Are Renewable Generators Price Responsive?

Harrison Fell[*] and J Scott Holladay[†] and Daniel T. Kaffine[‡] [§]

Abstract

We examine the responsiveness of renewable generators to short-run electricity prices. Pairing high frequency generation data for individual renewable generators in Texas with spatially-detailed meteorological data, we find that wind-powered facilities increase their output in response to high day-ahead market prices conditional on local meteorological conditions and accounting for curtailment. By contrast, we find no evidence of price-response by solar-powered facilities when controlling for curtailment. Using instrumental variable methods to address endogeneity concerns, we find an average short-run price elasticity for wind generators to be on the order of 0.10 for wind-powered facilities. Price-responsiveness is greatest for unsubsidized wind farms and those with more market exposure, consistent with theoretical predictions. Our estimates reveal an overlooked intensive margin of operational adjustments by renewable firms. We illustrate the quantitative importance of this margin in the marginal value of public funds theoretical framework and in two simulation-based exercises: A shorter-term scenario where there is a sharp increase in demand but fixed generation capacity (e.g. recent data center build-out) and a longer-term scenario with high deployment of renewable generation capacity and reduced wholesale electricity prices.

[*]Department of Agricultural and Resource Economics, North Carolina State University; hfell@ncsu.edu

[†]Department of Economics, University of Tennessee; jhollad3@utk.edu

[‡]Department of Economics and Renewable and Sustainable Energy Institute, University of Colorado Boulder; daniel.kaffine@colorado.edu

[§]Disclosure: The authors did not receive any outside research support associated with this study. The authors have no significant financial relationships or other potential conflicts of interest that apply to this research.

1 Introduction

Many countries provide some form of subsidies for wind and solar generation, which has helped propel substantial growth in the deployment of these technologies over the last two decades. Likewise, lapses in these subsidies have been met with noticeable declines in wind and solar investment, such as in the U.S. with the uncertainty in the production subsidy for wind in 2013 and the recent repeal of Inflation Reduction Act (IRA) subsidies.[1] This pattern suggests a positive long-run or extensive margin (capacity deployment) price elasticity of supply for wind and solar generators.

Less attention has been paid to the short-run supply elasticity of these renewable technologies. As wind and solar generators are often infra-marginal producers in the electricity production supply curve, it is typically assumed the short-run price elasticity of supply for renewable generation along the "operating" margin is perfectly inelastic - renewable generation is simply a function of technology and meteorology (the type of wind turbine/solar panel, the number of wind turbines/solar panels, the realized wind speed/solar irradiance, etc). However, renewable operators are profit-maximizing entities that own and operate complicated capital equipment, and we would expect them to strive for as much production as possible when prices are favorable. In this paper, we utilize high frequency, wind and solar farm-level data to estimate the extent to which generation from renewable firms respond to daily price signals. In essence, when electricity prices are high, do renewable firms respond

[1] Following the expiration of the U.S. production tax credit for wind generation, wind capacity additions plummeted in 2013 to about 1 gigawatt (GW) after adding roughly 13 GW in 2012. The credit resumed in late 2014 and capacity additions rebounded quickly (Plumer 2014). Following the removal of clean energy production and investment subsidies under the U.S. "One Big Beautiful Bill", over \$24 billion in clean energy investment were withdrawn in 2025 compared to the announcement of \$11 billion in new clean energy projects over the same span (DiGangi 2025).

by increasing their wind and solar generation?

To answer this question, we use production data from individual wind and solar generation facilities in the Electric Reliability Corporation of Texas (ERCOT) market over the 2016 - 2024 time period. To control for local weather-driven changes in production, we derive generator-specific meteorological controls, including wind speed and solar irradiance measures, based on spatially gridded weather reanalysis data. We then pair this generator-specific data with regional electricity price data and other market data necessary to conduct various panel and instrumental variable estimation techniques.

Using daily-level data, we find robust evidence that production from wind generating units responds positively to increasing daily average electricity prices, implying a short-run supply elasticity on the order of 0.10.[2] Furthermore, as expected, we find wind generators that no longer receive the federal production subsidy are more responsive to electricity prices than those facilities still receiving the subsidy. Likewise, we find some evidence that wind generators with more exposure to electricity market prices are more price-responsive than those facilities with the majority of their production sold via power purchase agreements (PPAs). In contrast, once we account for curtailment, we find no evidence of short-run price-responsive supply from solar generators.

To assess the quantitative implications of the positive short-run price elasticity of supply for wind generators, we first calculate the additional benefits from this intensive margin response in the marginal value of public funds framework (Hahn et al. 2026). We find an

[2] Our instrumental variable approach leveraging Henry Hub natural gas prices addresses the the key identification challenge arising from the simultaneity between electricity prices and aggregate renewable generation, which will bias OLS estimates negatively. We also address the potential for curtailment (i.e. the forced reduction of generation when local supply exceeds demand), and we find that production from wind generating units increases in response to higher prices when curtailment is absent.

additional \$4.4 billion in social benefit from the approximately \$10 billion in annual federal subsidy expenditure (assuming marginal effects scale linearly with spending levels). We then conduct two additional simulation-based exercises. In the first exercise, we build a least-cost dispatch simulation model for ERCOT to understand the short-run effects of data center load additions with and without price-responsive wind generators. We find the price-response by wind generators can mitigate the increase in prices and emissions due to new load additions, reducing costs of serving new load by tens of millions of dollars annually and reducing CO_2 impacts by several hundred thousand tons per year. In the second exercise, we calculate predicted changes in capacity factor from wind generators when firms face lower average prices but higher volatility, consistent with projected expansion of renewable generation shares. We find the price-response by wind generators can substantially reduce capacity factors, by as much as 2 percentage points.

This research contributes to the literature in several ways. First, supply curves and their elasticities are fundamental objects of interest in economics (Shea 1993), particularly for emerging and disruptive technologies. More directly related to the energy sector, there have been several studies demonstrating how operating margins of fossil-fuel and nuclear power plants respond to changing incentives (e.g. Fabrizio et al. (2007), Cicala (2015), Davis and Hausman (2016), Chan et al. (2017), Doyle and Fell (2018)), but have not addressed similar responses from wind and solar generators. Instead, as noted above, most of the work looking at how renewables respond to prices or other incentives has focused on extensive margin responses via capacity additions (e.g. Fell and Linn (2013), Hollingsworth and Rudik (2019), Deschenes et al. (2023), Belyak et al. (2024)).

An exception to this extensive margin focus is recent work on the effect of production subsidies on generation from existing wind generators. Specifically, Aldy et al. (2023) find wind generators who claim the production subsidy are about 10% more productive than those that claim an investment subsidy and, similarly, Ricks and Kay (2025) find that facilities whose production subsidy has expired produce about 10% less than similar facilities still receiving the subsidy. This research suggests incentives do affect the production from renewable firms, and that there exist margins of adjustment by operators that can increase output. Here, we examine if this production response extends more generally to daily variation in price incentives and not just via significant, one-time changes in incentives from tax credits. Furthermore, our data allow us to directly address curtailment, such that supply elasticity estimates only reflect price-induced increases in output, and not changes in output due to curtailment effects.[3]

The results of the analysis are important from a policy perspective for several reasons. First, the supply-side elasticities for renewable energy (both short-run and long-run) are critical determinants of the costs of meeting the decarbonization and net-zero goals that many countries and local governments aspire to. Second, understanding how existing renewable generators respond to financial incentives is crucial for assessing policy-induced generation additionality — that is, whether policy-based incentives actually increase renewable output.[4] Third, as grid reliability concerns increase, it is important to understand how existing generators, including renewable generators, respond to production incentives in order to support

[3] Both Aldy et al. (2023) and Ricks and Kay (2025) use monthly generation data, which limits the ability to identify curtailed conditions for a generator. Both papers find suggestive evidence that their estimated output responses to prices are not solely due to curtailment. Our high-frequency generation and curtailment data at the generator level allows us to more definitively show that output is responding to price variation and provide additional confirmation that their results are not driven by curtailment.

[4] Concerns that such incentives do little to boost generation have fueled accusations of corporate "greenwashing" in REC markets and led policymakers to limit green hydrogen subsidies to hydrogen producers buying energy from newer renewable sources (Bjørn et al. 2022; Fell et al. 2025)

grid stability. Fourth, continued subsidies for near zero-marginal-cost renewables will almost certainly increase their overall capacity, which in turn will likely reduce wholesale market prices. However, to the extent that production margins are responsive to wholesale prices, such investment subsidies cannibalize production from existing generators beyond the observed increased curtailment rates (see for example Novan and Wang (2024)).

2 Conceptual Framework of Maintenance Decisions

Why might we expect renewable firms to respond to electricity prices in the short-run? Our starting point is that wind and solar generators are complicated capital equipment, and just like any other capital-intensive firm, there are operational choices and expenses to keep the system running optimally. When wholesale electricity prices are, say, $1000 per megawatthour (MWh), being able to coax an extra MWh of production from a wind turbine or solar panel is very lucrative, and we would expect firms to respond to those incentives accordingly.[5] That is, given that there exist costly margins of adjustment related to the operation and maintenance of wind and solar generators, firms will be more willing to take on those costs when wholesale electricity prices are high.[6]

To formalize these notions, let $c(m)$ represent the cost for a unit of maintenance m, such that $c'(m) > 0$ and $c''(m) > 0$. Maintenance activities increase the generation from the wind

[5] As noted by wind power service company Windurance: "When it comes to replacing components, a turbine owner has the choice to replace the entire component instead of repairing. In general, while a quick repair could be more expensive, it does ensure that your turbine is quickly up, running, and generating income." https://blog.windurance.com/wind-turbine-component-replacement-vs-wind-turbine-repair

[6] Shine-up Solar notes those costs and trade-offs: "If water is not available because of drought, distance or costs then we can dust the modules. Dusting takes about the same time in some cases and can be a better spend per dollar. Dusting doesn't get the panels quite as clean as water ..., but if water costs 20-30% more it may be worth it." https://shineupsolar.com/utility-solar-cleaning/

or solar farm.[7] To fix ideas, consider a wind farm with MWh generation g given by:

$$g = \lambda(m) * v^3, \tag{1}$$

where v is wind speed and $\lambda(m)$ is an engineering productivity constant that translates wind speeds into generation for a given level of m. The standard assumption is that $\lambda(m)$ is fixed, whereby the output from the wind generator is simply determined by engineering features and wind speed. Here, we assume the productivity constant can depend on the firm's choice of maintenance efforts, such that $\lambda'(m) > 0$ and $\lambda''(m) < 0$ - intuitively these maintenance efforts can be thought of as costly actions to increase the productivity per turbine or increase the number of functioning turbines, increasing output for a given wind speed level (to diminishing effect).[8] In other words, increasing generation g via maintenance activities is costly, but may be worth it depending on the potential revenues earned.

The policy and market environment will influence a firm's revenue and thus maintenance decisions. Many renewable firms contract their generation output through arrangements like Power Purchase Agreements (PPAs), where some (or all) of the firm's generation has a fixed-price offtaker at some price $\bar{p}$. Other firms sell some (or all) of their generation on the wholesale electricity market, at price p subject to the fluctuations of the market. Let $0 \leq \gamma \leq 1$ represent the share of generation sold via fixed-price contracts, such that $1 - \gamma$

[7] These maintenance activities can broadly be thought of as capturing a wide array of operational decisions that renewable firms can make. For example, it might entail scheduling an additional maintenance crew, rush-ordering certain parts to get a few more turbines or panels up and running, delaying repairs to a keep a turbine or panel on the maintenance "watch list" running an extra day or two (at the expense of potentially larger repair costs in the future), or just simply devoting some extra time, energy and attention to ensuring all turbines and panels are performing as best as possible.

