The effect of wind power investments on rural labor markets.

Johannes Mauritzen
BI Norwegian Business School
johannes.mauritzen@bi.no
jmaurit.github.io

January 15, 2018

Abstract

I examine the labor market effects of wind power investment in rural counties in the US. I combine quarterly panel data on county employment and wages with data on all wind power plant investments larger than 1 mega-watt (MW). I argue that wind power investments can to a high degree be considered exogenous to outcomes in labor markets due to a dependence on average wind speeds. In addition, identification is achieved through a multilevel model where unobserved time-invariant county variables are controlled for. I find no significant effect on net employment, but find that on average, a mid-sized 200 MW wind farm leads to a permanent increase in wages of approximately 2.5 percent. The findings have implications for energy policy and provides a case-study that can inform broader questions and discussions of US labor market. I estimate the model with a Bayesian approach, using Markov Chain Monte-Carlo simulations.

1 Introduction

Economists have increasingly come to recognise that a large swath of the US population has been left out of the benefits of the last several decades of economic growth in the form of poor job prospects, stagnant wages and decreasing income mobility. This trend has been ascribed to increased automation and mechanisation, outsourcing to countries with cheaper labor costs, and lack of critical skills among the labor force [Autor et al., 2015, Acemoglu et al., 2015, Autor, 2014, Chetty et al., 2017]. In a related trend Case and Deaton [2015, 2017] document the dramatic reversal of health and morbidity outcomes among the working class, which are driven by substantial increases in "deaths of despair": Drug overdoses, suicides and alcohol-related disease.

While metro areas have not been immune to these trends, rural areas and small towns tend to have an over-representation of the demographics that have been most effected: non-hispanic white, middle aged, working class, and no college degree. Rural areas and small towns were particularly hard hit in the most recent recession and experienced steeper falls in both employment and wages [USDA, 2016].

But amidst a dearth of investments and deteriorating job possibilities in rural and small-metro counties, a bright spot has been investments in wind turbines and wind farms. The cost of wind power fell by 75 percent between 1984 and 2014 and is cost competitive in most locations in the US without subsidies [Trancik, 2015]. Wind power has then moved from being a niche and highly subsidized generation found mostly in rich states, to a competitive form of power generation that now makes up a significant portion of generation in states with substantial rural areas such as Iowa, South Dakota, Texas, and Wyoming. Decreasing costs and wider penetration has also meant that the wind power industry is playing a growing role in the US labor market as a whole. The U.S. Department of Energy (DOE) estimates that as of 2015, the wind power industry supported approximately 50,000 jobs. The DOE further extrapolates that if wind power penetration continues to grow, the industry could support up to 600,000 jobs by 2050.[1]

In contrast to traditional power plants, that tend to be relatively compact and are generally located close to population centres, modern land-based wind turbines are often over 80 meters tall with blade-lengths of over 100 meters [2]. Out of spatial necessity, investments in wind turbines tend to happen away from large population centres.

Investments in wind power will of course have an impact on economic growth and lead to job creation in both the manufacturing, installation and maintenance of the turbines.

[1]https://energy.gov/eere/articles/wind-energy-supporting-600000-jobs-2050.

[2]The wingspan of a 747 jumbo-jet is approximately 60 meters

They will also generate revenues for land-owners who either lease land for wind turbines or own the turbines directly, sometimes through a cooperative structure. However, it is ex-ante unclear how and to what extent these economic effects influence the local labor market.

Conceivably, it may make sense to employ skilled labor from outside the county hosting a wind power plant for both the initial build-out as well as subsequent maintenance and repair. Because wind turbine maintenance and repair is a skilled occupation, even if an in-county job is created, it is not clear to what extent this would be a net-increase in employment as opposed to a skilled worker moving from one position to another.

The role that leasing payments or profits from the sale of electricity has on local labor markets is also ex-ante unclear. Agricultural land is to a growing extent held by corporations or individuals who are not located in the same county or even state. The income from wind turbines in the county may, in many cases, end up flowing completely out of the county.

I use data from the Energy Information Agency form 860 on all wind power installations over 1 MW in the United States and match it with data on quarterly wage and employment data from the Bureau of Labor Statistics Quarterly Census to estimate the effect of wind power investments on wages and net-employment in rural counties.

The effects of wind power, and more generally renewable energy on economic growth and labor markets has been an active topic of research, especially in Northern Europe where generous subsidies led to early and sustained investment in renewable energy [Lehr et al., 2012, 2008, Ejdemo and Söderholm, 2015]. Studies of the US have been more sparse [Haerer and Pratson, 2015, Wei et al., 2010]. A common element of these studies is that they tend to be aggregated to the regional or national level, without considering the geographic distribution of economic effects. The results are often based on only partially empirical methods–large scale input-output models calibrated to aggregated data on investments and penetration, but highly dependent on modelling assumptions.

Comparisons can be made to the local economic effects of another recent energy boom. The Shale oil and gas boom, driven by technological advances in "fracking" [Gold, 2014] also primarily affected rural areas in the major petroleum-containing formations in the US.[3]. Komarek [2016], Weber [2012] and Brown [2014] all find substantial increases in employment and wages in counties that experience a boom in oil and gas extraction. Importantly though, these economic and labor market effects often retreat or disappear as oil and gas wells run dry.

Wind turbines, on the other hand, tend to have a mechanical life of over 20 years. More so, older wind power sites tend to get re-powered–that is the turbines get replaced by newer, more efficient turbines–as the wind resources and transmission infrastructure make such sites ideal for continued investment [Mauritzen, 2014]. The local economic effects of wind power may not be of the initial magnitude of petroleum finds, but the effects can be expected to be more permanent.

Unique to studies of investments on labor markets, wind power investment decisions can be seen as exogenous to labor markets. The most important factor in the profitability of a wind farm is the average wind speed of a location. However, there remains good reason to believe that labor market outcomes and investment in wind power could be partly endogenously determined. On the margin, counties more likely to attract wind power projects could, for example have the necessary transmission infrastructure in place, or have local governments that are more investment friendly, with stream-lined processes for permits and approvals. These unobserved variables could also be correlated with labor market outcomes.

Further aiding identification, I use a panel of data with 30 quarterly observations on labor market outcomes for every county in the United States. Intercepts and both deterministic and stochastic trends are allowed to vary by county. I then compare outcomes before and after a wind power investment, both measured as a temporary jump and permanent shift. These parameters are also allowed to vary by county. In turn, these county parameters are themselves modelled together with relevant county-level variables. An average effect is then estimated through a common distribution for the county parameters with associated

[3]The Bakken formation in North Dakota and Montana, The Marcellus in the north-east, and the Barnett in Texas

meta-parameters. From this multilevel structure, I can estimate an average treatment effect of wind power investment across counties while allowing for varying intercepts and trends by county. The multilevel model also has the attractive feature of automatic shrinkage of county coefficients through partial pooling. The model is fit using Bayesian Markov-Chain Monte Carlo (MCMC).

The results indicate that wind power investments have no significant effect on employment in rural counties. However, a significant effect on wages is found. A 200 megawatt (MW) wind farm–approximately the capacity of 60 modern turbines–leads to a median permanent increase in wages of 2.5 percent in rural counties.

Wind power has traditionally received subsidies–in the form of federal tax credits as well as various state and local incentives. Environmental and climate externalities are often given as the justification for such subsidies. This article informs renewable energy policy by suggesting a distributive effect of wind power policy. Claims that wind power will lead to significant job gains in struggling rural areas do not appear to have strong support in the data, though wind power investments do appear to press up wages.

The case of wind turbines also has the potential to inform broader questions about labor markets. The results of this research are consistent, though not exclusively so, with a skills-based explanation of the weak rural US labor market. When a lack of a skilled work force is the constraining factor, new job opportunities due to an exogenous placement of a wind farm would be expected to have limited effects on the number of employed, while still pressing up wages. This could be a promising route for researchers with access to detailed employment micro-data.

2 Data

I combine data from three sources. Data on investments in wind energy plants is from the U.S. Energy Information Agency (EIA) form 860.[4] This data provides yearly information

[4] https://www.eia.gov/electricity/data/eia860/

	County Rural-Urban Continuum Codes (RUCC)
1	Counties in metro areas of 1 million population or more
2	Counties in metro areas of 250,000 to 1 million population
3	Counties in metro areas of fewer than 250,000 population
4	Urban population of 20,000 or more, adjacent to a metro area
5	Urban population of 20,000 or more, not adjacent to a metro area
6	Urban population of 2,500 to 19,999, adjacent to a metro area
7	Urban population of 2,500 to 19,999, not adjacent to a metro area
8	Completely rural or less than 2,500 urban population, adjacent to a metro area
9	Completely rural or less than 2,500 urban population, not adjacent to a metro area
	Aggregated categories
1	Metro counties (1,2,3)
2	Non-metro with urban population, adjacent to a metro area (4,6)
3	Completely rural, or small urban population not adjacent to metro area. (5, 7, 8, 9)

Table 1: Rural Urban Continuum Codes obtained from the Department of Agriculture Economic Research Service (ERS) are aggregated into three broader categories.

on every power plant and planned power plant with capacity of over 1 MW in the United States. Data is at the generator level. Variables include the date of first operation, size of generator, county of generator, ownership, and grid connection.

Data on quarterly county-level labor market outcomes is from the Bureau of Labor Statistics (BLS) Quarterly Census of Employment and Wages.[5] Variables include average weekly wages and employment for each of the 3223 U.S. counties.

