

where $\Delta\text{Utilization}_s$ is utilization rate in state s under one of the three reallocation scenarios (“Intensive Margin”, “Near Grid” or “Full Sample”) minus the observed utilization rate.

Table 5 presents the results. Column 1 uses changes in utilization under intensive margin reallocation.¹⁴ The first cell shows that a ten percent increase in the share of deregulated generators is associated with a 4 percentage point reduction counterfactual investment intensity. A ten percent increase in state RPS share is associated with a 0.7 percentage point increase in a counterfactual capacity, although the standard error is much larger. The second column repeats this exercise with the Near Grid reallocation. The point estimate for deregulated share is now 0.3, while the estimate for RPS share remains statistically indistinguishable from zero. Column 3 presents the results under extensive margin reallocation. Here a ten percent increases in the share of deregulated plants reduces counterfactual investment intensity by only 1 percentage point. Column 4 presents the results from reallocation where output is constrained to match output by season and time of day. The point estimate on Deregulated share lies in between the near grid and full sample cases.

¹³Thank you to Arik Levinson for providing us with this data.

¹⁴Three states without any observed wind turbines are excluded from this sample.

Table 5: State Reallocation Regressions

	(1)	(2)	(3)	(4)
Deregulated Share	-0.397 (0.176)	-0.293 (0.086)	-0.098 (0.028)	-0.174 (0.065)
RPS Share	0.069 (0.427)	-0.058 (0.241)	-0.021 (0.080)	-0.053 (0.181)
Observations	39	42	42	42
R ²	0.125	0.244	0.250	0.172
Scenario	Intensive	Near Grid	Full Sample	Near Grid, By Time

Notes: The dependent variable is the share of potential capacity under a reallocation scenario minus the observed share.

Across these four models, we find that portfolio reallocation reallocates investment away from states with a larger share of deregulated plants. This is consistent with the idea that regulated plants are more likely to internalize the impact of their entry on system variability. Conversely, we find no evidence that state level renewable policy is meaningfully shifting the location of wind investment across states. It should be emphasized that these are cross-sectional regressions, purely meant as a descriptive exercise. Any omitted factors that effect wind investment and are correlated with either measure will confound a causal interpretation.

Although the variance-minimization framework employed here highlights the diversification benefits of spreading wind capacity across space, it omits several forces that can rationally encourage developers to cluster turbines spatially. Local land-use ordinances and siting rules or costs could be spatially correlated. While we have relied on the analysis of [Hodge \(2016\)](#), which attempted to remove locations with *observable* land-use restrictions, unobservable restrictions could persist in a manner that is spatially correlated. Second, fixed investment costs shared across turbines could generate increasing returns to scale at the plant level. One important example of this is grid interconnection costs ([Johnston et al., 2023](#)). Pulling in the opposite direction, however, are aerodynamic wake interactions. Modern wind turbines are massive, and they reduce wind speeds for miles in their wake. This can impose sizable productivity losses on downstream machines, meaning our simple expected output calculation, which does not account for wake effects, likely understates the physical cost of spatial concentration.

6 Conclusion

This paper quantifies the excess variability arising from decentralized investment in wind power in the United States, driven by the unpriced externalities resulting from correlated intermittency across locations. By combining detailed high-frequency wind speed data with observed investment patterns, we document that actual entry decisions systematically overlook the covariance of output among sites, resulting in higher than necessary variability in aggregate wind power production. Through a portfolio optimization analysis, we demonstrate that significant reductions in aggregate variability—holding total investment constant—could have been achieved by reallocating wind capacity from heavily concentrated areas to locations with less correlated wind profiles.

While some excess covariability is unavoidable under decentralized entry due to imperfect information, policy could help minimize excess variance during the energy transition. First, market-design changes could sharpen the price signals facing developers: ancillary-service charges could be levied directly on generators in proportion to their contribution to balancing costs; curtailment protocols could grant dispatch priority to earlier entrants; and capacity-market compensation could be more directly tied to each site’s empirical availability profile. Second, policymakers could ease land-use restrictions at and extend transmissions towards locations that are particularly uncorrelated with the existing VRE fleet. Future work should quantify the gains from such policies across alternative time periods and, with turbine-specific power curves, refine estimates of the socially optimal spatial deployment of wind capacity. It would also be interesting to explore whether the advent of cheap battery storage mitigates the costs of spatial concentration.

References

- Aldy, Joseph E, Todd D Gerarden, and Richard L Sweeney**, “Investment versus Output Subsidies: Implications of Alternative Incentives for Wind Energy,” *The Journal of the Association of Agricultural and Resource Economists*, 2023, 10 (4), 981–1018.
- Amir, Rabah, Luciano De Castro, and Leonidas Koutsougeras**, “Free entry versus socially optimal entry,” *Journal of Economic Theory*, November 2014, 154, 112–125.
- Borenstein, Severin and James Bushnell**, “The US Electricity Industry After 20 Years of Restructuring,” *Annual Review of Economics*, 2015, 7 (Volume 7, 2015), 437–463.
- Covert, Thomas R and Richard L Sweeney**, “Winds of change: Estimating learning by doing without cost or input data,” Technical Report, Working paper 2022.

- Feldman, Rachel and Arik Levinson**, “Renewable portfolio standards,” *The Energy Journal*, 2023, 44 (5), 1–20.
- Hodge, Bri-Mathias**, “Final Report on the Creation of the Wind Integration National Dataset (WIND) Toolkit and API: October 1, 2013 - September 30, 2015,” Technical Report NREL/SR-5D00-66189, 1247462, NREL April 2016.
- Hogan, William W**, “Electricity wholesale market design in a low-carbon future,” in “Harnessing Renewable energy in electric power systems,” Routledge, 2010, pp. 113–136.
- Hollingsworth, Alex and Ivan Rudik**, “External impacts of local energy policy: The case of renewable portfolio standards,” *Journal of the Association of Environmental and Resource Economists*, 2019, 6 (1), 187–213.
- Johnston, Sarah, Yifei Liu, and Chenyu Yang**, “An empirical analysis of the interconnection queue,” Technical Report, National Bureau of Economic Research 2023.
- Joskow, Paul L**, “Challenges for wholesale electricity markets with intermittent renewable generation at scale: the US experience,” *Oxford Review of Economic Policy*, April 2019, 35 (2), 291–331.
- Mankiw, N. Gregory and Michael D. Whinston**, “Free Entry and Social Inefficiency,” *The RAND Journal of Economics*, 1986, 17 (1), 48–58.
- Novan, Kevin and Yingzi Wang**, “Estimates of the marginal curtailment rates for solar and wind generation,” *Journal of Environmental Economics and Management*, 2024, 124, 102930.
- U.S. Department of Energy**, “Land-Based Wind Market Report: 2023 Edition,” Technical Report 2023.
- Weber, Paige and Matt Woerman**, “Intermittency or uncertainty? impacts of renewable energy in electricity markets,” *Journal of the Association of Environmental and Resource Economists*, 2024, 11 (6), 1351–1385.
- Wolak, Frank A.**, “Level versus Variability Trade-offs in Wind and Solar Generation Investments: The Case of California,” *The Energy Journal*, 2016, 37, 185–220. Publisher: International Association for Energy Economics.