[8] The cubic relationship between wind speeds and power generation holds over the range of wind speeds between 4 m/s to 11 m/s typically observed (Kaffine and Worley 2010). Below 4 m/s or so, no generation is produced and above 11 m/s, the unit produces at nameplate capacity. The term $\lambda(m)$ captures both engineering features reflecting generation per wind turbine as well as the number of operating wind turbines. It is straightforward to map the above setting to a solar farm, where the productivity of the solar farm for a given level of solar irradiance can be increased via costly maintenance efforts.

represents the exposure to the wholesale market.[9] Finally, wind generators also receive a Production Tax Credit (PTC) of τ per MWh of generation.[10]

Putting this together, the profit maximization problem for the firm can be expressed with respect to maintenance choice m as:

$$\max_m \pi = (1-\gamma)(p+\tau)\lambda(m)v^3 + \gamma(\bar{p}+\tau)\lambda(m)v^3 - c(m), \tag{2}$$

which has first-order conditions:

$$(1-\gamma)(p+\tau)\lambda'(m)v^3 + \gamma(\bar{p}+\tau)\lambda'(m)v^3 - c'(m) = 0. \tag{3}$$

The solution to the above, $m^*(p,\tau,\gamma)$, is the optimal maintenance effort by the renewable firm, which intuitively balances the additional return from increased productivity $\lambda'(m)$ against the increased costs $c'(m)$. How does that optimal maintenance effort vary with the wholesale electricity price? By the Implicit Function Theorem, we have that:

$$\frac{dm^*}{dp} = -\frac{(1-\gamma)\lambda'(m)v^3}{\lambda''(m)v^3(\tau+(1-\gamma)p+\gamma\bar{p}) - c''(m)} > 0, \tag{4}$$

which given our assumptions on $\lambda(m)$ and $c(m)$ can be signed as positive. Intuitively, increasing wholesale prices increases the return on maintenance efforts and the firm will do more of it, and thereby increase generation when prices are higher.

Equation 4 reveals two other key features. First, a firm with no wholesale market exposure ($\gamma = 1$) will intuitively not respond to changes in market prices ($\frac{dm^*}{dp} = 0$), while a firm with some market exposure ($\gamma < 1$) has some responsiveness to prices $\frac{dm^*}{dp} > 0$. For a fully exposed

[9] In practice, PPA contracts may contain a variety of obligations and terms. For example, a renewable firm may have quantity delivery targets, and failure to meet them may require purchase on the spot market. Such a firm would still face wholesale market exposure.

[10] The federal Production Tax Credit in the U.S. provides a production subsidy of approximately $24/MWh for wind turbines over the facility's first 10 year of operation.

firm ($\gamma = 0$), their response is $\frac{dm^*}{dp} = -\frac{\lambda'(m)v^3}{(p+\tau)\lambda''(m)v^3 - c''(m)}$. Second, increasing the PTC (τ) will decrease the responsiveness of the firm to wholesale prices (in the limit $\frac{dm^*}{dp} = 0$). A firm receiving the PTC is already incentivized to increase production (Aldy et al. 2023; Ricks and Kay 2025) and thus has less incentive to undertake additional costly efforts to increase production further, assuming diminishing returns.

From the above, three empirical predictions emerge: 1) Renewable generators will respond to higher wholesale market prices by increasing production (holding wind speed or solar irradiance constant); 2) Generators that are more contracted through Power Purchase Agreements will respond less to electricity prices; 3) Generators that are still receiving the PTC will be less responsive to electricity prices. We examine these predictions in the sections that follow.

3 Data and methods

Below we detail the data and methods used to empirically examine how wind and solar generation facilities respond to electricity prices in the ERCOT region. The ERCOT electricity market provides a good context to study this price-responsiveness issue for several reasons. First, ERCOT has a significant quantity of wind generation capacity (25% of annual generation in 2022) and increasing solar capacity (6% of annual generation in 2022) located throughout the market region, which allows for significant variation in meteorological conditions and local market conditions. Second, ERCOT's wind generation has been a significant presence for nearly two decades, which allows for many years of analysis and also allows us to estimate the price responsiveness of wind farms whose federal production subsidy expires after 10 years of operation. Third, the ERCOT market operators provide detailed, publicly-available data on the operations of individual generating units and general market conditions, which allow us to explore within-generator responses under a variety of empirical settings.[11]

3.1 Data

We collect data on generation, meteorological conditions, and electricity prices for all wind and solar generators in ERCOT. Our dependent variable of interest is the capacity factor (generation divided by capacity) from wind and solar generators, which we obtain for 2016 - 2024 from ERCOT through their Security Constrained Economic Dispatch (SCED) reports. This SCED data has 15-min production data by "resource" which we aggregate to the daily level in our primary specifications.[12] We create a cross-walk between SCED "resource code" identifiers and Energy Information Administration (EIA) Plant Code and Generating Unit identifiers, allowing us to pair EIA-860 facility information on capacity, latitude/longitude and generation start date with the SCED generation data. Finally, we collapse the SCED data to the EIA Plant Code-Generating Unit level, which smoothes out some of the arbitrary disaggregation within some solar/wind farms in the SCED data.

Given a generator's latitude and longitude, we match the generator to gridded meteorological data provided in the National Oceanic and Atmospheric Administration's North Ameri-

[11] Other data sets such as EIA-923 provide monthly generation data for individual generators across the US, while EIA-930 provides hourly generation data for aggregate generation types across the US. In contrast, ERCOT market data is high-frequency at the individual generator level.

[12] Our analysis is at the daily level, which is a plausible time scale for firms to respond to price incentives. That said, we do utilize the hourly data for various robustness checks and data restrictions (e.g. identifying hours of potential curtailment). In the appendix, we also include results from an hourly aggregation, which matches to the frequency of our meteorological data.

can Regional Reanalysis (NARR) project (see https://psl.noaa.gov/data/gridded/data.narr.html). NARR provides backcasted weather predictions for a variety of meteorological variables over a 32km grid of North America for eight, evenly spaced hourly observations each day from 1979 to the near present. We use a bilinear interpolation of this gridded data to create a weather observation for each generator and then interpolate the 8-times-daily observations to form hourly weather data for each generator over the 2016-2024 span. The meteorological variables serve as key controls for local generation and include u- and v-wind components at the 10m level (which allow us to form wind speed and direction variables), precipitation, temperature, barometric pressure, and downward shortwave radiation flux (DSWRF, a measure of solar irradiance).[13]

ERCOT also provides more aggregated regional and system data. For each generator, we use regional identifiers provided in the ERCOT SCED data for generating resources to assign hourly, regional electricity prices from the day-ahead market (DAM) to each generator.[14] ERCOT gives regional load (consumption) data for realized and predicted load, which, as described below, we use in our instrumental variable (IV) estimation approach.[15] Finally, our IV approach also uses natural gas prices, for which we use daily Henry Hub prices as

[13] For a solar irradiance control, we also use NREL's National Solar Radiation Database (NSRDB), which is a collection of satellite-derived measurements of solar radiation (global horizontal, direct normal, and diffuse horizontal irradiance) at the daily level on a 2km grid. This data is not available for the full 2016-2024 sample and results over the shorter time sample using the NSRDB measure of solar irradiance do not differ materially from the NARR-derived metric.

[14] ERCOT clears DAM and real-time market prices at a nodal level. These nodal prices are then aggregated and averaged to a hub-region level, of which there are five regions - Houston, North, South, West, and Panhandle. We match these regional hub prices for the DAM to the generating resource based on the resource's listed region in the SCED data.

[15] Realized and predicted load are given by eight ERCOT-defined "Weather Zone" regions (Coast, East, Far-west, North, North-central, Southern, South-central, and West). Based on the county-location of the generating resource and the maps of the weather zones provided by ERCOT (https://www.ercot.com/files/docs/2016/01/06/2015ercotconstraintsandneedsreport.pdf) we can match the generating resource to a weather-zone load value.

reported by the EIA (https://www.eia.gov/dnav/ng/hist/rngwhhdm.htm). See Table 1for wind and solar generator summary statistics.

3.2 Empirical strategy

Similar to Fell and Kaffine (2018), our unit of analysis is daily capacity factor by generating unit, and we adopt a similar specification in regressing daily capacity factors on daily price movements. Our base regression specification takes the following form:

$$\log(CF_{iht}) = \beta \log(P_{ht}) + \sum_{j=t-1}^{t} \gamma_j X_{ij} + \alpha_i + \eta_{my} + \delta_{dow} + \epsilon_{iht}, \tag{5}$$

where CF_{iht} is the capacity factor of generating unit i in hub h on day t and P_{ht} is the electricity price for hub h at t. The term X_{ij} represents a flexible vector of other control variables including levels and squared values of meteorological controls (wind speed for the analysis of wind generators and DSWRF for solar, along with temperature, pressure, and precipitation). To control for other unobservables at the system and generator level, we also include generating unit fixed effects α_i, month-by-year fixed effects η_{my}, and day-of-week fixed effects δ_{dow}. Finally, in the base specifications, standard errors are clustered at the plant code and day-of-sample level. Note, in our analysis we conduct separate methods of estimating 5 for wind and solar generators. We estimate the price responsiveness of these technologies separately as wind and solar generators are obviously quite different technologies with different potential adjustment margins.

The key coefficient of interest is β, representing the elasticity of generation with respect to price, identified off of within-firm variation in generation output within month-years, conditional on meteorological variables and other controls. Per our theoretical model, we expect the sign on β to be positive, as conditional on wind speed or solar irradiance and other control variables, renewable firms should respond to high price incentives by undertaking efforts to increase output. However, identification in this context can be difficult. The primary concern is somewhat of a mechanical simultaneity bias with aggregate renewable production. More specifically, as renewable generation is typically inframarginal (i.e. it effectively has zero marginal cost), an increase in generation from renewables effectively shifts the price-setting portion of the electricity supply curve outward and thus lowers equilibrium prices. This effect is particularly pronounced for aggregate wind generation given its scale relative to the total ERCOT market. On average, any given wind or solar generators' production is positively correlated with total generation from that technology type. Given this simultaneity issue, an OLS estimation of 5 would bias the estimate down.[16]

We take two approaches to addressing this bias issue. First, we employ an IV approach using demand shifters and supply shifters that are plausibly unrelated to renewable generation. For a supply shifter, we use the one- and two-day lag of the log of Henry Hub natural gas prices. Natural gas generators are typically the marginal (i.e. price-setting) technology in ERCOT, so the fuel price of that technology will greatly affect electricity prices. Henry Hub prices are the standard price metric for the U.S. natural gas market and are affected by gas market conditions across the country (and world). As such, they are not unduly influenced by activities in the ERCOT market alone. In addition, to further ensure there is no influence of contemporaneous ERCOT market conditions on Henry Hub prices, we use lagged prices as they are likely correlated with current prices and can additionally influence

[16] Note controlling for aggregate wind or solar generation in ERCOT is not a reasonable strategy as the interpretation of the parameter of interest in such a case is difficult at best. That is, the interpretation of the average responsiveness of a facility-level generation metric to prices while holding aggregate production from that technology constant holds no particular economic meaning.

the marginal cost of natural gas-powered generators via multi-day gas contracting arrangements. For a demand shifter, we adhere to the often-used assumption that the demand (or load) for electricity at the daily level is largely exogenous. However, grid disruptions (e.g. transmission/distribution line outages or power plant outages) may affect observed load and renewable generation. To decouple this process, we use load predictions for day t made 48 to 72 hours before t. Thus, in addition to the log of the lagged Henry Hub prices, we include the log of predicted loads for ERCOT in total and for the given generator's region.

The second approach we take to break this mechanical simultaneity bias is to replace $\log(P_{ht})$ with $\log(P_{ht-1})$ while also controlling for ERCOT day $t-1$ aggregate renewable generation from the given technology (W_{t-1} for wind and S_{t-1} for solar). This specification exploits the observation that daily average prices are correlated and thus $t-1$ electricity prices serve as a reasonable proxy for t prices. However, day t aggregate wind/solar generation is also highly correlated with $t-1$ wind/solar generation, so controlling for W_{t-1} or S_{t-1} breaks the mechanical simultaneity bias concern while still retaining a useful interpretation of the estimator as the supply elasticity. Another possible motivation for this specification is that if the actions taken by renewable operators to adjust output take time to complete, and if day $t-1$ DAM prices are a good predictor of day t prices, then the $t-1$ prices may in fact be the price signal that operators are responding to.