I classify the counties based on U.S. Department of Agriculture's Economic Research Service (ERS) Rural-Urban Continuum Codes (RUCC).[6] from 2013. County designations are updated every 10 years based on decennial Census data. RUCC codes go from 1-9, as defined in table 1. In order to simplify the analysis, I aggregate the designations into three broader categories, which are also shown in the lower pane of the table. The aggregated categories are meant to separate out metro areas and counties adjacent to metro areas from those that consist of rural areas and small towns not directly connected to a big city economy. From now on I will simply refer to these three categories as "Metro", "Adjacent metro" and

[5]https://www.bls.gov/cew/

[6]https://www.ers.usda.gov/data-products/rural-urban-continuum-codes/

"Rural".

Many recent analysis of the US labor market have used Commuting Zone (CZ), as developed by Tolbert and Sizer [1996], as the geographic unit. CZs approximate the labor markets associated with metro areas which often stretch across metropolitan and suburban counties. I do not, however, make use of CZ's as I am explicitly concerned with rural and small-town counties not adjacent to metro areas.

Additional data on county population and agricultural land values was obtained from the ERS [7]. These variables can clearly change over time, however they are only available at 10-year intervals, with the most recent year being 2013. In the analysis, these variables then appear as time-invariant county-level variables.

The upper pane of figure 1 shows the distribution of counties by rural-urban indicator. The lower pane of the figure shows the distribution of operating wind power plants. Rural counties have clearly seen a large share of wind power investments. Figure 2 shows that rural counties have been the location of nearly half the total wind power capacity and that capacity additions in rural areas more than doubled in the period studied.

Comparing the distribution of wind power in figure 1 to a map of average wind resources produced by the US National Renewable Energy Laboratory (NREL) (See figure 18 in the appendix) gives a visual impression of the high correlation between wind resources and the geographic investment decision. As mentioned, the most important factor in the capacity factor and in turn profitability of a wind turbine is the average wind speed of the turbine location. The physical relationship between power generation and average wind speed is cubic[8]. Average wind speed is then a dominating factor in the geographic investment decision and, arguably, exogenous to economic and labor market variables.

The upper pane of figure 3 shows that employment in non-metro areas has been largely stagnant since 2009 compared to metro areas. As the lower panel shows, however, wage

[7] https://www.ers.usda.gov/topics/farm-economy/land-use-land-value-tenure/farmland-value/

[8] A simplified equation for wind power output can be written $P = kC_p\frac{1}{2}\rho AV^3$, where P = Power output, C_p = Maximum power coefficient, ρ = Air density, A Rotor swept area, V = wind speed, k = a constant. [MacKay, 2016]

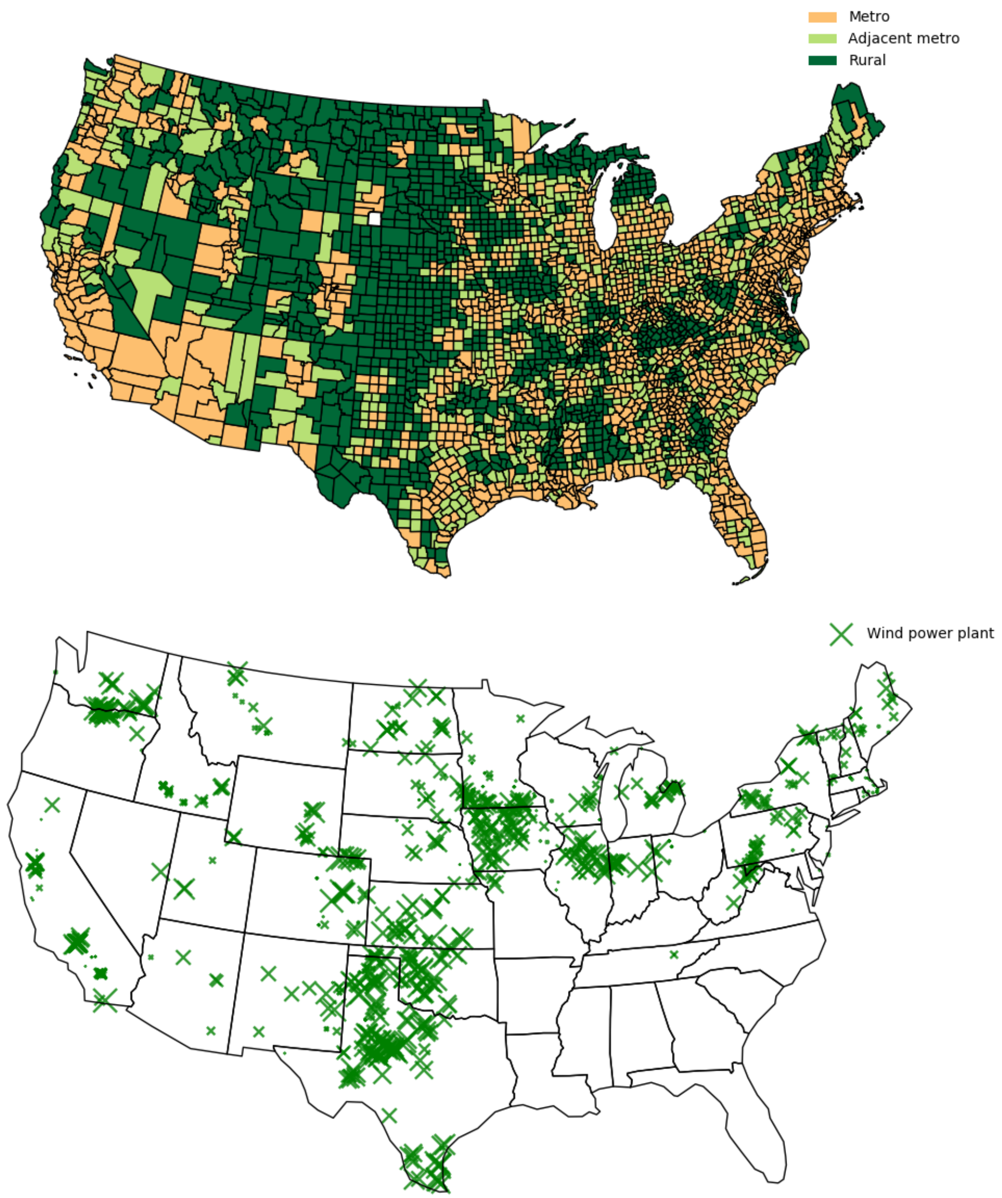

Figure 1: The upper pane of the figure shows the distribution of counties by the rural-urban indicator. The lower pane shows the distribution of wind power plants across the U.S. Wind power plants tend to be concentrated in rural counties.

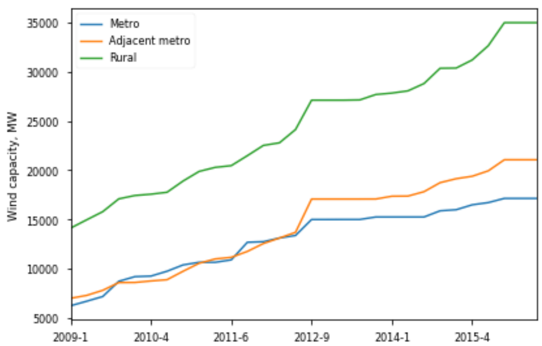

Figure 2: Almost half of all wind turbine capacity is located in rural areas. Wind power capacity more than doubled in the period studied.

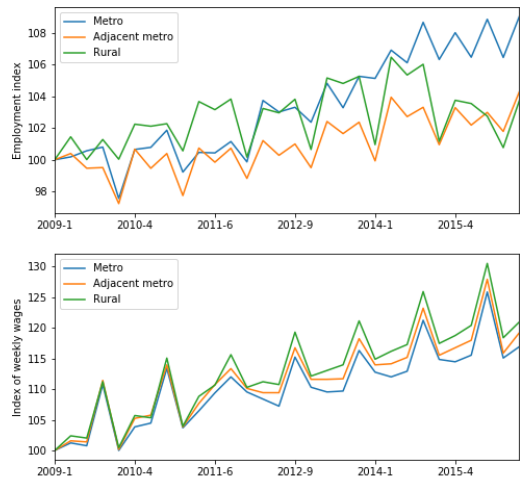

Figure 3: Employment growth in non-metro counties has lagged significantly behind employment in metro counties. Average wage growth, has, however been similar between metro and non-metro counties, through from a lower absolute level.

growth has been similar in both rural and metro areas, though rural wages fell more during the preceding recession.

3 A multilevel model of wind investment and labor market outcomes

In the model of the employment effects of wind power investment, the response variable is employment in county c at time t, $employment_{c,t}$. This variable is transformed by subtracting the county mean and dividing by two times the standard deviation of the county employment. This transformation allows for meaningful comparisons of counties of different populations and employment pools. Dividing by two standard deviations maintains coherence when comparing coefficients to binary variables. [Gelman and Hill, 2006].