Curtailment can also affect the interpretation of β in equation 5. Curtailment of renewables, which is dictated by the ERCOT market operators, typically occurs in periods of relatively low electricity demand and (potentially) high renewable generation. Thus, curtailment would lead to both lower prices P_{ht} and lower capacity factors CF_{iht}, leading to a positive β estimate. However, this curtailment response is a distinct mechanism from our interest in price-induced increases in supply from operation and maintenance activities. We thus undertake two approaches to ensure the supply response we are estimating is not simply a curtailment phenomenon. First, very low and possibly even negative clearing prices are a sign that curtailment of renewable generators is almost certainly occurring. As such, in all specifications we exclude observations with $P_{ht} < 5$ unless noted otherwise.[17] Second, the ERCOT SCED data also reports an hourly High Sustained Limit (HSL) value, which is defined as a production limit continuously updated and reported by the generating facility "...that describes the maximum sustained energy production capability of the Resource" (Electric Reliability Council of Texas (2024)). If the observed generation is significantly below this HSL value, it indicates the resource is being curtailed. We thus consider sample restrictions where we exclude observations for a given generator for a given day if the number of hours with a curtailment factor greater than 20% ($(HSL - Generation)/Capacity > 0.2$) is at least 12 hours, 6 hours, or 0 hours. These restrictions further reduce the likelihood that our results are driven by curtailment effects.

Additionally, we conduct a variety of heterogeneity and robustness checks based on the general specification in equation 5. More specifically, we interact the electricity price measures above with PTC expiration identifiers (based on generator age) and measures of exposure in the merchant market based on PPA contracting (from American Cleanpower). These heterogeneous treatment estimates will allow us to test if firms that are not receiving

[17] Negative hourly prices are possible because of the production subsidy received by some wind generators. Negative prices suggest wind generators are on the margin and, thus, curtailment is more likely. When examining the hourly data that underlie our daily analysis and with the restriction that average *daily* prices are greater than \$5/MWh, only about 1.4% of the individual hourly prices in the sample span are negative. With a restriction that daily average prices are greater than \$10, less than 1% of the hourly prices are negative. These restrictions then imply that hours in which renewables are on the margin and curtailment is occurring constitute a negligible portion of our sample.

a PTC subsidy and have substantial market exposure are more responsive to prices, per our theoretical model predictions. The robustness checks consider a variety of subsamples and control variations, which are discussed more thoroughly in the results section.

While the above regression in Equation 5 has a simple economic interpretation as an elasticity, to allow for more response flexibility we estimate the following price-bin model:

$$CF_{iht} = \sum_{k=1}^{B} \theta_k D_{kht} + \sum_{j=t-1}^{t} \gamma_j X_{ij} + \alpha_i + \eta_{my} + \delta_{dow} + \epsilon_{iht}, \tag{6}$$

where D_{kht} is an identifier variable such that $D_{kht} = 1$ if the price measure P_{ht-1} is in the k^{th} bin of the electricity price distribution.[18] The coefficients of interest θ_k correspond to the estimated differences in capacity factors when prices are in the k^{th} bin relative to the lowest bin. We explore two different binning techniques. First, we do simple $10-width bins, with varying number of observations per bin. Second, to create equal sample support by bin, we consider price-decile binning (see Table A.1 for decile ranges and mean prices).

4 Results

We begin by presenting the results for samples including only data for wind generators. We then present the results for solar generators.

4.1 Elasticities - Base Specifications

Table 2 presents the price elasticity estimates as derived from the estimation of Equation (5) for wind generators. The OLS estimate of the elasticity from Equation 5 in column (1)

[18] Given the number of indicator variables and their discreet values, an IV approach was not feasible, so we based the bin indicators on the lagged prices. We again control for lagged system-wide wind generation to limit endogeneity and simultaneity issues.

is negative due to the simultaneity bias discussed above. However, the instrumental variable approach in column (2) implies a small, but statistically significant price elasticity of supply of about 0.14. The elasticity estimate when proxying for the DAM price with its lagged value in column (3) gives a similar value of 0.10. To put these estimates in context, the PTC of approximately $25 per MWh is about 70% of the average price in our sample. Our elasticity estimates would then imply that receiving the PTC, relative to not having it, should increase capacity factors by about 5 to 7 percent, which is very near the PTC effects on capacity factor estimated in Aldy et al. (2023) and Ricks and Kay (2025). This suggests that the production adjustment mechanisms in response to large, long-term price changes for wind generators are likely similar (at least in effect size) to those used to adjust to shorter-term fluctuations in price.

As noted above, a positive relationship between price and production may be driven by curtailment of renewables in very low-demand/high-renewable production periods. While observations with prices below $5 were already excluded above, Table 3 presents results where we further try to minimize the curtailment response channel. More specifically, we consider samples where we exclude observations if $P_{ht} \leq \$10$ or $15 (columns (1), (2), and (3) of Table 3). We also consider samples where we exclude daily observations if the number of high curtailment hours ($(HSL - Generation)/Capacity > 0.20$) within the day for the given generator is greater than 12, 6, or 0 (columns (4), (5), (6), and (7) of Table 3). Results from these sample cuts are numerically quite similar to the base elasticity results in Table 2, suggesting the base specification results are not meaningfully driven by generation responses due to curtailment. Furthermore, these results, particularly column (6) where we eliminate daily observations for generators with any curtailed hours, provide confirmation that the estimated increases in output due to the PTC in Aldy et al. (2023) and Ricks and Kay (2025) are likely due to real price-driven increases in output by firms, and not simply curtailment or reshuffling of the merit order.

4.2 Robustness Checks

To further consider the robustness of this positive price elasticity estimate, we consider a battery of other sample restrictions and specifications. These alternatives include specifications that remove extreme high price days (keep only days with "$P_{ht} < 300$"), remove observations from generators that are consistently experiencing low average daily capacity factors in a given year (remove if the average of CF_{it} observations in a given year is "Avg. CF < 0.2"), remove generators from the sample that do not have at least 365 observations ("Obs. $>$ 365"), remove generators with capacity less than 80 MW ("Capacity $>$ 80"), remove generators that went into service prior to 2007 ("Start Yr. $>$ 2007"), consider specifications with current and lagged meteorological controls ("Lagged Meteo."), replace meteorological controls with the ERCOT-derived forecasted HSL for a given generator made 48 to 72 hours prior to day t ("Forecast. HSL Control"), and consider clustering of the standard errors at the EIA plant-code and Year-Week ("Cluster PlantCode Year-Week") and plant-code and Year-Month ("Cluster PlantCode Year-Month") levels. The results of these specifications for both the IV and lagged price approaches are given in Figure 1. Parameter estimates of the supply elasticity across all of these alternatives are relatively stable when compared to the base specification results in Table 2.

4.3 Dynamic Effects

While the elasticity estimates appear stable across various sample restrictions and control variables, one may be concerned that the mechanisms behind the production adjustments create a dynamic response to prices, such that higher prices do not generate *net* production gains over longer periods (e.g. intertemporal maintenance adjustments such that a small increase in output today is offset by an equal decrease later). We formalize this idea with a more dynamic maintenance model given in Appendix B. In that setting, renewable operators may engage in maintenance activities at t, possibly at the cost of some contemporaneous production ("downtime"), in anticipation of higher prices in the future, and maintenance done at t may positively affect production for several periods ("persistence"). As we discuss in Appendix B, the size of the maintenance downtime and persistence effects can impact the observed dynamic response to a price shock at t. To explore this empirically, we consider a more dynamic version of Equation 5 by including leads and lags of P_{ht}. Results of these specifications are given in Table 4. To begin, we consider various specifications with price leads (i.e. $\log(P_{ht+1})$, $\log(P_{ht+2})$, or $\log(P_{ht+3})$ added to equation 5 to make inferences on the severity of the downtime-effect.[19] If the downtime effect is large, then high future prices would be met with large reductions in production today. Across columns (1)-(3) of Table 4 we find negative parameters on the future price variables, suggestive of a downtime effect. However, the effect is small in size and the summed effect of the parameters on the contemporaneous and future $\log(P_{ht})$ variables is close to, and not statistically different from,

[19] Note that in all specification in Table 4 we instrument for $\log(P_{ht})$ as detailed above, but not for leads or lags of prices because they do not present the direct mechanical simultaneity bias we discuss above. We do, however, include aggregate wind controls based on the dynamic specification used (e.g. we include W_{t+1} when including $\log(P_{ht+1})$, W_{t+2} when including $\log(P_{ht+2})$, and so on).

the base specification elasticity results given in Table 2.

If there was a significant persistence effect from maintenance, then we would expect contemporaneous generation to be affected by past prices and have a more muted response to contemporaneous prices. The results presented in columns (4)-(6) of Table 4 all estimate a positive parameter on the lagged $\log(P_{ht})$ variable, consistent with persistence, but that effect is always smaller than the parameter on the contemporaneous price control.[20] This suggests that maintenance persistence does not dampen future responses to prices. Combined, the results of Table 4 suggest that wind generators positively respond to prices in ways that lead to true generation gains in high price periods that are not counteracted by some dynamic reshuffling of maintenance activities.

4.4 Heterogeneous Effects - PTC Expiration and Market Exposure

We further explore the heterogeneity of price responsiveness by examining interaction models with price subsidies (i.e., PTC's) and wholesale market exposure (i.e. how much of a firm's generation is contracted under PPAs) for wind generation facilities. As modeled in Section 2, receiving a production subsidy and having less exposure to market prices should reduce the price responsiveness of a given generator. To empirically explore this, we consider versions of equation 5 where we interact the $\log(P_{ht}$ term with either an indicator variable that equals

[20] The results in column (4) of Table 4 show a large effect on $\log(P_{ht})$ than on $\log(P_{ht-1})$ though the parameter $\log(P_{ht-1})$ is much more precisely estimated. However, the 1-period differentiated prices are highly correlated, so separate identification of these effects will be difficult. In addition, as noted above, if the adjustment mechanisms take time to complete and lagged prices are good proxy for future prices, generation at t may be responding to a lagged price signal not because of a maintenance persistent effect, but rather because of speed of adjustment restrictions.

one if the generator is no longer receiving the PTC ($1(\text{PTC EXP})_{it}$) or one if the generator likely faces some wholesale market exposure ($1(\text{Merchant})_i$). For $1(\text{PTC EXP})_{it}$, we set the indicator equal to one if day t is ten years past the initial operating year and month (assuming a mid-month start day) as defined in the EIA 860 data. To create $1(\text{Merchant})_i$, we use data from CleanPower Association. Unfortunately, this data is only available for 2022, giving us a snapshot of the merchant status of the generators and requiring the implicit assumption that the merchant status of firms is relatively stable.

These results are given in Table 5. Columns (1) and (2) of Table 5 show the estimates of an alternative version of equation 5 that allows the price responsiveness to vary for those facilities whose PTC has expired. We find that those facilities that no longer receive the PTC are considerably more price responsive than those still receiving the PTC (approximately 40-80% more responsive, depending on estimation approach). This is consistent with predictions of Section 2. Additionally, the PTC expiration indicator variable ($1(\text{PTC EXP})_{it}$) is measured with some error as we only know the month and year of the start date for the facility and not the precise PTC activation date, so the difference in price responsiveness across PTC and non-PTC generators may be larger.

Similarly, columns (3) and (4) of Table 5 show the results for an altered equation 5 that allows price responsiveness of wind facilities to be different for those that, at least for 2022, have some exposure to the wholesale market ("merchant" generators) versus those that are listed as having their production pre-sold via contracts such as PPA's. Despite the measurement error in our "merchant" designation, we do find these merchant facilities have price elasticities that are somewhat larger than the PPA-selling counterparts, but the effect size is, as expected, muted and imprecisely measured.

We explore several other dimensions of heterogeneity. To begin, as can be seen in Appendix Figure A.3, average wind generation follows a diurnal pattern with more generation in the night and early morning hours and less in the middle of the day. We therefore consider the response of daily wind generation to the average price over hours where wind generation is typically high (0-6 and 19-23), P^{Hi}_{ht-1} and over hours when it is low (7-18), P^{Lo}_{ht-1} and estimate the response to each price. Results of this specification are given in Appendix Table A.2. We find wind generation is slightly more responsive to prices over the low-wind generation periods, which is also when price is highest, though the difference in the responsiveness to the two prices are not statistically different from one another. We also explore a specification where we interact the price term with wind speed at the generator level. These results are also given in Appendix Table A.2, in column (2). Here we find the price responsiveness is higher in higher wind periods, suggesting that higher wind environments enhance the return to maintenance (potentially due to the cubic relationship between output and wind speed).