All other continuous variables are transformed in a similar manner though where the mean and standard deviation are calculated from the entire data range and not only within each county. These transformation have the added benefit of aiding the convergence of the MCMC algorithm.[Gelman et al., 2013]. Binary variables are not transformed.

$$log\_empl_{c,t} = \frac{log(employment_{c,t}) - mean(log(employment_{c,t}))}{2 * std(log(employment_{c,t}))} \tag{1}$$

The likelihood of each response is modelled as a normal random variable with mean $\widehat{y_{c,t}}$ and standard deviation σ^y.

$$log(empl_{c,t}) \sim normal(\widehat{y^e_{c,t}}, sigma^{y,e}) \tag{2}$$

Analogously, for the wage effects of wind power investment, the response variable is log of wages in county c at time t, $log(wages_{c,t})$. The likelihood of each response is modelled as a normal random variable with mean $\widehat{y_{c,t}}$ and standard deviation $sigma^y$.

$$log(wages_{c,t}) \sim normal(\widehat{y^w_{c,t}}, sigma^{y,w}) \tag{3}$$

The fitted response data, $\widehat{y_{c,t}}$ is in turn modelled simultaneously by two hierarchical equations. To attain the final specifications for the models of employment and wages, I followed a process of estimating the models and checking the fit, making particular use of posterior predictive checks [Gelman et al., 2013] to be discussed further below, and Watanabe-Akaiki Information Criterion (WAIC).

Equation 4 describes the model at the observation level in a log-linear form. The fitted values for log employment are modelled as an intercept term, α_c, a stochastic trend θ_c and four time-varying covariates with corresponding parameters β^i_c where $i \in \{0, 1, 2, 3\}$. As the c indexing indicates, all of these variables are allowed to vary by county. A vector of quarterly dummy variables, $\mathbf{quarter_t}$ are included to control for seasonality, as rural counties tend to have a high proportion of seasonal workers.

The covariates include an indicator for the period, $period_{c,t}$ in quarterly intervals. The estimated parameter $\beta_c^{0,e}$ then represents a linear time trend on the log employment. The variable $capacity\_addition_{c,t}$ indicates the amount of capacity that was installed in a given county at time t. The parameter $\beta_c^{1,e}$ represents the immediate but temporary effect on employment of an investment. Because the data indicates when an investment is completed, A forward lag, $capacity\_addition_{c,t+1}$ is also included in order to capture employment effects for work-in-progress. Additional forward lags are not included as they can not be shown to add predictive power to the model.

The variable $capacity_{c,t}$ indicates the total wind power capacity in county, c in period t. The parameter $\beta_c^{3,e}$ is then an indication of the permanent employment effects of a wind power investment. Figure 4 shows a simplified illustrative diagram of the model for the log of employment over time with a wind power investment at time t. For the sake of simplicity, the diagram does not show potential effects of a forward lag parameter nor a stochastic trend.

$$\widehat{y_{c,t}^e} = \alpha_c^e + \theta y_{c,t-1}^e + \beta_c^{0,e} period_{c,t} + \beta_c^{1,e} capacity\_addition_{c,t} + \beta_c^{2,e} capacity\_addition_{c,t+1} \\ + \beta_c^{3,e} capacity_{c,t} + \zeta^{\mathbf{q,e}} \mathbf{quarter_t} \quad (4)$$

For the model of wages, the specification for the observation level equation is similar to that of employment. The main difference is the absence of a stochastic trend term (equation 5). The log wage series have less variance within a county and a linear trend and seasonal dummies are sufficient to obtain a satisfactory fit.[9]

[9]The empirical observation that wages have less variance than employment is well known theoretically grounded in the macroeconomics literature.

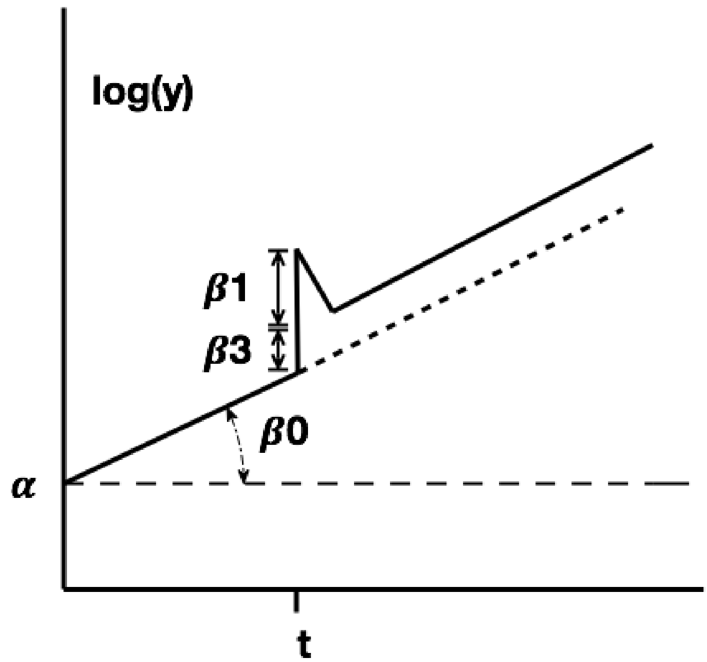

Figure 4: The diagram illustrates the observation level equations for each county. The diagram excludes illustrations of forward lags and a stochastic trend for simplicity.

$$\widehat{y_{c,t}^{w}} = \alpha_c^w + \beta_c^{0,w} period_{c,t} + \beta_c^{1,w} capacity\_addition_{c,t} + \beta_c^{2,w} capacity\_addition_{c,t+1} + \beta_c^{3,w} capacity_{c,t} + \zeta^{\mathbf{q,w}} \mathbf{quarter_t} \quad (5)$$

For both the wage and employment models, the c underscript in the parameters α_c and β_c^i signifies that these parameters are allowed to vary by county. Estimating this equation then involves estimated several parameters for each of 1140 rural counties plus the quarterly dummies.

This large pool of parameters is only meaningful if they themselves are modelled at a higher level by meta-parameters: A.) In order to avoid overweighting of outliers and false inference from multiple comparisons, and more importantly B.) To provide average inference and identification across counties. Equation 7 shows that the α_c parameters are modelled by an average intercept term,Γ^a and by two time-invariant county-level covariates, the county population, *population*$_c$ from the 2010 Census and the log of agricultural land value.[10] with corresponding parameters Φ^1 and Φ^2. A county random effect, α_c^{re} is included to model idiosyncratic county-level variation.

The β_c^i parameters are modelled simply as a pooled mean effect, Γ^i and an idiosyncratic county-level random effect $\beta_c^{i,re}$.

The hierarchical form of the model then allows each of the county-level coefficients and intercepts to be decomposed into a pooled average effect, as well as an idiosyncratic county-level random-effect. Such a "partial-pooling" model avoids undue influence by outliers by pulling them towards a grouped mean. Importantly they also allow for inference on average effect, while controlling for geographic variation and naturally taking into account issues of multiple comparisons through parameter shrinkage. For more in-depth discussions of mutlilevel models[11] and partial pooling I refer to Gelman and Hill [2006] and McElreath [2015].

$$\alpha_c = \Gamma^a + \Phi^1 population_c + \Phi^2 log\_agg\_land\_value_c + \alpha_c^{re} \tag{6}$$

$$\beta_c^i = \Gamma^i + \beta_c^{i,re} \qquad i \in \{0, 1, 2, 3\} \tag{7}$$

4 Model fitting with Bayesian MCMC

I use Bayesian Markov Chain Monte-Carlo (MCMC) simulation to fit the model using the Stan probabilistic programming language [Stan Development Team, 2014], which utilizes Hamiltonian MCMC (see MacKay [2003, ch. 30]) and a No-U-Turn Sampler [Homan and Gelman, 2014] for efficient sampling in high-dimensional probability space.

[10] Here invariant refers to the available data. Population and land values, of course, can and will change over time, however I do not have available quarterly or yearly estimates of either variable.

[11] Multilevel models are also referred to as random effects models, hierarchical models, and in the case of linear models: linear mixed models.

Weakly informative Cauchy priors[12] are assigned to the parameters as shown in equations 8. The corresponding variance terms, σ are themselves assigned half-Cauchy priors with with location parameter 0, and variance parameter of 5. Weakly informative priors have the effect of focusing the initial draws of the MCMC algorithm to reasonable values of the parameters, with the fatter tails of the Cauchy distribution, as opposed to a normal distribution, allowing for a non-negligible probability of outliers. The priors do not, however, impose any strong assumption of prior information on the model results. Use of the Cauchy prior distribution also allows for inference in the case of complete separation by covariates [Gelman, 2006].

$$
\begin{aligned}
\Gamma^i &\sim cauchy(0, \sigma^i) & i &\in \{0,1,2,3,4\} \\
\Gamma^\alpha &\sim cauchy(0, \sigma^\alpha) & i &\in \{0,1,2,3,4\} \\
\phi^j &\sim cauchy(0, \sigma^\phi) & j &\in \{0,1\} \\
\alpha_c^{re} &\sim cauchy(0, \sigma^{\alpha^{re}}) & & \\
\beta_c^k &\sim cauchy(0, \sigma^{qtr}) & k &\in \{5,6,7,8\} \\
\beta_c^{l,re} &\sim cauchy(0, \sigma^{re}) & & \\
\sigma^m &\sim \textit{half-cauchy}(0, 5) & m &\in \{0-4, y, \alpha, \alpha^{re}, re, \phi, qtr\}
\end{aligned} \tag{8}
$$

The Hamiltonian MCMC routine was run with four chains and 1000 iterations. R-statistics of 1 indicated convergence of the simulation to the target probability [Gelman et al., 2013].

[12]The Cauchy distribution is a t-distribution with 1 degree of freedom.

5 The effect of wind power investment on Labor Market outcomes

Since parameters are treated as random variables in Bayesian analysis, I present coefficients as distributions in the form of histograms and percentiles. Presenting the 6000-plus β and α distributions is space prohibitive. More so, the parameters of interest are the meta-parameters that provide average inference across counties. In particular, the Γ parameters provide inference for the variables of interest: wind power investment and capacity.