Next, we explore the heterogeneity of price responsiveness by wind generator characteristics. We consider dimensions of characteristics which may reasonably affect maintenance costs or returns. Specifically, we consider varying responses by quartile of the number of turbines and the height of wind turbines in the wind farm. Additionally, EIA-860 data provides a measure of the quality of the wind resource where the wind generator is built, ranging from 1-4, with wind quality 1 regions having the highest average winds and wind quality 4 having the lowest. The results of these specifications are given in Appendix Table **??**. With respect to the number of turbines, facilities in the lowest quartile of turbine count have lower responsiveness than those with a higher number of turbines, which may be as expected given more turbines give more margins of maintenance. With respect to turbine height, we find facilities in the highest quartile of turbine height are less responsive than those with lower turbine height. This may be as expected if it is more costly to conduct generation-improving maintenance on taller turbines. However, the newer turbines are built considerably taller than older versions, so the effect of high turbines is also picking up the PTC effect discussed above. Finally, with respect to wind quality class, we find facilities built in higher wind-quality regions (1 and 2) are more responsive than those in lower wind-quality places (3 and 4). This is consistent with the specification where we interact price measures with wind speed and suggests that maintenance is marginally more valuable in high wind regions.

4.5 Alternative Specification - Price Bin Responses

Given the wide variation in prices observed, a constant elasticity form may not hold throughout the price distribution. Additionally, the constant elasticity specification is undefined when prices are negative, a condition that may become more frequent as renewable penetration increases. The bin specification in 6 can be used to assess responses across price ranges. To begin, we create $10-width bins over the price range from 0-$140/MWh and then a top-coded bin for prices over $140 and a bottom-coded bin for negative prices.[21]

Figure 2 provides a plot of the estimated parameters associated with the bin indicators for wind generators (using the negative price bin as the reference group). In general, the response pattern shows an initial and substantial jump in capacity factor as prices go from the negative price bin to the 0-$10 and $10-$20 bins. The capacity factor response continues

[21] This binning width and top-coding for the upper bins creates a sample with at least 1000 observations per bin, thus providing flexibility along with sufficient variation for estimation.

to increase until about the \$40/MWh price and then remains flat until around \$100/MWh, where it experiences another jump. This suggests that much of the price responsiveness is coming at the extremes of the price distribution.

We consider another version of equation 6 by forming price deciles of the price distribution. Again we see a pattern of large responses in capacity factor going from the lowest decile to the mid-range decile and then another jump as prices go into the highest decile. Furthermore, we also considered specifications of equation 6 where we interact the bin indicators with the PTC-expiration indicator. The pattern remains relatively the same, though it is accentuated for the PTC-expired facilities. Plots of these estimates are given in the appendix (see Appendix Figures A.1 and A.2).

The results highlight the importance of how changes in the electricity price distribution, beyond just the change in average price level, may affect generation from existing generators. More specifically, with more renewable generation, electricity prices are expected to decrease on average, but may also see more extreme high and low prices (Mallapragada et al. 2023). The results of this decile and price-bin analysis highlight the importance of how changes in the tails of the price distribution can have significant impacts on overall production levels from existing facilities, particularly those that no longer receive the PTC. We return to this point below in our discussion of the implications of the estimation results.

4.6 Elasticity Estimates - Solar Generators

We now turn our focus to estimating the price responsiveness of solar generators. Solar generation is growing rapidly in ERCOT, but over the sample we analyze, daily, ERCOT-wide solar generation averaged about 4% of ERCOT load, whereas wind generation averaged about 23% of daily load.[22] As a result, a small percentage change in aggregate solar generation will have a smaller effect on prices. In addition, the correlation coefficient between individual solar generators capacity factor and aggregate solar generation is much smaller for solar than it is for wind generators (0.26 compared to 0.54). This suggests that the mechanical simultaneity bias that is a major concern for wind is likely less of an issue for solar. As a result, we focus mainly on OLS estimation of equation 5, though we also explore a similar IV strategy for solar generators. Note also, because we are conducting a daily analysis, but solar generation is relevant for approximately half the day, we restrict our aggregation of hourly observations to the daily level over the typical daylight hours of 7-18.

Column (1) of Table 6 gives the OLS elasticity estimate for solar generators. Unlike the corresponding estimate for wind generators, we find a positive elasticity estimate from OLS for solar, again highlighting the diminished concern for simultaneity bias for solar. However, once we remove daily observations from solar generators that experience a relatively high number of high-curtailment hours, the elasticity estimate falls dramatically in magnitude and is no longer statistically different from zero.[23] Similarly, the elasticity estimate based on the IV approach is also small in magnitude and not statistically different from zero. These results suggest that solar generators are not meaningfully responsive to daily electricity prices in the same way that wind generators are, or if they are responsive it is primarily via forced reductions in generation due to localized supply/demand imbalances and not because of solar

[22] Even if one only considers the predominantly daylight hours (hours 7 - 18), the aggregate solar generation in these hours as a percent of aggregate load is only about 12%)

[23] Note we focus on high-curtailment event restrictions to deal with curtailment instead of restricting the sample to higher minimum prices because the 7-18 hour span sees relatively few negative price hours. However, despite regional prices during this span rarely going negative, very specific nodal prices may still be negative due to localized supply/demand imbalances that can force curtailment events.

operators adjusting production margins to take advantage of high prices.

5 Implications of price responsiveness

How important is the intensive margin response identified above? We first extend the analysis of wind production subsidies in Hahn et al. (2026) to include the intensive margin response in the theoretical calculation of the marginal value of public funds (MVPF). We then conduct two simulation-based exercises leveraging the intensive margin estimates of price-response above to illustrate its potential importance. First, we use estimated supply elasticities for wind generators in an economic dispatch model of the ERCOT market. This exercise compares generation, price, and emission outcomes under scenarios with and without price responsiveness for wind generators in the context where total generating capacity remains constant, but electricity demand is increasing. This simulation is relevant to recent considerations about environmental and electricity price effects associated with recent data-center-driven demand growth projections. The second exercise is a more long-run analysis where we consider the effect that prices under a renewable-dominated electricity grid may have on existing generators. More specifically, we leverage the price analysis of the U.S.-based Inflation Reduction Act (IRA) by Bistline et al. (2023). While the significant clean energy provisions of the IRA have largely been repealed, the analysis of Bistline et al. (2023) can also be viewed as a detailed study on the electricity price effects in the U.S. from a large influx of heavily subsidized renewables.

5.1 MVPF of wind production subsidies - intensive margin

In Hahn et al. (2026), wind production subsidies have a large MVPF of 4.63, implying an additional \$3.63 of social benefit per dollar of government expenditure. Their calculation uses an extensive margin elasticity of 1.13, derived from linear probability model estimates of wind adoption due to production incentives in Hitaj (2013). In Appendix B, we derive the MVPF inclusive of our estimated IV intensive margin price response elasticity of 0.14 in Table 2 and calculate an MVPF of 4.971, implying an extra \$0.34 per dollar of government expenditure. Thus, through the intensive margin response, for every \$10 billion in production subsidy expenditure (roughly the annual US federal expenditure on wind production subsidies), an additional \$3.4 billion in social benefits are generated, assuming the marginal effects scale linearly.

However, Hitaj (2013) notes that the effects of production incentives may be endogenous and thus estimates a series of IV models, with a mean IV point estimate of a 5.01% increase in new MW capacity per 1 cent/kWh increase in production subsidy. Using the implied extensive margin elasticity of 0.509 from these IV estimates gives an MVPF of 2.855, or \$1.855 in additional social benefit per dollar of government expenditure. In this case, incorporating the 0.14 intensive margin elasticity gives an MPVF of 3.295, implying an extra \$0.44 per dollar of government expenditure from the intensive margin response.

5.2 Implication Exercise 1 - Short-term Demand Growth

The Energy Information Administration (EIA) forecasts electricity demand growth of up to 3% a year through 2030, driven in large part by data center growth. New generation sources are projected to come online, but current forecasts suggest they will lag the growth in load. This imbalance will increase wholesale electricity prices, and the results above suggest wind generators will respond by increasing generation. Failing to account for price responsive wind generation in forecasting the impact of load growth will lead to an overestimate of the impact of data centers on prices and pollution emissions.

We develop a simple dispatch model for the ERCOT electric grid that incorporates price responsive wind. We simulate the grid every hour for the year 2023 and demonstrate that our baseline model accurately represents wholesale prices and fuel mixes across ERCOT zones. Our baseline scenario predicts an average marginal cost of \$43.87, just below the 2023 actual marginal cost of around \$47.[24]

We then model three counterfactual scenarios based on ERCOT's large-load interconnection queue, which suggests that 14.5 GW of new data center load will enter the market by the end of 2026.[25] We use the capacity of the proposed site and interconnection point to identify the changes to each zone's load levels, and the three scenarios then consider how that load enters the grid. First, we obtain data on hourly demand for a data center in MISO observed over two months and parameterize our first counterfactual so that the 14.5 GW of new data center capacity is dispatched in the same way. Because this data center's dispatch profile is relatively flat, the second scenario considers a case in which 50% of capacity matches the observed data center load and 50% is load-following. The final scenario assumes all load-following, so that the new capacity essentially replicates the shape of the current hourly load. We then run our dispatch algorithm for each hour of the year to identify the

[24] Our dispatch model does not incorporate within-zone transmission constraints. This simplification likely accounts for the difference in estimated marginal cost between our baseline and the data.

[25] We use ERCOT's interconnection queue to identify projects in the construction phase with an expected in-service date before the end of 2026.

impact of large loads on wholesale prices, fuel mix, wind generation by zone, and pollution emissions by zone. We then allow for price-responsive wind, using a supply elasticity of 0.14 based on the baseline IV estimates reported in Table 2. Comparing counterfactuals with and without price-responsive wind illustrates the importance of allowing for wind response. Table 7 presents the results.

Data center load adds 126 TWh to total demand over the course of a year. This additional load increases total costs by $1.6-2.6 billion depending on the scenario, and CO_2 emissions increase by 16-30 million tons per year. The shape of load affects the price impact: as data center load becomes more load-following, both the level and variance of prices increase.

Ignoring wind response overstates the magnitude of these effects. Higher prices induce approximately 1% more wind generation relative to the no-wind-response baseline. This reduces the cost of serving data center load by $20-30 million per year depending on the scenario, yielding cost reductions of roughly 25 cents per megawatt-hour and CO_2 reductions of 130-210 thousand tons per year. These results suggest that incorporating price-responsive wind supply into counterfactual analyses of load growth has meaningful impacts on estimated total effects.

5.3 Implication Exercise 2 - Long-run Growth in Renewables

Our second exercise leverages model-based projections of regional electricity price distributions inclusive of the impacts from the Inflation Reduction Act from Bistline et al. (2023).[26] Again, while the IRA is now essentially defunct, the primary driver of the price changes

[26] While this exercise provides a useful look at how generators would respond to plausible counterfactual price distributions, we do require the hearty assumption that the estimated wind generator price-responsiveness in ERCOT generally applies to other regions.

that are estimated in Bistline et al. (2023) comes from the massive expansion of renewable generation capacity projected to come online with the IRA-based subsidies.[27]

The underlying model in Bistline et al. (2023) US-REGEN solves for 120 representative hours for 16 different U.S. regions every 5 years through 2050, yielding hourly price distributions as shown for Texas for 2020 and 2050 in Figure 3. Note that while average prices decline by approximately \$15/MWh, the distribution shifts considerably for 2050, with more frequent higher and lower prices due to the influx of renewables. Given these price distributions, for each region and year we map prices to the \$10-width price bins used in our analysis to create distribution weights ω_k^y, which essentially describe the frequency in which the prices from the given year y fall into the k^{th} bin. We then calculate the change in capacity factor from 2020 to 2050, using the price-bin effects θ_k displayed in Figure 2, for each region's price forecasts as:

$$\bar{CF}^{2050} - \bar{CF}^{2020} = \sum_{k=1}^{B} \omega_k^{2050}\theta_k - \sum_{k=1}^{B} \omega_k^{2020}\theta_k \tag{7}$$

We similarly estimate the change in capacity factor implied by a constant elasticity of supply (CE) function (i.e. $CF = AP^{\beta}$), using the elasticity estimate (β) given in column (2) of Table 2.[28] In this procedure, we parameterize A based on the 2020 average price and normalize CF to one. Then, using the other projected prices for 2020 and for 2050, we can calculate how capacity factors move relative to the initial normalized condition. To get an estimate of the average change in capacity factor across the years, we use the relative deviation estimated at each price point and multiply it by the average capacity factors for

[27] The Bistline et al. (2023) predicts average annual deployments of wind, solar, and storage capacity of about 51GW/year through 2035 under the IRA, compared to only 27GW/year of additions for these technologies over the same span in a scenario without the IRA.

[28] We use the estimate based on lagged prices for consistency because the price-bin effect estimates use bins based on lagged prices.

the region as calculated for 2020. We then essentially have a capacity factor estimate for each price point and we can then calculate implied differences. Note, however, in the CE specification, the negative price region is undefined so we simply bottom code the projected prices at \$5/MWh, and as such this specification may underestimate the effect of more negative price hours.

Table 8 reports the results of this exercise. For the eight major wind producing regions in the US-REGEN model, Table 8 displays the average 2020 and predicted 2050 prices under the IRA (the $\bar{P}_{2020}$ and $\bar{P}_{2050}$ columns, respectively), along with the calculated average capacity factor for wind generation in each region ($\bar{CF}$) based off the EIA's 2020 state-level generation and capacity information (https://www.eia.gov/electricity/data/state/).[29] Column (4), "ΔCF_{bin}", of the table gives the average predicted change in capacity factor for a given facility calculated from equation (7). Column (5), "ΔCF_{CE}", gives the average change in capacity factor from the relatively low- to high-renewable case using the constant elasticity formulation of supply.

Several key points emerge from the analysis. First, the drop in capacity factor, relative to mean capacity factor values, associated with moving to the high-renewable scenario ranges from about a one- to a seven-percent decrease in productivity depending on the region and method of estimation. To further put this in perspective, if the average capacity factor of Texas wind generators fell by 0.019 points (roughly 5.8% of the mean), that would equate to an annual loss of generation of approximately 7,000 GWh of electricity based on 2024 capacities. Second, the implied losses in generation are slightly higher based on the estimates

[29] A map of the market regions from the US-REGEN model is given in Appendix Figure A.4. We exclude eight of these regions because wind generation is small and predicted to remain so, given the minimal wind generation potential.

from the binned-price regressions in regions that experience large increases in low-price extremes. For example, as can be seen from Figure 3, the projected prices with a large influx of renewables (the 2050 price) in Texas has many more low-price hours compared to the 2020 prices. This large increase in low-price periods further reduces capacity factor for the binned-price response. For other regions however, like MISO-North (MISO-N), the shift in the price distribution comes more from a drop in the mid-level prices, which elicits less of a response in capacity factor when based on the binned-price estimates.

6 Conclusion

Generally, it is assumed that renewable generators do not respond on the intensive margin to financial incentives. However, recent analysis of the effect of wind generation production subsidies on generation from existing generators challenges this notion.

In this analysis, we use high-frequency, spatially-detailed generation and weather data from Texas to show that wind generators respond to short-run variation in electricity prices. This finding is robust to a variety of specification and data restrictions, with an estimated short-run supply elasticity on the order of 0.10. Furthermore, through a variety of sample restrictions, this output response to prices does not appear to be driven by ERCOT-system-operator decisions to curtail wind generation in periods of low-demand/high-renewable-production. Additionally, the estimated price responsiveness leads to production changes similar to those predicted by the inclusion of the U.S. PTC (Aldy et al. 2023) and the loss of the PTC (Ricks and Kay 2025). This suggests similar mechanisms may be used by firms in response to transitory price movements as those for more permanent price movements.

Finally, consistent with theoretical predictions, we also find evidence that wind generators no longer receiving production subsidies and generators more exposed to market prices are more responsive to daily price changes.

We examine two simulation-based implications of this price-responsiveness. First, using a dispatch model of ERCOT, we show that price-responsive wind generators mitigate the price and emissions effects of new data center entry by tens of millions of dollars in costs and hundreds of thousands of tons of CO_2 annually. Second, using price projections based on high future levels of renewable deployment, we find high renewable penetration can lead to significant reductions in wind capacity factor, even in cases where the average price change from low- to high-renewable penetration is small. This highlights how high-renewable penetration may have significant impacts on generation from existing facilities, beyond just a curtailment mechanism. More generally, our results indicate that operation margin decisions by existing renewable firms may be increasingly important as the share of renewables on the grid increases.

References

Aldy, J. E., T. D. Gerarden, and R. L. Sweeney (2023). Investment versus output subsidies: Implications of alternative incentives for wind energy. *Journal of the Association of Environmental and Resource Economists 10*(4), 981–1018.

Belyak, N., S. A. Gabriel, N. Khabarov, and F. Oliveira (2024). Renewable energy expansion under taxes and subsidies: A transmission operator's perspective. *Journal of Cleaner Production 451*, 141955.

Bistline, J. E., N. R. Mehrotra, and C. Wolfram (2023). Economic implications of the climate provisions of the inflation reduction act. *Brookings Papers on Economic Activity 2023*(1), 77–182.

Bjørn, A., S. M. Lloyd, M. Brander, and H. D. Matthews (2022). Renewable energy certificates allow companies to overstate their emission reductions. *Nature Climate Change 12*, 508–509.

Chan, H. R., H. Fell, I. Lange, and S. Li (2017). Efficiency and environmental impacts of electricity restructuring on coal-fired power plants. *Journal of Environmental Economics and Management 81*, 1–18.

Cicala, S. (2015). When does regulation distort costs? lessons from fuel procurement in us electricity generation. *American Economic Review 105*(1), 411–444.

Davis, L. and C. Hausman (2016). Market impacts of a nuclear power plant closure. *American Economic Journal: Applied Economics 8*(2), 92–122.

Deschenes, O., C. Malloy, and G. McDonald (2023). Causal effects of renewable portfolio standards on renewable investments and generation: The role of heterogeneity and dynamics. *Resource and Energy Economics 75*, 101393.

DiGangi, D. (2025). $1.6b in clean energy investments were rolled back in september: E2. *Utility Dive*.

Doyle, M. and H. Fell (2018). Fuel prices, restructuring, and natural gas plant operations. *Resource and Energy Economics 52*, 153–172.

Electric Reliability Council of Texas, I. (2024, June). Nodal protocols: Section 2 – definitions and acronyms (effective june 28 2024). Word document (docx). Accessed: March 27, 2026.

Fabrizio, K. R., N. L. Rose, and C. D. Wolfram (2007). Do markets reduce costs? assessing the impact of regulatory restructuring on us electric generation efficiency. *American Economic Review 97*(4), 1250–1277.

Fell, H., S. P. Holland, and A. J. Yates (2025). Optimal subsidies for green hydrogen production. *Journal of the Association of Environmental and Resource Economists 12*(1), 33–63.

Fell, H. and D. T. Kaffine (2018). The fall of coal: Joint impacts of fuel prices and renewables on generation and emissions. *American Economic Journal: Economic Policy 10*(2), 90–116.

Fell, H. and J. Linn (2013). Renewable electricity policies, heterogeneity, and cost effectiveness. *Journal of Environmental Economics and Management 66*(3), 688–707.

Hahn, R. W., N. Hendren, R. D. Metcalfe, and B. Sprung-Keyser (2026). A welfare analysis of policies impacting climate change. *American Economic Review Forthcoming*.

Hitaj, C. (2013). Wind power development in the united states. *Journal of Environmental Economics and Management 65*(3), 394–410.

Hollingsworth, A. and I. Rudik (2019). External impacts of local energy policy: The case of renewable portfolio standards. *Journal of the Association of Environmental and Resource Economists 6*(1), 187–213.

Kaffine, D. T. and C. M. Worley (2010). The windy commons? *Environmental and Resource Economics 47*, 151–172.

Mallapragada, D. S., C. Junge, C. Wang, H. Pfeifenberger, P. L. Joskow, and R. Schmalensee (2023). Electricity pricing challenges in future renewables-dominant power systems. *Energy Economics 126*, 106981.

Novan, K. and Y. Wang (2024). Estimates of the marginal curtailment rates for solar and wind generation. *Journal of Environmental Economics and Management 124*, 102930.

Plumer, B. (2014). The rise and fall of wind power in america. *Vox*.

Ricks, M. and O. Kay (2025). Time-limited subsidies: Optimal taxation with implications for renewable energy subsidies. *Journal of Political Economy*.

Shea, J. (1993). Do supply curves slope up? *The Quarterly Journal of Economics 108*(1), 1–32.

Table 1: Wind and solar generator summary statistics ERCOT 2016-2024

Wind	mean	sd	min	max	N
Capacity(MW)	176	78.7	30	525	499476
Capacity Factor	0.33	0.22	3.75e-07	1.24	499,476
DAM Price ($/MWh)	36.61	54.13	5.02	1077.74	499,476
Age (months)	93.75	64.08	0.03	312.54	499,476
Solar					
Capacity(MW)	154	92.6	24.80	600	115,443
Capacity Factor	0.50	0.24	1.40e-06	1.24	115,443
DAM Price ($/MWh)	53.73	107.03	5.05	2321.10	115,443
Age (months)	32.96	28.78	0.03	156.23	115,443

Notes: Unit of observation is a wind or solar generator (as defined by the EIA) by day. Averages are based on sample from column (1) of Tables 2 and 6. The solar "DAM Price" and "Capacity Factor" values are daily averages based on hourly observations from hours 7 to 18. There are 208 wind generators in total and 116 solar generators.

Table 2: Wind generation elasticity with respect to DAM prices

Dep var: $\log(CF_{iht})$	(1)	(2)	(3)
$\log(P_{ht})$	-0.130***	0.142***	
	(0.0160)	(0.0501)	
$\log(P_{ht-1})$			0.102***
			(0.01393)
Observations	499,476	436,532	495,591
IV	N	Y	N
R^2	0.521	0.406	0.530

Notes: Coefficient represents the elasticity of capacity factor to DAM prices. Control variables include quadratic meteorological variables, generating unit fixed effects, month-by-year and day-of-week fixed effects. Instruments in the "IV-Y" column are the log of weather-zone-specific and ERCOT-wide forecasted load made 48 to 72 hours ahead of day t and the log of one- and two-day lagged Henry Hub natural gas prices. First stage Kleibergen-Paap rk Wald F statistic is 84.6. Standard errors, clustered by EIA plant code and day, are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Robustness checks - Curtailment Exclusion

Dep var: $\log(CF_{iht})$	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\log(P_{ht})$	0.147***	0.159***		0.143***	0.139***	0.132**	
	(0.0518)	(0.0557)		(0.0511)	(0.0518)	(0.0534)	
$\log(P_{ht-1})$			0.108***				0.0980***
			(0.0163)				(0.0149)
IV	Y	Y		Y	Y	Y	
Hi Curt.	None	None	None	≥ 12	≥ 6	0	0
min. P_{ht}	10	15	15	5	5	5	5
Observations	421,549	391,186	445,713	429,475	419,639	391,105	446,856

Notes: "Hi. Curt" events are defined as an hour where (HSL- Generation/capacity of the generating resource exceeds 0.2. These events are summed to the EIA plant-code/generator level and summed over the day. Columns (4) and (5) drop daily observations where intraday summed curtailment events are greater than or equal to 12 and 6, respectively. Columns (6) and (7) drop daily observations that had any hour with a high curtailment event. Instruments in the "IV-Y" column are the log of weather-zone-specific and ERCOT-wide forecasted load made 48 to 72 hours ahead of day t and the log of one- and two-day lagged Henry Hub natural gas prices. Control variables include meteorological variables, regional load, ERCOT-wide generation by technology, generating unit fixed effects, month-by-year fixed effects, and day-of-week fixed effects. Standard errors, clustered by EIA plant code and day, are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Dynamic Effects

Dep var: $\log(CF_{iht})$	(1)	(2)	(3)	(4)	(5)	(6)
$\log(P_{ht})$	0.137**	0.132***	0.141***	0.0894	0.132**	0.137***
	(0.0562)	(0.0502)	(0.0496)	(0.0570)	(0.0521)	(0.0507)
$\log(P_{ht+1})$	-0.0209*					
	(0.0124)					
$\log(P_{ht+2})$		-0.00415				
		(0.00772)				
$\log(P_{ht+3})$			-0.00527			
			(0.00853)			
$\log(P_{ht-1})$				0.0484***		
				(0.0127)		
$\log(P_{ht-2})$					0.0125*	
					(0.00686)	
$\log(P_{ht-3})$						0.00580
						(0.00714)
Sum Effect	0.116***	0.128***	0.136***	0.138**	0.145***	0.142**
	(0.0467)	(0.0477)	(0.0483)	(0.0495)	(0.0499)	(0.0498)
Observations	435,984	435,549	435,289	436,532	436,532	435,831