5.1 The effect of wind power investment on employment

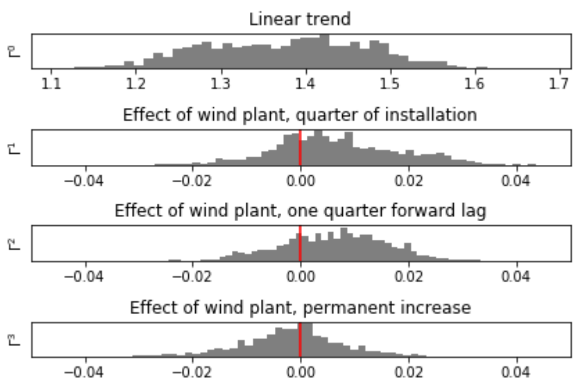

Figure 5: Estimated distribution of $\Gamma^{i,e}$ parameters from top to bottom panel: Cross county average of ($i = 0$) trend, ($i = 1$) temporary effect in quarter of first operation ($i = 2$) temporary effect, 1-quarter forward lag, ($i = 3$) permanent effect. No significant measurable evidence of temporary nor permanent effect on employment.

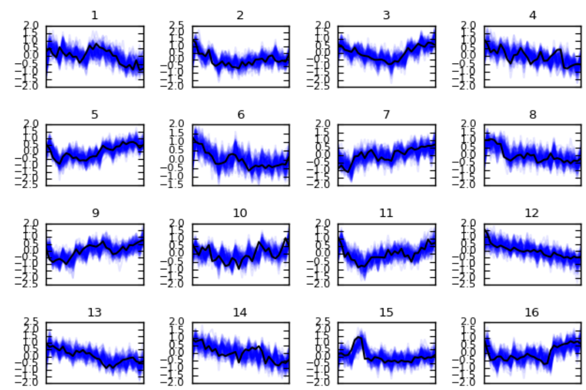

Figure 6: Posterior predictive check of the rural employment model. The blue lines are simulations from the probability model from a selection of rural counties. The black lines are the actual series. The probability model has a good fit to the realised time series.

Figure 5 shows summary histograms of the estimated distributions over the $\Gamma^{i,e}$ parameters for the regression on county employment. Table 2 shows percentiles for the $\Gamma^{i,e}$ as well as other meta-parameters. $\Gamma_{0,e}$ is the variable on the period variable, and thus represents the trend of log employment. A positive coefficient then represents a positive average growth rate of employment across rural counties.

Percentile	2.5	15	50	85	97.5
$\Gamma^{0,e}$	1.192	1.265	1.382	1.483	1.552
$\Gamma^{1,e}$	-0.015	-0.004	0.005	0.020	0.030
$\Gamma^{2,e}$	-0.013	-0.004	0.007	0.016	0.024
$\Gamma^{3,e}$	-0.025	-0.011	-0.001	0.008	0.018
$\Gamma^{a,e}$	-0.218	-0.123	0.029	0.202	0.347
$\Phi^{1,e}$	-0.007	-0.003	-0.000	0.003	0.006
$\Phi^{2,e}$	-0.007	-0.003	-0.000	0.002	0.006

Table 2: Percentiles of main parameters for model of rural employment.

Of more interest, the distributions of $\Gamma^{1,e}$ and $\Gamma^{2,e}$ parameters, which represent the immediate average effect on employment of a wind power investment and a forward lag of one quarter respectively, have distributions centered around zero. This can be interpreted to mean that there is no evidence for a immediate temporary positive effect on local employment from the initial build-out of a wind power farm. Arguably, construction of a wind farm could have an effect on employment more than 1 quarter before project completion, however specifications with additional forward lags also failed to show any significant temporary effect and did not increase the explanatory power of the model.

The distribution on the $\Gamma^{3,e}$ parameter represents the permanent effect of a wind power plant investment on employment. This distribution is also found to be centered around zero.

Figure 6 shows the posterior predictive check of the model specification for a sample of counties. The blue lines represent 100 draws from the posterior distribution over the in-sample data, giving an impression of the uncertainty in the model. The black lines represent the actual realised data on log employment for the counties. Visually, the model gives adequate coverage for the realised data. Watanabe-Akaiki Information Criterion (WAIC) were also used to compare model specifications.

5.2 The effect of wind power investment on wages

Figure 7 gives an overview of the estimated distributions over the $\Gamma^{i,w}$ parameters for the model of rural wages. The $\Gamma^{0,w}$ parameter, which represents the average trend line of wage growth in rural counties, shows a distribution over positive values. Interestingly, comparing the average estimate wage growth rate to that of metro counties, which can be found in the appendix, rural wage growth has actually been higher than in metro areas. However, average wages were at an absolute lower level and had fallen more during the preceding recession in rural counties.

The $\Gamma^{1,w}$ and $\Gamma^{2,w}$ parameter distributions, which represent the immediate temporary effect of a new wind power plant on wages from construction and planning activities, are again centered around zero. The construction of wind farms do not appear to have a measurable effect on rural county wages.

However, the distribution of the $\Gamma^{3,w}$ parameter, which represents the permanent effect of a wind power plant on wages, has a median value of 0.055 with a 95% confidence interval between 0.003 and 0.107. Interpreting from the mean value of wages across counties, this means that a mid-to-large sized wind farm with capacity of 200 MW[13] in a rural county is estimated at the median of the posterior to permanently raise wages by approximately 2.5 percent, an economically significant increase.

Figure 8 again shows the posterior predictive check of the model specification for a sample of counties. The linear trend in the model appears to be adequate for the realised wage series in most counties.

A testable hypothesis and robustness check for the result that wind power increases average wages in rural counties, is that the same will not be true in metro counties. In metro counties, with varied industries and a large employment pool, even a large wind power plant will be expected to have a negligible impact on average wages. Therefor, if a significant

[13] A wind farm consisting of approximately 60-70 modern turbines would have a total capacity of approximately 200 MW

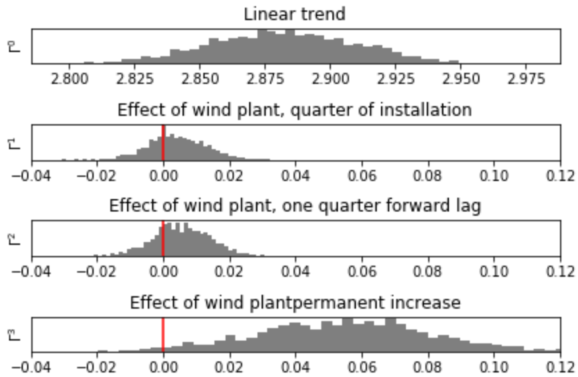

Figure 7: Estimated distribution of $\Gamma^{i,w}$ parameters: Cross county average of ($i = 0$) trend, ($i = 1$) temporary effect in quarter of first operation ($i = 2$) temporary effect, 1-quarter forward lag ($i = 3$) permanent effect. No measurable evidence of temporary effect on wages, but significant permanent effect on wages.

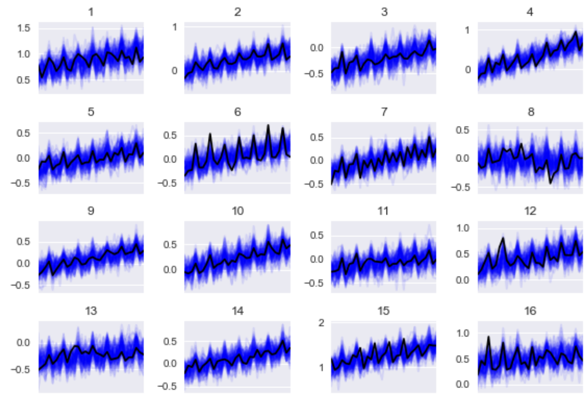

Figure 8: Posterior predictive check of the rural county wages model. The blue lines are simulations from the probability model from a selection of rural counties. The black lines are the actual series. The probability model has a good fit to the realised time series.

positive coefficient is estimated also for metro counties, then an unobserved variable is likely being confounded with wind power investment in the estimation of the $\Gamma^{3,w}$ parameter. Results from a full wage model with all counties can be found in section A in the appendix. I do not find a significant effect of wind power investment on wages in either metro counties nor counties adjacent to metros.

A presentation of all the individual county coefficients is space prohibitive, but figure 9 presents a summary of the estimated $\beta_c^{3,w}$ coefficients–representing the permanent effect of a

	2.5	15	50	85	97.5
$\Gamma^{0,w}$	2.827	2.853	2.882	2.913	2.936
$\Gamma^{1,w}$	-0.012	-0.004	0.004	0.014	0.022
$\Gamma^{2,w}$	-0.013	-0.003	0.005	0.014	0.022
$\Gamma^{3,w}$	0.003	0.027	0.055	0.082	0.107
$\Gamma^{a,w}$	-0.192	-0.135	-0.106	-0.043	0.030
$\Phi^{1,w}$	0.267	0.285	0.309	0.331	0.351
$\Phi^{2,w}$	-0.138	-0.117	-0.096	-0.071	-0.049

Table 3: Summary distribution of main parameters for model of rural employment.

wind power investment on wages–for the 120 rural counties that experienced a wind power investment in the period studied. The black vertical lines in panel A show 95% credible intervals of the $\beta_c^{3,w}$ parameters ordered in descending order by the absolute increase in wind power capacity in the county. For reference, the estimated median value of the meta-parameter, $\Gamma^{3,w} = .055$ is also shown. The partial-pooling property of the hierarchical model has the effect of shrinking the county $\beta_c^{3,w}$ parameters towards the mean $\Gamma^{3,w}$ parameter. As discussed earlier, this provides both a natural method of obtaining average, cross-county inference that takes into account the clustering of the data, as well as a addressing the multiple-comparison problem when estimating many parameters.

Because of the transformations of the data, the estimated parameters do not provide a intuitive interpretation. Therefor, for each county, c, and corresponding parameter, $\beta_c^{3,w}$, I calculate a county total average effect, $effect_c$ on wages of the total wind power capacity in the county, cap_c. These estimates can be interpreted as the permanent effect on wages of the installed wind power capacity in the county in the form of a ratio. For example, a county with a calculated total effect ratio of 1.05 and 100 MW of wind power, can be interpreted to mean that 100 MW of wind power capacity has increased wages by 5% compared to the counter-factual of no wind power invested.

The formula for this calculation is shown in equation 9. Here $f(\beta_c)$ represents some function of the distribution of β_c, for example, the median value or another percentile. I then multiply by $2\widehat{\sigma^w}$ to adjust to the normalization of the log wage data in the model. I multiply by the standardised capacity, st_cap_c[14] of wind power in the county in MW units. Finally, I take the exponent in order to interpret the results as ratios. Notice, that the $effect_c$ is then an out-of-sample prediction as the parameters β_c^3 were estimated based on the increase in wind power capacity within a county in the period studied not the total amount of wind power.

Panel B shows these calculated total effects in the form of median point estimates with

[14] recall that to standardise the variable by demeaning and dividing by 2 times the standard deviation

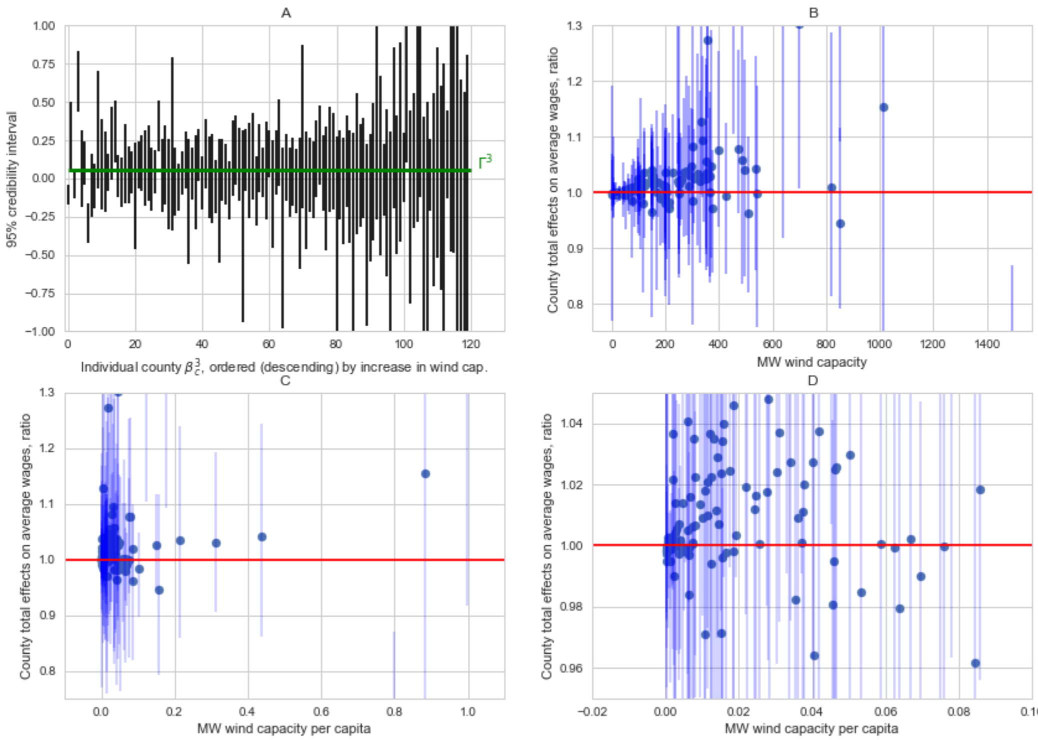

Figure 9: Panel A shows the individual county coefficients, $\beta^{3,w}$, for the permanent effect of wind power capacity additions on rural wages in the form of 95% credible intervals. The counties are ordered from most wind power to least wind power capacity additions. Panels B, C and D show a out-of-sample, total effect estimate for the 120 counties that experienced added investment in wind power capacity. The point estimates are median values of the estimated distributions and the bands represent 95% credible intervals. The estimates represent, in the form of a ratio, the estimated permanent effect on wages in each county of the total amount of wind power capacity installed. The estimates are in the form of a ratios, where 1 is no effect. In Panel B the estimates are plotted against total MW of installed wind power. In C and D the estimates are plotted against MW per capita. Panel D is scaled to better see the cluster of points near 0. The total effect of wind power investment on wages is in most counties modest, though not insignificant: below 5%. Some counties are estimated to experience substantially larger total effects, above 10%.

95% credibility bands, ordered by the size of the wind power capacity in the county. Panels C and D show the calculated total effects ordered by wind power capacity per capita, where panel D zooms-in on the cluster between 0 and .10 MW wind power capacity per capita. From the figure, it is apparent that most of the magnitudes are modest, with effect sizes below 5%. Some effect sizes are estimated to be substantially larger, above 10%, driven by a large wind farm capacity in the county and a relatively high median β_c^3 parameter estimate.

$$effect_c = exp(f(\beta_c)2\widehat{\sigma^w}st\_cap_c) \tag{9}$$

6 Validation by Simulated Data

In this section, I validate the model by simulating data with known parameters and distributions, and then estimating the model with the simulated data. I also compare the results of the Bayesian multilevel model with a commonly used fixed effects model with time and county fixed effects estimated with maximum likelihood. This is a technical section, and readers uninterested in the technical estimation of the results can skip to the next, concluding section.

6.1 Simulation of investment occurrence and magnitude.

The investment decision can be modelled as a sequential random process.[15] First, whether an investment in a certain county in a certain period is made can be modelled as a poisson process with mean, μ. Thus investment, I, is a random variables where the realised investment decision can be written $i_{c,t}$, representing how many investments where made in count c in period t. Following the definition of a poisson process, the probability mass function can be written as in equation 10, where λ is the average number of events per county-period.

$$P(i_{c,t}) \sim exp(-\lambda)\frac{\lambda^i}{i!} \tag{10}$$

Once a decision to invest is made, then a magnitude is chosen. I choose to model the magnitudes as a log-normal distribution. The log-normal distribution has a positive range[16]

[15] I refer to random here in the sense that the investment decision is exogenous to the observable variables, and therefor can be seen as a randomly assigned treatment effect.

[16] While still allowing the the distribution to integrate to one.

and has a rightward skew, fitting the pattern of the data that shows mostly small- to mid-sized wind farms but with a handful of very large wind farms. The PMF can be written as in equation

$$f(m,s) = \frac{1}{sm\sqrt{2\pi}} exp(-\frac{1}{2}(\frac{log(m)}{s})^2) \tag{11}$$

$$m = \frac{x - loc}{scale} \tag{12}$$

Here the random variable M, magnitude, is parameterised such that, $m = \frac{x-shift}{exp(\mu)}$ and $s = \sigma$, where μ and σ are from a standard normal distribution X, such that $exp(X) = M$.

I assign the following parameters to the log-normal distribution: $\mu = 20$, $\sigma = 1$, and $shift = 10$. The generated distribution is shown in figure 10. The peak of the distribution is at around 25, but with a long positive tail representing a few wind farms with large capacity.

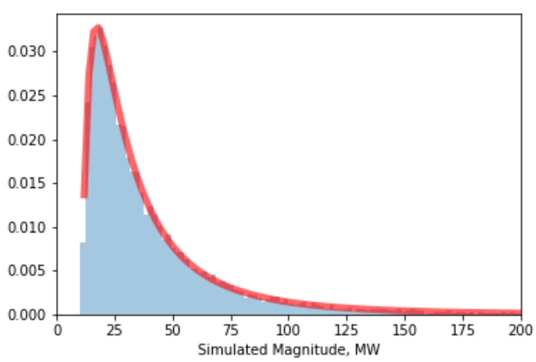

Figure 10: Distribution of simulated magnitude values for wind power plants.

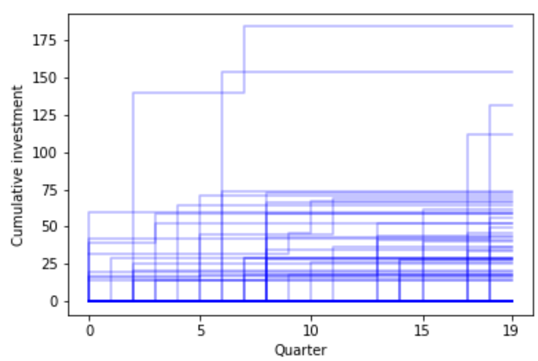

Figure 11: Simulated wind power capacity over time for each of 100 counties.

Combining the investment decision generation with the magnitude generation, I create a QxC matrix, **I** representing Q periods for each of C counties, and where every non-zero entry represents a wind power investment of size m. Take the cumulative sums across periods for each county, I create a QxC matrix **CI** representing the cumulative investment for each county C in period Q. Figure 11 shows an instance of the simulated cumulative capacity data.