Notes: "Sum Effect" gives the sum of the parameter estimates on contemporaneous, plus leads and lags, $\log(P_h)$. Standard errors of this sum are given in parentheses below. The parameter on $\log(P_{ht})$ is an IV estimate with instruments of log of weather-zone-specific and ERCOT-wide forecasted load made 48 to 72 hours ahead of day t and the log of one- and two-day lagged Henry Hub natural gas prices. Control variables include meteorological variables, regional load, ERCOT-wide generation by technology, generating unit fixed effects, month-by-year fixed effects, and day-of-week fixed effects. Standard errors, clustered by EIA plant code and day, are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Heterogeneous effects for wind generators by PTC and market exposure status

Dep var: $\log(CF_{iht})$	(1)	(2)	(3)	(4)
$\log(P_{ht})$	0.121**		0.116**	
	(0.0507)		(0.0518)	
$\log(P_{ht})$*1(PTC EXP)$_{it}$	0.105***			
	(0.0317)			
$\log(P_{ht})$*1(Merchant)$_i$			0.0584*	
			(0.0312)	
$\log(P_{ht-1})$		0.0910***		0.0953***
		(0.0146)		(0.0156)
$\log(P_{ht-1})$*1(PTC EXP)$_{it}$		0.0404**		
		(0.0143)		
$\log(P_{ht-1})$*1(Merchant)$_i$				0.0173
				(0.0140)
Observations	436,532	495,591	432,841	491,914

Notes: Heterogeneous treatment effects depend on PTC expiration status: 1(PTC EXP)$_{it}$ $= 1$ if the generator is not receiving the PTC, and market exposure: 1(Merchant)$_i$ $= 1$ if the generator sells any of their output as a merchant generator into the wholesale market. Control variables include quadratic meteorological variables, generating unit fixed effects, month-by-year fixed effects, and day-of-week fixed effects. Standard errors, clustered by EIA plant code and day, are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Solar generation elasticity with respect to DAM prices

Dep var: $\log(CF_{iht})$	(1)	(2)	(3)	(4)	(5)
$\log(P_{ht})$	0.159***	0.0245	-0.00635	-0.0179	-0.0421
	(0.0285)	(0.0181)	(0.0165)	(0.0167)	(0.0729)
IV	N	N	N	N	Y
High Curt. Restr.	None	≥6	≥3	0	None
Observations	115,443	110,658	106,477	100,166	115,050

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: "Hi. Curt" events are defined as an hour where (HSL-Net Output)/capacity of the generating resource exceeds 0.2. These events are summed to the EIA plant-code/generator level and summed over the day. Columns (1) and (2), drop daily observations where intraday summed curtailment events are less than or equal to 6 and 3, respectively. Columns (3) and (4) drop daily observations that had any hour with a high curtailment event. Instruments in the "IV-Y" column are the log of weather-zone-specific and ERCOT-wide forecasted load made 48 to 72 hours ahead of day t and the log of one- and two-day lagged Henry Hub natural gas prices. Control variables include meteorological variables, regional load, generating unit fixed effects, month-by-year fixed effects, and day-of-week fixed effects. Standard errors, clustered by EIA plant code and day, are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Effect of Price-Responsive Wind Across Data Center Load Profiles

	Flat DC		**50% Load-Following**		**100% Load-Following**	
	No Wind Response	With Wind Response	No Wind Response	With Wind Response	No Wind Response	With Wind Response
Panel A: Grid Outcomes						
Total Cost ($B)	10.89	10.86	10.89	10.86	9.83	9.81
Avg. Price ($/MWh)	45.69	45.41	46.54	46.22	44.55	44.30
Price Std. Dev. ($/MWh)	21.33	20.97	22.28	21.87	20.82	20.48
Total CO_2 (M tons)	110.9	110.7	112.0	111.8	98.0	97.8
Avg. Wind (MW)	12,305	12,434	12,305	12,462	12,305	12,392
Panel B: Wind Response Effect						
Wind Change (%)	+1.06%		+1.28%		+0.72%	
Cost Reduction ($M)	27		32		19	
Price Reduction ($/MWh)	0.28		0.32		0.25	
CO_2 Reduction (M tons)	0.20		0.21		0.13	

Notes: Wind supply elasticity = 0.14. Flat Data Center (DC) operates at constant 14.5 GW. 50% Load-Following combines baseload with demand-correlated load while holding total energy consumption constant. 100% Load-Following scales entirely with demand. Wind response is stronger with load-following scenarios due to higher price volatility during peak hours.

Table 8: Capacity Factor effects of changes in the price distribution

Region	(1) $\bar{P}_{20}$	(2) $\bar{P}_{50}$	(3) $\bar{CF}$	(4) ΔCF_{bin}	(5) ΔCF_{CE}
Calif.	33.4	31.7	0.28	-0.018	-0.011
MISO-E	24.0	25.6	0.33	-0.005	-0.001
MISO-N	23.5	13.4	0.37	-0.021	-0.023
Mtn-N	24.0	15.1	0.33	-0.022	-0.017
Mtn-S	27.8	22.9	0.34	-0.019	-0.012
Pacific	28.5	28.0	0.28	-0.015	-0.007
SPP	22.1	10.0	0.36	-0.027	-0.025
Texas	37.0	19.3	0.33	-0.019	-0.010

Notes: "Region" identifiers correspond to regions used in Bistline et al. (2023), $\bar{P}_{20}$ and $\bar{P}_{50}$ are average prices in 2020 and 2050 respectively, and $\bar{CF}$ is average regional capacity factors. ΔCF_{bin} are calculated changes in wind capacity factors based on IRA price distributions from Bistline et al. (2023) and the bin responses for PTC-Expired facilities given in Figure 2. ΔCF_{CE} are calculated changes in capacity factor given elasticity estimates from Table 2 column (2).

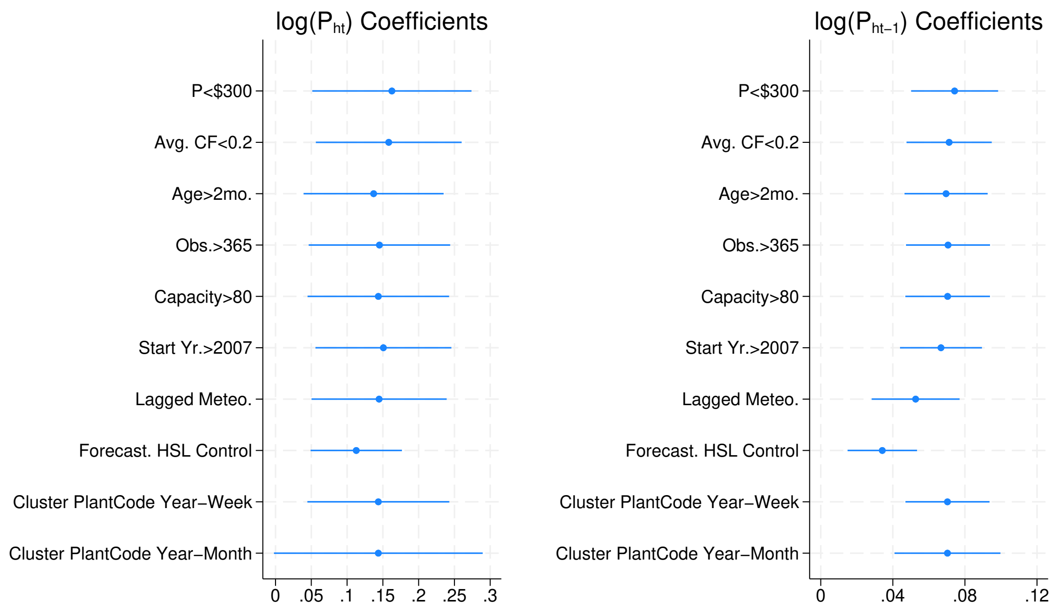

Figure 1: Robustness Checks of Supply Elasticity

Notes: Panels show point estimates of the supply elasticity and 95% confidence intervals based on standard errors clustered at EIA plant code and day of sample, unless otherwise noted. Left panel are IV estimates using instruments log of weather-zone-specific and ERCOT-wide forecasted load made 48 to 72 hours ahead of day t and the log of one- and two-day lagged Henry Hub natural gas prices. All specification includes control variables of quadratic meteorological variables, generating unit fixed effects, month-by-year fixed effects, and day-of-week fixed effects.

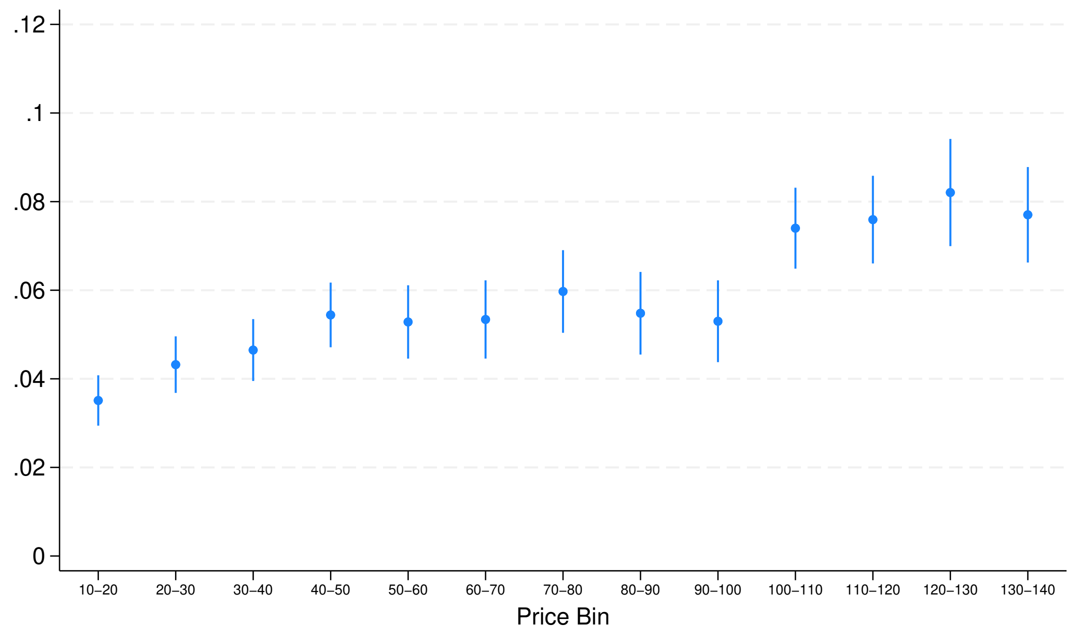

Figure 2: Price Bin Capacity Factor Response

Notes: The plot displays the parameter estimates and 95% confidence intervals of the θ_k parameters given in equation (6). The bin indicators are based on the P_{ht-1} values. The regression also included control variables of quadratic meteorological variables, lagged aggregate wind generation, generating unit fixed effects, month-by-year fixed effects, and day-of-week fixed effects. Standard errors were clustered by EIA plant code.

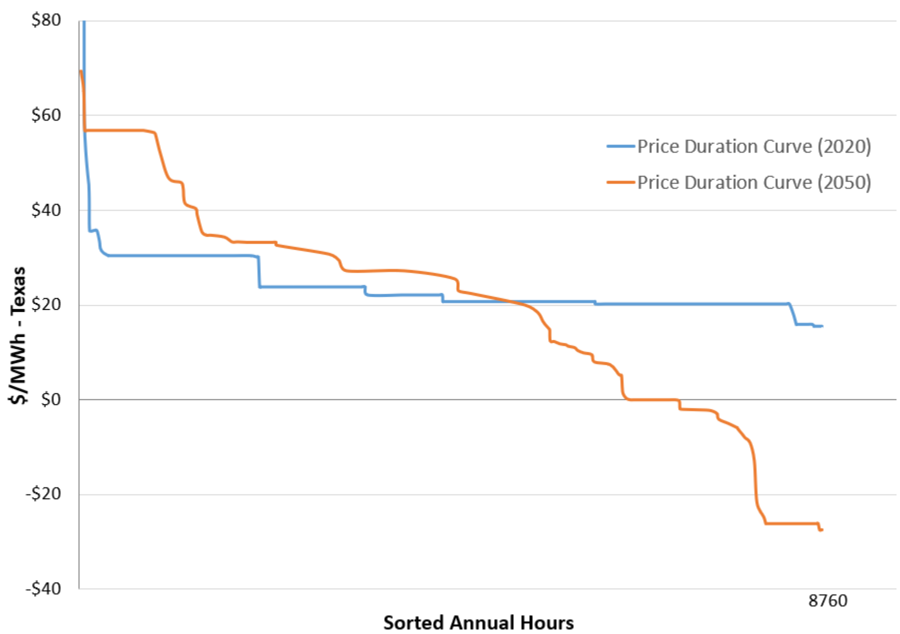

Figure 3: Modeled number of hourly prices for Texas in 2020 versus 2050 from Bistline et al. (2023)

A Appendix Tables and Figures

Table A.1: Price Deciles (DAM) ERCOT 2016-2024

Decile	Mean	Min	Max
1	11.53	5.09	15.06
2	16.84	15.07	18.40
3	19.66	18.40	20.72
4	21.84	20.72	22.89
5	23.96	22.90	25.10
6	26.40	25.10	27.75
7	29.53	27.76	31.49
8	34.49	31.49	38.58
9	45.66	38.59	56.03
10	92.57	56.05	288.72

Notes: The table shows the summary of the lagged DAM prices for the decile analysis.

Table A.2: Heterogeneous Response: Wind Timing and Windspeed

Dep var: $\log(CF_{iht})$	(1)	(2)
$\log(P^{Hi}_{ht-1})$	0.0412***	
	(0.0137)	
$\log(P^{Lo}_{ht-1})$	0.0669***	
	(0.0134)	
$\log(P_{ht-1})$		0.0261
		(0.0301)
$\log(P_{ht-1})$*$Windspeed_{it}$		0.0156***
		(0.0055)
Constant	21.60	14.67
	(17.71)	(16.17)
Observations	469,019	495,591
At Mean		0.0962
		(0.0140)
At 95th		0.140
		(0.0192)
At 5th		0.0621
		(0.0197)

Notes: $\log(P^{Hi}_{ht-1})$ refers to the log of daily average of hourly prices over hour 0-6 and 19-23, lagged one period. $\log(P^{Lo}_{ht-1})$ refers to the log of the daily average of hourly prices over hours 7-18. "At Mean", "At 95th", and "At 5th" refer to the marginal effect of $\log(P_{ht-1})$ evaluated at the sample mean, 95th percentile, and 5th percentile of generator-level windspeeds. The regression also included control variables of quadratic meteorological variables, lagged aggregate wind generation, generating unit fixed effects, month-by-year fixed effects, and day-of-week fixed effects. Standard errors were clustered by EIA plant code and day-of-sample.

Table A.3: Price Responsiveness by Plant Characteristic

	No. of Turb.	Turb. Height	Wind Qual. Class
Q1	0.058**	0.096***	0.104***
	(0.0229)	(0.0125)	(0.0278)
Q2	0.104***	0.133***	0.098***
	(0.0155)	(0.0178)	(0.0072)
Q3	0.078***	0.075***	0.060***
	(0.0150)	(0.0170)	(0.0159)
Q4	0.101***	0.041**	0.065*
	(0.0113)	(0.0156)	(0.0407)

Notes: "No. of Turbines", "Turb. Height", and "Wind Qual." refer to number of turbines at the generator, the height of the turbines, and the wind quality classification of where the generator is located, respectively. For "No. of Turb." and "Turb. Height", Q1 - Q4 designate the price responsiveness of geneators in the 1st-4th quartiles of the given variable. For "Wind Qual.", Q1-Q4 refer to the price responsiveness of generators built in areas of wind quality 1 - 4. The regression also included control variables of quadratic meteorological variables, lagged aggregate wind generation, generating unit fixed effects, month-by-year fixed effects, and day-of-week fixed effects. Standard errors were clustered by EIA plant code and day-of-sample.

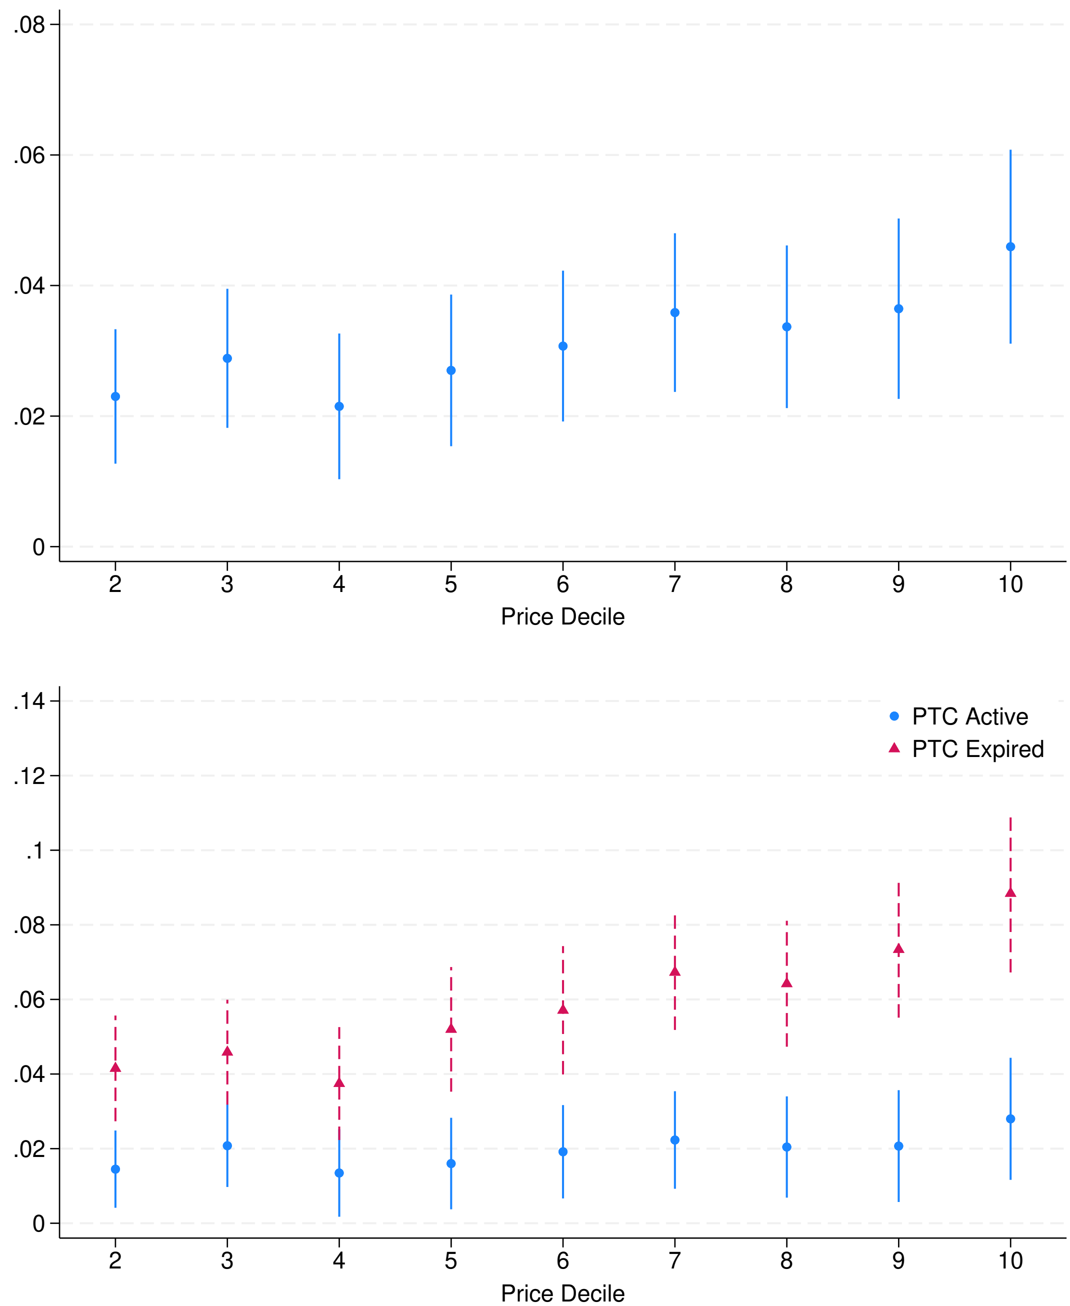

Figure A.1: Price Decile Capacity Factor Response

Notes: The plot displays the parameter estimates and 95% confidence intervals of the θ_k parameters given in equation (6), where the binning of the prices is by the decile of the price distribution. The bin indicators are based on the P_{ht-1} values and the horizontal axis refers to the decile value. The bottom panel gives results when the bin indicators are also interacted with the PTC-expiration indicator and thus estimates responses for PTC-active and PTC-expired facilities. The regression also included control variables of quadratic meteorological variables, lagged aggregate wind generation, generating unit fixed effects, month-by-year fixed effects, and day-of-week fixed effects. Standard errors were clustered by EIA plant code and day of sample.

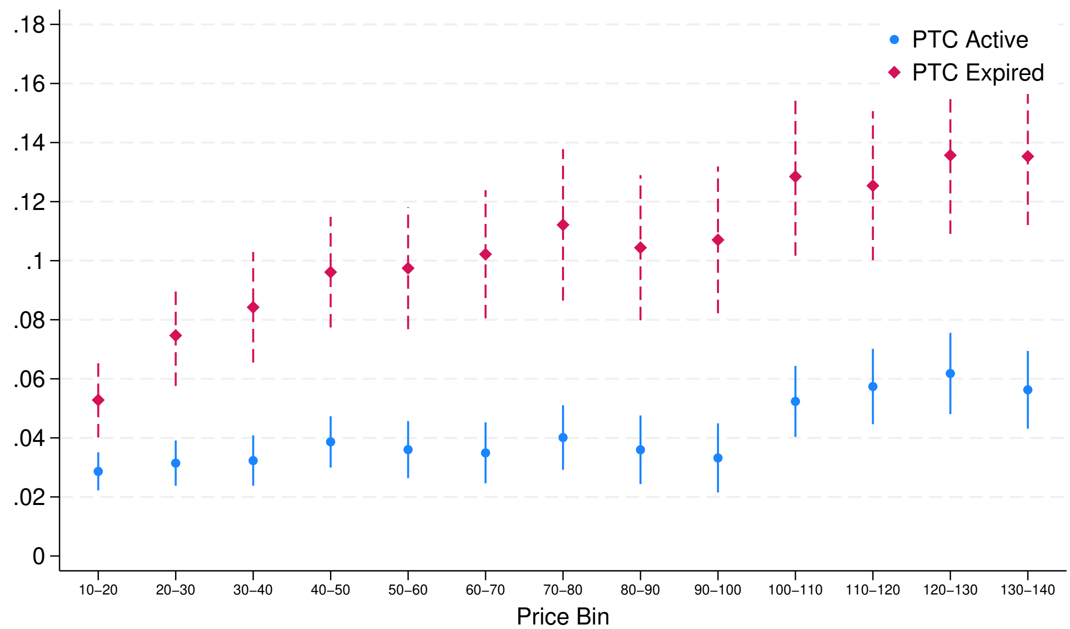

Figure A.2: Price Bin Capacity Factor Response - PTC Interaction

Notes: The plots show parameter estimates and 95% confidence intervals of capacity factor responses to binned prices based on the estimation of (6), with the indicator variable is 1 if the price is given in the price bin range given on the horizontal axis. The results are from a specification of (6) where the bin indicator variables are interacted with the PTC-expiration indicator, leading to separate estimates for PTC-active and PTC-expired facilities. The price-bin indicators are based off of P_{ht-1} values. The regression also included control variables of quadratic meteorological variables, lagged aggregate wind generation, generating unit fixed effects, month-by-year fixed effects, and day-of-week fixed effects. Standard errors were clustered by EIA plant code level.