6.2 Simulating hierarchical dependence

With the exogenous data simulated, I now create the dependent variable through a hierarchical structure. In particular, the QxC matrix $\mathbf{W}$ representing the log wage for each county C in period Q is calculated as in equation 13.

$$\mathbf{W} = \alpha + \underbrace{\mathbf{Q} \otimes \beta^{\mathbf{0}}}_{\textbf{trend}} + \underbrace{\mathbf{I^T} \circ \beta^{\mathbf{1}}}_{\textbf{temporary}} + \underbrace{\mathbf{CI^T} \circ \beta^{\mathbf{3}}}_{\textbf{permanent}} \tag{13}$$

The vectors of coefficients α, $\beta^{\mathbf{0}}$, $\beta^{\mathbf{1}}$, and $\beta^{\mathbf{3}}$, which represent the county intercepts, time trend, temporary effect of investment and permanent effect of investment are generated as a composite of an overall average effect, Γ^i_t and individual county random effects, $\alpha^r e_c$, $\beta^{i,re}_c$. The Γ coefficients are all distributed as log normal, which give them a positive range. Referring back to the parameterisation in equation 12, and letting y be realisations random variable Y, then the Γ distributions are generated as follows:

$$\Gamma^\alpha \sim lognormal(s = 1, loc = 5, scale = .05) \tag{14}$$

$$\Gamma^0 \sim lognormal(s = 1, loc = .02, scale = .01) \tag{15}$$

$$\Gamma^1 \sim lognormal(s = .5, loc = .001, scale = .0005) \tag{16}$$

$$\Gamma^3 \sim lognormal(s = .5, loc = .001, scale = .0005) \tag{17}$$

Histograms of 100 draws from each of the distributions is shown in figure 12

The random effect components are all generated as normal random variables with mean 0 and standard deviation .1, .005, .001, .001 for α^{re} and the $\beta^{i,re}$ distributions where $i \in 0, 1, 3$. The frequencies of a 100 draws from the α and β coefficients is shown in figure 13

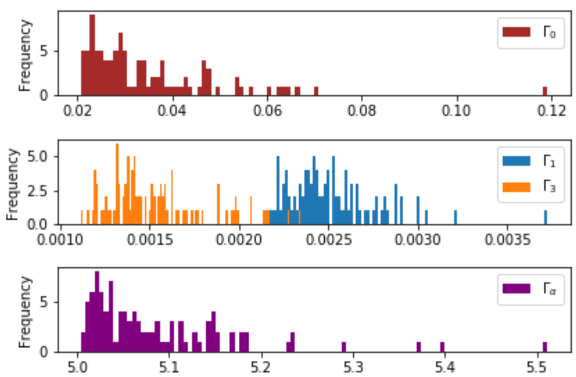

Figure 12: Simulated draws from Γ coefficients

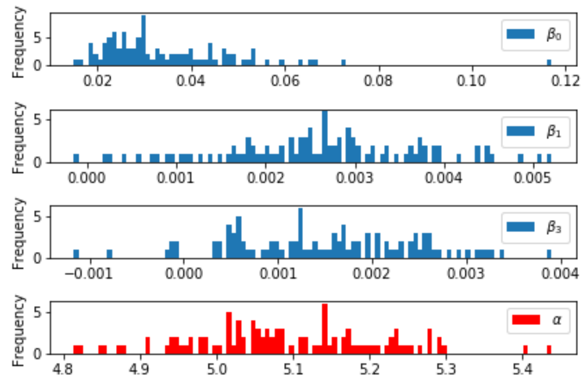

Figure 13: distribution of 100 draws from composite α and β coefficients.

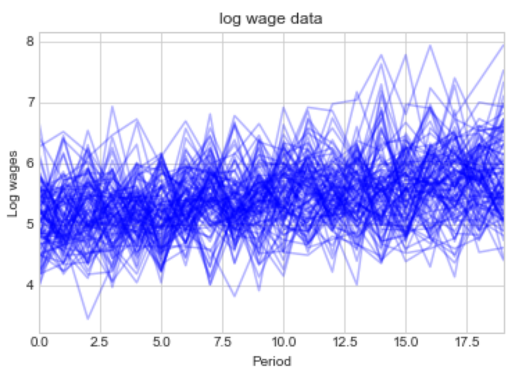

Figure 14: Simulated log wage data for 100 counties over 20 periods.

The resulting simulated log wage data is shown figure 14

6.3 Results from simulated data

Now that the simulated data is generated I use the multilevel model, as detailed above to estimate parameters from the generated data. I also compare results to a standard fixed effects model estimated with maximum likelihood with both time and county fixed effects and where the standard errors are corrected for clustering. This is a common model for estimating panel data. Here I do a Monte-Carlo style study: regenerating the random effects components of the data 100 times and then estimating the point estimates and standard errors of the parameters.

Figure 15 shows the results of the Monte-Carlo study of the fixed effects model. The parameter $\Gamma^{F,1}$ represents the estimated parameter on the investment variable, thus representing the immediate temporary effect of investment, while $\Gamma^{F,3}$ represents the estimated parameter on the cumulative capacity, representing the permanent effect of the investment. The results are scaled to reflect a 100 MW investment.

Comparing the results, which are presented with scale associated with the effect of a 100 MW investment, to figure 12, the point estimates appear to coincide fairly well with the actual Γ distributions. However, in a hypothesis testing frameworks, in many of the simulated regressions, the point estimates would not be rejected as significantly different from zero, as shown by the histograms of the positive and negative cut-off points of a 95% confidence interval, shown in grey. In some ways, this is not surprising. This reflects the conservatism of the cluster-adjusted standard errors of the fixed effects model.

We can compare with the results of the multilevel model, shown in figure 16. The results capture to a close approximation the magnitude and variance of the distribution of the gamma coefficients.

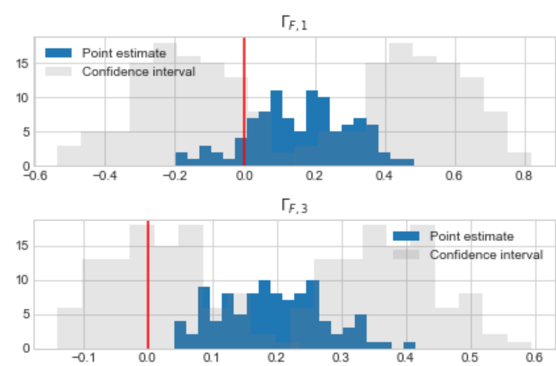

Figure 15: Results of monte-carlo simulation of a fixed-effects model. Coefficients are scaled to reflect effects from a 100 MW investment

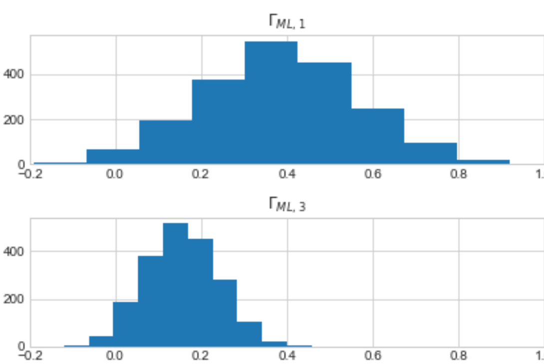

Figure 16: Results of multilevel estimation from simulated data. Coefficients are scaled to reflect effects from a 100 MW investment

Some discussion and perspective of the simulation results is warranted. The main purpose of this section was to validate the results from the multilevel model: To show that the model is capable of sensibly estimating the distribution of the parameters of interest. The comparison to the fixed effects model is meant purely as a reference point. A rigorous methodological comparison of fixed effects models with multilevel models is well beyond the scope of this article. I can however refer to several excellent discussions of this subject. Gelman and Hill [2006], McElreath [2015] and Kruschke [2014] recommend using multilevel models when modelling data with natural groupings and hierarchies and provide extensive explanation and evidence for this recommendation.

Another important point about this simulation is that the generated data is simplified. I do not, for example model in observed or unobserved county variables that could be correlated with investment, and which my empirical model takes into account. Nor do I model between-county differences in variance, which is also allowed for in my empirical specification. The more restrictive fixed-effects estimation may be expected to perform worse under these assumptions.

7 Discussion and Speculation

In summary, I find that wind power investments in rural counties have no measurable effect on employment, but a positive and economically significant permanent effect on wages. A mid-sized wind farm located in a rural county raises wages by 2.5% on average.

I argue that the model setup provides adequate identification of the causal effect of wind power investment on employment outcomes. This identification comes partly from the exogenous nature of wind power investments, which are heavily dependent on average wind speeds of a location. In addition, a panel data set with a multilevel model setup provides identification in the presence of unobserved variables correlated with the probability of wind power investment. The identifying assumption being that these variables are time-invariant.

However, even if there do exist unobserved variables that are correlated with the probability of a wind power investment and which are time-varying, there is reason to believe that identification will still hold. The reason is that the timing between when an investment decision is made and when a wind plant is built out has high variance and can be seen as an instrument inserting extra randomness into the treatment variable. Consider a hypothetical two counties which, due to an unobserved time-varying covariate–say, a change in local governments–both experience a commitment to build out a wind farm at the same time. This covariate could also independently affect employment and wages. A wind farm is build out in one county after a year, while it takes two years to build out the wind farm in the other county due to extra environmental concerns. So while the time-varying covariate–the change in local government–confounds the decision to invest, the effect of the actual realised investment is still identified by the multilevel model.

While I believe that the model is well identified, the underlying mechanism of how wind power investment affects average wages but not employment in rural counties is not clear. Several plausible mechanisms are consistent with the results.

One potential mechanism is that the skilled workforce required to install and maintain a wind turbine is difficult to find in rural counties, and that therefor the workforce commutes into the county for occasional maintenance work. The effect on wages could then be explained through increased wealth in the county that accumulates due to lease payments or ownership stakes from the wind turbines that in turn affects earnings.

An alternative explanation is that the addition of skilled positions maintaining and operating wind turbines leads to a transfer of already employed, skilled labor from one position in the county to another, leaving a vacant position in another industry in its wake. If a skilled workforce is the constraining factor, then rising wages but stagnant overall employment would be the expected result.

The research question and results in this article are interesting in their own right. Investments in energy generation and the related effects on labor markets are relevant to current public policy debates. In fact, they even played a significant role in the narrative of the presidential election of 2016.[17]

However, the degree to which investments in wind power can be considered exogenous to the local labor market is rare. When making a industrial investment, most firms explicitly or implicitly take into account the local labor market as a major factor. Skilled work force, labor costs, and local demand for the product are important factors in the expected profitability of most industrial investment. Profitability for wind power is however overwhelmingly determined by the average wind speeds of a given location.

This provides the prospect of wind power investments serving as an exogenous shock or instrument, and therefor setting up a type of natural experiment for some of the most important questions in the economics of labor markets. One important topic has been the trend of labor market "polarisation" in the last four decades, where employment has increased for low-skilled work and high skilled work, but real wage growth in these two categories has diverged, with low-skilled work actually experiencing a sustained real wage decline [Autor, 2014, Autor and Dorn, 2013]. Semi-skilled employment, such as a turbine technician, has traditionally defined the middle class. This category of employment has however stagnated in terms of both number of jobs and wages. Whether this stagnation is due to trade, technology or lack of necessary skills has been a active research topic [Autor et al., 2015, Acemoglu et al., 2015]. This article only gives hints about this larger debate, though researchers making use of more detailed register and tax data, of which I do not have access, could extract more robust insights.

8 Software and Replication Resources

For the analysis, I use the scientific computing environment for python: Numpy, Scipy, IPython and Jupyter [Walt et al., 2011, Oliphant, 2007, Perez and Granger, 2007]. The package Pandas was used for cleaning, formatting and descriptive analysis of the data [Wes

[17] `http://www.washingtonpost.com/news/energy-environment/wp/2017/03/29/trump-promised-to-bring-back-coal-jobs-that-promise-will-not-be-kept-experts-say`

Mckinney, 2010]. Figures were created using the package matplotlib [Hunter, 2007]. The Bayesian multilevel model was coded and computed using the excellent Stan probabilistic programming language and engine [Stan Development Team, 2014]. All of these software packages are open source and freely available. My Stan code and cleaned dataset are available on my website[18]. Other code used for preparation of data and descriptive analysis is available upon request.

References

Daron Acemoglu, David Autor, David Dorn, Gordon H. Hanson, and Brendan Price. Import Competition and the Great US Employment Sag of the 2000s. *Journal of Labor Economics*, 34(S1):S141–S198, December 2015. ISSN 0734-306X. doi: 10.1086/682384. URL `http://www.journals.uchicago.edu/doi/abs/10.1086/682384`.

D. H. Autor. Skills, education, and the rise of earnings inequality among the ”other 99 percent”. *Autor via Kate McNeill*, May 2014. ISSN 0036-8075. URL `http://dspace.mit.edu/handle/1721.1/96768`.

David Autor and David Dorn. The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5):1553–97, 2013. URL `http://econpapers.repec.org/article/aeaaecrev/v_3a103_3ay_3a2013_3ai_3a5_3ap_3a1553-97.htm`.

David H. Autor, David Dorn, and Gordon H. Hanson. Untangling Trade and Technology: Evidence from Local Labour Markets. *Econ J*, 125(584):621–646, May 2015. ISSN 1468-0297. doi: 10.1111/ecoj.12245. URL `http://onlinelibrary.wiley.com/doi/10.1111/ecoj.12245/abstract`.

Jason P. Brown. Production of Natural Gas From Shale in Local Economies: A Resource Blessing or Curse? *Economic Review - Federal Reserve Bank of Kansas City; Kansas City*, pages 5–33, 2014. ISSN 01612387. URL `http://search.proquest.com/docview/1532504720/abstract/E14B3AA81CF84058PQ/1`.

Anne Case and Angus Deaton. Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century. *Proceedings of the National Academy of Sciences*, pages 1–6, 2015.

Anne Case and Angus Deaton. Mortality and morbidity in the 21st century. *Brookings Papers on Economic Activity, Spring 2017, forthcoming.*, 2017.

Raj Chetty, David Grusky, Maximilian Hell, Nathaniel Hendren, Robert Manduca, and Jimmy Narang. The fading American dream: Trends in absolute income mobility since

[18] `http://jmaurit.github.io#wind_dev`

1940. *Science*, 356(6336):398–406, April 2017. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.aal4617. URL http://science.sciencemag.org/content/356/6336/398.

Thomas Ejdemo and Patrik Söderholm. Wind power, regional development and benefit-sharing: The case of Northern Sweden. *Renewable and Sustainable Energy Reviews*, 47:476–485, July 2015. ISSN 1364-0321. doi: 10.1016/j.rser.2015.03.082. URL http://www.sciencedirect.com/science/article/pii/S136403211500235X.

Andrew Gelman. Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Anal.*, 1(3):515–534, September 2006. ISSN 1936-0975, 1931-6690. doi: 10.1214/06-BA117A. URL http://projecteuclid.org/euclid.ba/1340371048.

Andrew Gelman and Jennifer Hill. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press, Cambridge ; New York, 1 edition edition, December 2006. ISBN 978-0-521-68689-1.

Andrew Gelman, John Carlin, Hal Stern, David Dunson, Aki Vehtari, and Donald Rubin. *Bayesian Data Analysis, Third Edition*. Chapman and Hall/CRC, Boca Raton, 3 edition edition, November 2013. ISBN 978-1-4398-4095-5.

Russell Gold. *The Boom: How Fracking Ignited the American Energy Revolution and Changed the World*. Simon & Schuster, reprint edition edition, April 2014.

Drew Haerer and Lincoln Pratson. Employment trends in the U.S. Electricity Sector, 2008–2012. *Energy Policy*, 82:85–98, July 2015. ISSN 0301-4215. doi: 10.1016/j.enpol.2015.03.006. URL http://www.sciencedirect.com/science/article/pii/S0301421515001123.

Matthew D. Homan and Andrew Gelman. The No-U-turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.*, 15(1):1593–1623, January 2014. ISSN 1532-4435. URL http://dl.acm.org/citation.cfm?id=2627435.2638586.

John D. Hunter. Matplotlib: A 2d Graphics Environment. *Computing in Science & Engineering*, 9(3):90–95, May 2007. ISSN 1521-9615. doi: 10.1109/MCSE.2007.55. URL http://aip.scitation.org/doi/abs/10.1109/MCSE.2007.55.

Timothy M. Komarek. Labor market dynamics and the unconventional natural gas boom: Evidence from the Marcellus region. *Resource and Energy Economics*, 45:1–17, August 2016. ISSN 0928-7655. doi: 10.1016/j.reseneeco.2016.03.004. URL http://www.sciencedirect.com/science/article/pii/S092876551630077X.

John K. Kruschke. *Doing Bayesian Data Analysis, Second Edition: A Tutorial with R, JAGS, and Stan*. Academic Press, Boston, 2 edition edition, November 2014. ISBN 978-0-12-405888-0.

Ulrike Lehr, Joachim Nitsch, Marlene Kratzat, Christian Lutz, and Dietmar Edler. Renewable energy and employment in Germany. *Energy Policy*, 36(1):108–117, January 2008.

ISSN 0301-4215. doi: 10.1016/j.enpol.2007.09.004. URL http://www.sciencedirect.com/science/article/pii/S0301421507003850.

Ulrike Lehr, Christian Lutz, and Dietmar Edler. Green jobs? Economic impacts of renewable energy in Germany. *Energy Policy*, 47:358–364, August 2012. ISSN 0301-4215. doi: 10.1016/j.enpol.2012.04.076. URL http://www.sciencedirect.com/science/article/pii/S0301421512003928.

David MacKay. *Sustainable Energy – without the hot air*. UIT Cambridge, May 2016.

David J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, Cambridge, UK ; New York, 1 edition edition, October 2003. ISBN 978-0-521-64298-9.

Johannes Mauritzen. Scrapping a Wind Turbine: Policy Changes, Scrapping Incentives and Why Wind Turbines in Good Locations Get Scrapped First. *The Energy journal*, 35 (2):157–181, 2014. ISSN 0195-6574. URL http://cat.inist.fr/?aModele=afficheN&cpsidt=28427728.

Richard McElreath. *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. Chapman and Hall/CRC, Boca Raton, 1 edition edition, December 2015. ISBN 978-1-4822-5344-3.

T. E. Oliphant. Python for Scientific Computing. *Computing in Science Engineering*, 9(3): 10–20, May 2007. ISSN 1521-9615. doi: 10.1109/MCSE.2007.58.

F. Perez and B. E. Granger. IPython: A System for Interactive Scientific Computing. *Computing in Science Engineering*, 9(3):21–29, May 2007. ISSN 1521-9615. doi: 10.1109/MCSE.2007.53.

Stan Development Team. *Stan Modeling Language Users Guide and Reference Manual, Version 2.2*. 2014. URL http://mc-stan.org/.

C. M. Tolbert and M. Sizer. U.S. COMMUTING ZONES AND LABOR MARKET AREAS: A 1990 UPDATE. September 1996. URL https://trid.trb.org/view.aspx?id=471923.