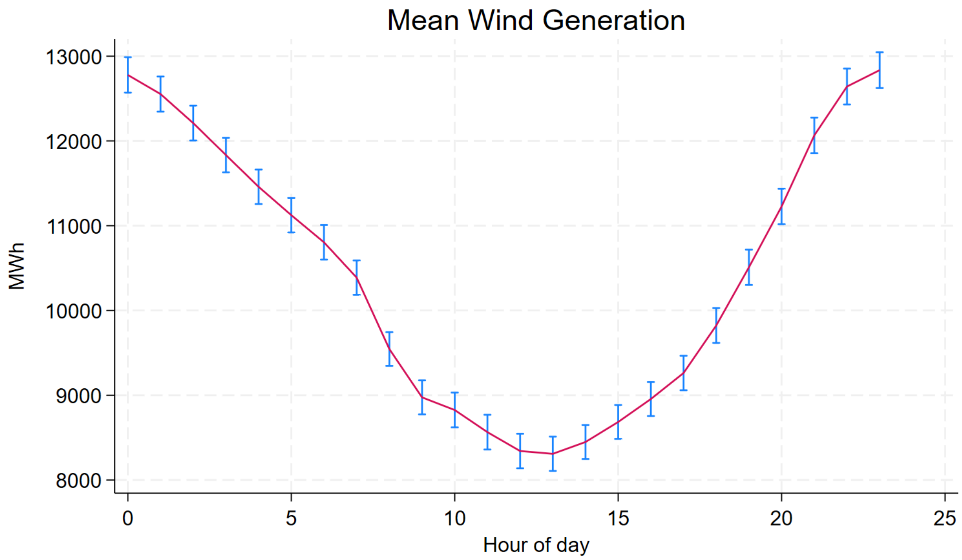

Figure A.3: ERCOT Mean Hourly Wind Generation

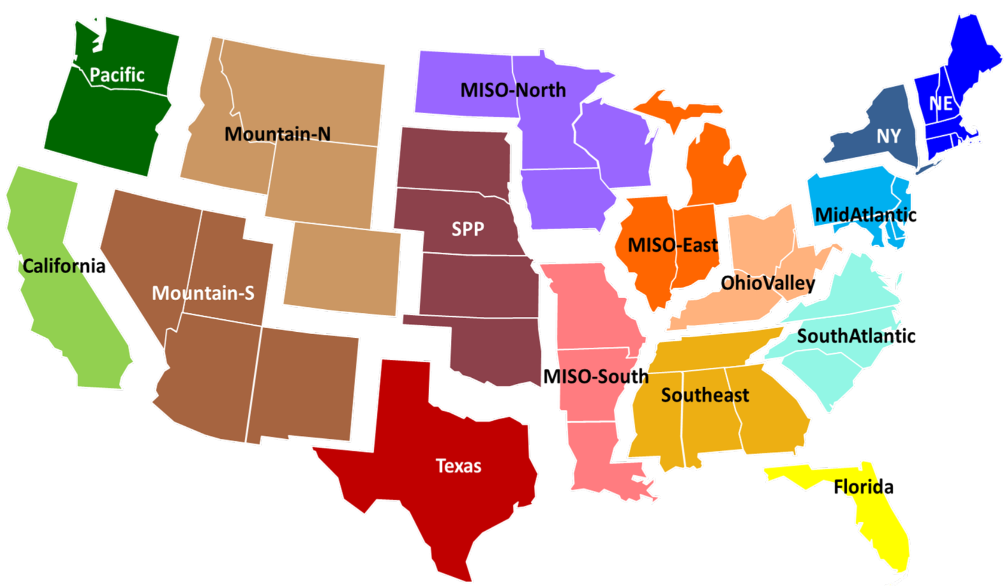

Figure A.4: US-REGEN Market Regions

B Appendix Discussion: MVPF of PTC with intensive margin response

Adapting the framework in Hahn et al. (2026), the basic MVPF (social value per dollar of government expenditure) calculation starts from:

$$MVPV = \frac{xd\tau + Vdx}{xd\tau + \tau dx} \tag{8}$$

where x is MWh, τ is the PTC per MWh, and V is the social value per MWh (GHG emissions and local emissions avoided). This can be expressed as:

$$MVPV = \frac{1 + \frac{V}{p}\varepsilon_{mwh}}{1 + \frac{\tau}{p}\varepsilon_{mwh}} \tag{9}$$

The elasticity term $\varepsilon_{mwh} = \frac{dMWh}{dp}\frac{p}{MWh}$ is the elasticity of MWh produced with respect to price p, inclusive of both intensive and extensive margin responses. In order to combine our estimated intensive margin elasticity of 0.14 with extensive margin estimates such as 1.13 from Hitaj (2013), we need to relate generation (MWhs) to capacity (MWs) as a function of prices. We can express the total lifetime discounted generation $MWh(p)$ over T years from given MWs of total capacity as:

$$MWh(p) = \sum_{t=0}^{T}(8760 * CF(p) * MW(p))e^{-rt} \tag{10}$$

$$= H * CF(p) * MW(p) \tag{11}$$

where H just converts a MW into the discounted lifetime MWh's produced (e.g $H = 116,762$ for $T = 20$ and $r = 0.05$).[30] Totally differentiating with respect to τ gives:

$$\frac{dMWh}{dp} = H[\frac{\partial CF}{\partial p} * MW + CF * \frac{\partial MW}{\partial p}] \tag{12}$$

[30] For a single megawatt of wind capacity, this implies about $1.5 million in present value revenue at an average price of $40 and a capacity factor of 0.33

Plugging into the elasticity formula, this gives:

$$\varepsilon_{mwh} = \frac{dMWh}{dp}\frac{p}{MWh} = H[\frac{\partial CF}{\partial p} * MW + CF * \frac{\partial MW}{\partial p}] * \frac{p}{MWh} \tag{13}$$

but $CF = MWh/(H * MW)$ and $\frac{\partial CF}{\partial p} = \frac{\partial MWh/\partial p}{H*MW}$, so we have:

$$\begin{aligned}
\varepsilon_{mwh} &= H[\frac{\partial MWh/\partial p}{H * MW} * MW + MWh/(H * MW) * \frac{\partial MW}{\partial p}] * \frac{p}{MWh} && (14)\\
&= \frac{\partial MWh}{\partial p} * \frac{p}{MWh} + \frac{\partial MW}{\partial p} * \frac{p}{MW} && (15)\\
&= \varepsilon_{int} + \varepsilon_{ext} && (16)
\end{aligned}$$

which just means that $\varepsilon_{mwh} = \varepsilon_{int} + \varepsilon_{ext} = 0.14 + 1.13 = 1.27$.

Back to the original MVPF formula, we can compare with a baseline estimate that only includes the extensive margin Hitaj (2013) elasticity. Reverse-engineering the values in Hahn et al. (2026), the MVPF can be expressed as:

$$MVPV = \frac{1 + 4.340 * \varepsilon_{mwh}}{1 + 0.244 * \varepsilon_{mwh}} \tag{17}$$

Using just the extensive margin elasticity $\varepsilon_{mwh} = \varepsilon_{ext} = 1.13$ from Hitaj (2013) gives an MVPF of 4.628, per the baseline Hahn et al. (2026) estimate. Including the intensive margin, increasing the elasticity to $\varepsilon_{mwh} = 1.27$, gives an MVPF of 4.971, implying roughly an additional \$0.34 of social benefit per \$1 of expenditure, on top of the extensive margin effects.[31] For comparison, if we instead use the IV specifications in Hitaj (2013) and corresponding extensive margin elasticity of $\varepsilon_{ext} = 0.509$, this gives an MVPF of 2.855 with just the extensive margin response, compared to an MVPF of 3.295 when the intensive margin response is included, adding an additional \$0.44 per dollar of expenditure.

[31] While an additional social benefit of \$0.34 might seem small, baseline estimates for the *total* social benefit of EV subsidies are similar at \$0.30 per \$1 of expenditure.

C Appendix Discussion: Dynamic Maintenance

In this section we develop a dynamic model of optimal maintenance. Again, we assume that maintenance affects the productivity constant of the wind generator in the same way as with the static model and that there is some cost of maintenance that varies with the maintenance level. In this dynamic formulation, we allow for past maintenance to, positively, affect the productivity constant such that $\lambda'(m_{t-j}) \geq 0$ and $\lambda''(m_{t-j}) \leq 0$ for any $j > 0$.

Given this, we form a profit maximization model over some planning horizon T as:

$$\max_{m_t} \pi = \sum_{t=1}^{T} \left[p_t \lambda_t(m_t, m_{t-1}, ..., m_{t-L}) v_t^3 - c(m_t) \right] \tag{18}$$

With the first order condition of:

$$\frac{\partial \pi}{\partial m_t} = \sum_{l=0}^{T} \left[p_t \frac{d\lambda_{t+l}}{dm_t} v_t^3 - c'(m_t) \right] = 0, \forall t \tag{19}$$

Can make different assumptions about L, $\frac{d\lambda_{t+l}}{dm_t}$, etc:

Consider a T-period profit maximization model for renewable generator operator. The operator takes electricity price, p_t as given, but can choose maintenance (m_t) at cost $c(m_t)$. Maintenance in period t lowers output in that period (q_t), but adds to the stock of maintenance in the next period (S_{t+1}) and that increases future output q_{t+1}. The maintenance stock decays at a constant rate, γ. Formally, the model is:

$$\begin{aligned} &\max_{m_t} \sum_{t=0}^{T} p_t q_t(m_t, S_t, w_t) - c(m_t) \\ &\text{s.t. } S_{t+1} = \gamma S_t + m_t \end{aligned} \tag{20}$$

With:

$$
\begin{aligned}
\frac{\partial q}{\partial m} = q_m < 0,\ \frac{\partial^2 q}{\partial m^2} = q_{mm} \leq 0 \\
\frac{\partial q}{\partial S} = q_S > 0,\ \frac{\partial^2 q}{\partial S^2} = q_{SS} \leq 0 \\
\frac{\partial q}{\partial w} = q_w > 0,\ \frac{\partial^2 q}{\partial w^2} = q_{ww} \geq 0 \\
\frac{\partial c}{\partial m} = c'(m) > 0,\ \frac{\partial^2 c}{\partial m^2} = c''(m) \geq 0 \\
0 \leq \gamma < 1
\end{aligned}
\tag{21}
$$

The Bellman equation of this model is:

$$
V_t(S_t) = \max_{m_t} \{p_t q(m_t, S_t, w_t) - c(m_t) + V_{t+1}(S_{t+1})\} \tag{22}
$$

The first order condition of this problem is:

$$
p_t q_m(m_t, S_t, w_t) - c'(m_t) + V'_{t+1}(S_{t+1}) = 0 \tag{23}
$$

The envelope theorem result is:

$$
V'_t(S_t) = p_t q_{S_t}(m_t, S_t, w_t) + \gamma V'_{t+1}(S_{t+1}) \tag{24}
$$

Stepping the envelope condition forward one period:

$$
V'_{t+1}(S_{t+1}) = p_{t+1} q_{S_{t+1}}(m_{t+1}, S_{t+1}, w_{t+1}) + \gamma V'_{t+2}(S_{t+2}) \tag{25}
$$

The Euler equation is determined by substituting V'_{t+1} from envelop condition into FOC and re-arranging:

$$
\begin{aligned}
c'(m_t) - p_t q_{m_t}(m_t, S_t, w_t) &= p_{t+1} q_{S_{t+1}}(S_{t+j}, w_{t+j}, m_{t+j}) + \gamma V'_{t+2}(S_{t+2}) \\
&= \sum_{j=1}^{T} \gamma^{j-1} p_{t+j} q_{S_{t+j}}(S_{t+j}, w_{t+j}, m_{t+j})
\end{aligned}
\tag{26}
$$

The Euler equation shows that the MC of maintenance at t, inclusive of the lost production cost, is equal to the full marginal benefit of maintenance at t.

From these conditions, it can be shown that $\frac{dm_t}{dp_t} < 0$ and $\frac{dm_t}{dp_{t+1}} > 0$. These responses would result in output having a positive response to prices and the greater the spread between p_t and p_{t-1}, the greater will be the observed response of quantity to prices. Furthermore, the model would imply $\frac{\partial}{\partial \gamma}(\frac{dq_t}{dp_t}) < 0$, the response of output to prices declines as the persistence in maintenance increases. The model would also imply that a level-shift in prices would result in a level shift in production in the same direction, consistent with the results of the PTC-effect in (Aldy, Gerarden, and Sweeney 2023) and (Ricks and Kay 2025). Likewise, a one-time price spike in this model, also leads to more generation in aggregate as opposed to a simple deference of maintenance that would imply no increase in actual production. The degree to which a price increase leads to a net increase in production depends, however, on the persistence of the maintenance effect on the stock of maintenance and the downtime cost of maintenance.