Jessica Trancik. Technology improvements and emissions reductions as mutually reinforcing efforts. Technical report, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, November 2015. URL http://trancik.scripts.mit.edu/home/wp-content/uploads/2015/11/Trancik_INDCReport.pdf.

USDA. Rural America at a Glance. Technical report, 2016.

Stéfan van der Walt, S. Chris Colbert, and Gaël Varoquaux. The NumPy Array: A Structure for Efficient Numerical Computation. *Computing in Science & Engineering*, 13(2):22–30, March 2011. ISSN 1521-9615. doi: 10.1109/MCSE.2011.37. URL http://aip.scitation.org/doi/abs/10.1109/MCSE.2011.37.

Jeremy G. Weber. The effects of a natural gas boom on employment and income in Colorado, Texas, and Wyoming. *Energy Economics*, 34(5):1580–1588, September 2012. ISSN 0140-9883. doi: 10.1016/j.eneco.2011.11.013. URL http://www.sciencedirect.com/science/article/pii/S0140988311002878.

Max Wei, Shana Patadia, and Daniel M. Kammen. Putting renewables and energy efficiency to work: How many jobs can the clean energy industry generate in the US? *Energy Policy*, 38(2):919–931, February 2010. ISSN 0301-4215. doi: 10.1016/j.enpol.2009.10.044. URL http://www.sciencedirect.com/science/article/pii/S0301421509007915.

Wes Mckinney. Data Structures for Statistical Computing in Python. *PROC. OF THE 9th PYTHON IN SCIENCE CONF. (SCIPY 2010)*, 2010. URL https://pdfs.semanticscholar.org/f6da/c1c52d3b07c993fe52513b8964f86e8fe381.pdf.

A Full model of wages

A robustness check to the result finding a significant effect of wind power investments in rural counties is to test the same model with the full data on all counties, including metro counties and counties adjacent to metro counties. For metro counties – with many industries and a large employment pool – even a large wind power investment plant should have a negligible effect on wages. We would not expect to see any significant effect of wind power investment on wages given that is what the model is capturing. Figure 17 provides a visual summary of the results. The estimated distributions for Γ_3, the parameter representing permanent increase in wages, is not significantly different from zero in metro counties and counties adjacent to metros.

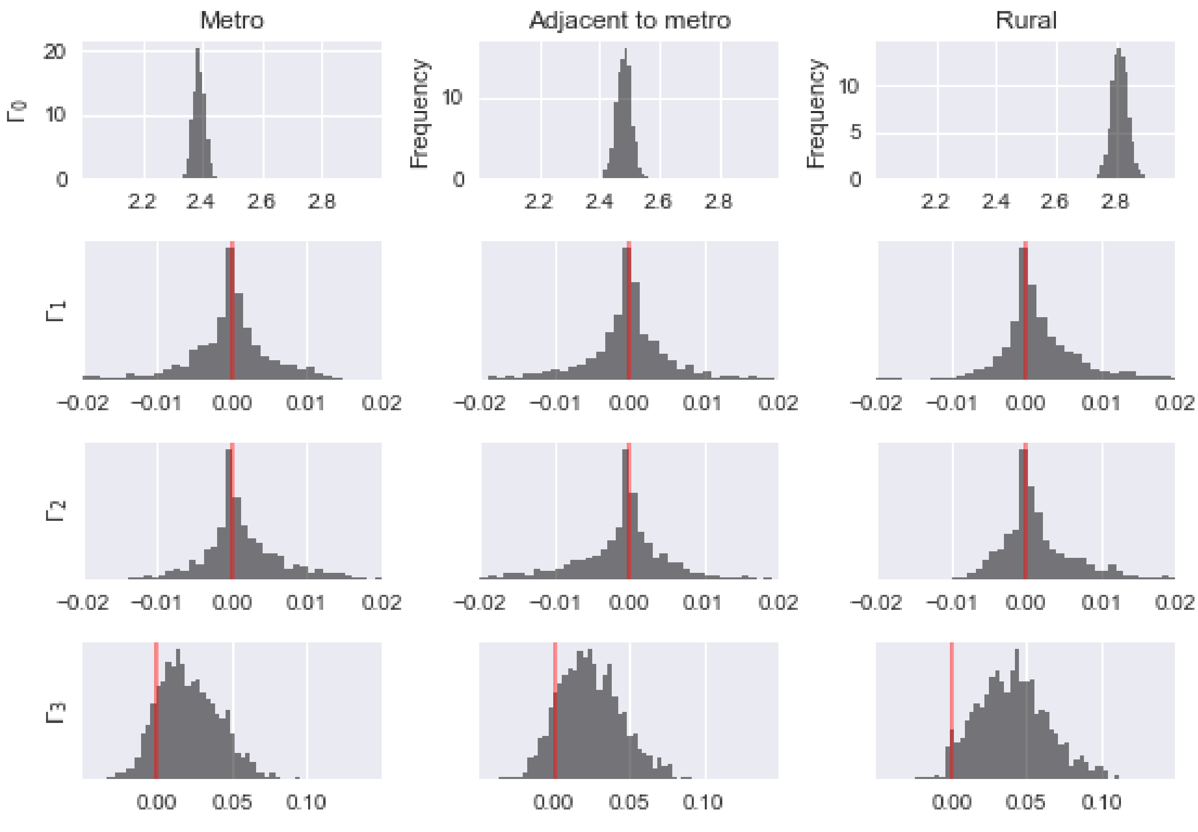

Figure 17: Summary distributions from full model of wages with all counties. Only rural counties show a statistically significant effect of wind power investment on wages.

param	var	2.5	15	50	85	97.5
Γ^0	Metro	2.353	2.368	2.389	2.409	2.427
	Small Metro	2.434	2.457	2.483	2.506	2.530
	Rural	2.759	2.785	2.813	2.841	2.869
Γ^1	Metro	-0.013	-0.005	-0.000	0.004	0.010
	Small Metro	-0.014	-0.005	-0.000	0.004	0.013
	Rural	-0.007	-0.002	0.001	0.006	0.015
Γ^2	Metro	-0.009	-0.003	0.000	0.006	0.013
	Small Metro	-0.016	-0.007	-0.000	0.004	0.011
	Rural	-0.006	-0.003	0.000	0.007	0.015
Γ^3	Metro	-0.014	0.000	0.019	0.043	0.061
	Small Metro	-0.014	0.002	0.022	0.046	0.070
	Rural	-0.001	0.015	0.040	0.065	0.092
Γ^a	Metro	-0.192	-0.176	-0.157	-0.134	-0.099
	Adj. Metro	-0.203	-0.183	-0.164	-0.140	-0.104
	Rural	-0.195	-0.178	-0.156	-0.131	-0.042
Φ^1	Metro	0.605	0.620	0.646	0.670	0.690
	Adj. Metro	0.313	0.354	0.404	0.454	0.486
	Rural	0.358	0.384	0.410	0.436	0.460
Φ^2	Metro	0.071	0.098	0.122	0.141	0.160
	Adj. Metro	-0.015	0.012	0.038	0.066	0.095
	Rural	-0.132	-0.119	-0.102	-0.079	-0.053

Table 4: Table from full model of wages with all counties. Wind power investments do not have statistically significant effects in metro areas and adjacent metros.

B Map of wind resources

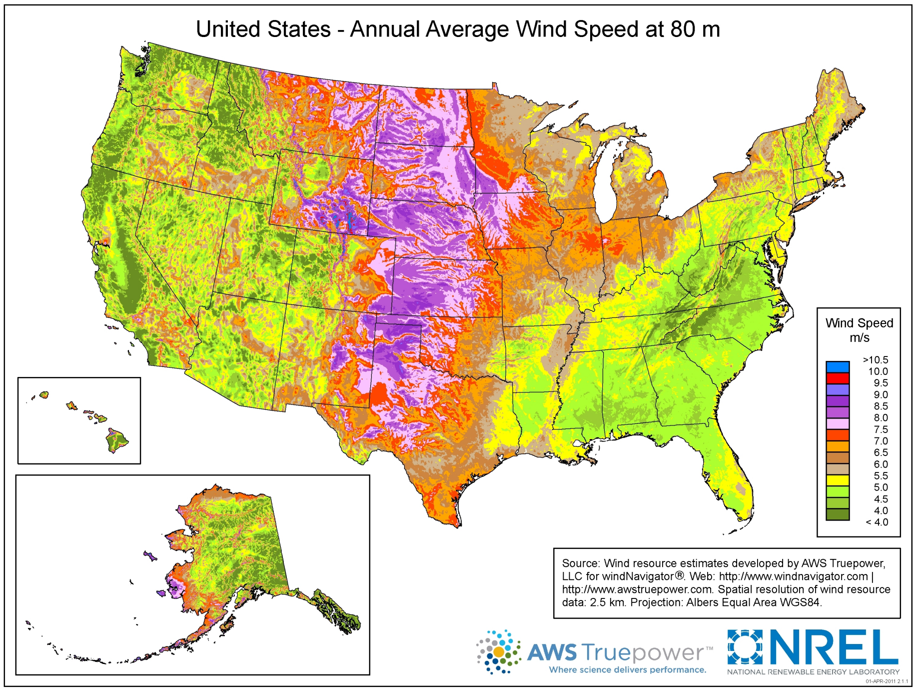

Figure 18: The National Renewable Energy Laboratory's wind resources map. Wind power investments, as shown in figure 1 are concentrated in the wind-rich spine of the US running from Texas up through North Dakota. Wind power investment decisions can to a certain extent be seen as exogenous