The causal impact of solar electricity generation on the electricity price

Nolan Ritter[1] and Sonja Rinne[2]

[1]German Institute for Economic Research (DIW Berlin), Mohrenstraße 58, 10117 Berlin, Germany, nritter@diw.de.

[2]Westfälische Wilhelms-Universität Münster, Scharnhorststraße 100, 48151 Münster, Germany, sonja.rinne@wiwi.uni-muenster.de.

June 10, 2016

Abstract

We estimate the causal effect of solar electricity generation on the electricity price in Germany, for the years 2011 through 2014. In our quasi-experimental approach, we split our sample. Observations without solar electricity generation are placed in the control group. The treatment group consists of observations with solar electricity generation. Coarsened exact matching identifies counterfactual observations in the control group that match the observations in the treatment group. Unconditional quantile regression is relied upon to relax the assumption that an average effect accurately describes the treatment across the unconditional distribution of the electricity price.

Keywords: electricity price, solar generation, causal effects, quantile regression

JEL Classification: C21, C23, I31, J16

1 Introduction

In Germany, renewable energies such as solar and wind receive considerable subsidies, at present in the form of feed-in tariffs. The motivation for the subsidization is to increase the share of non-CO_2 emitting electricity production. Under feed-in tariff schemes, each kilowatt hour of electricity produced is reimbursed at a fixed, technology specific nominal rate over the course of 20 years. The feed-in tariffs are financed by uplifting the cost unto the electricity bill of private households.

Although feed-in tariffs for newly installed systems have declined in recent years, the annual total subsidies under the feed-in tariff scheme have increased and reached an annual level of about 20 billion Euro in Germany. While information regarding the cost of the subsidies is readily available, causal estimates regarding their impact on electricity prices are not. Therefore, it is the aim of this paper to estimate the causal impact of solar electricity generation on the electricity price on the day-ahead market in Germany. Specifically, we are interested in finding out whether there exist circumstances for which the effect of the subsidization on the electricity cost exceeds the magnitude of the subsidies. To this end, we combine data obtained from the European Power Exchange (EPEX) with price information made available by DataStream.

In our quasi-experimental approach, we split our sample exploiting the exogeneity of solar insulation. In case that there is no solar electricity generation, the observations are placed in the control group. The treatment group consists of observations with solar electricity generation. We use coarsened exact matching to identify counterfactual observations in the control group with characteristics that match the observations in the treatment group. The relevant characteristics for matching are the prices for fuel as well as for CO_2 certificates, the level of wind electricity generation, the total system load, ant the capacity of electricity production out of service.

Our analysis goes beyond the mean by estimating unconditional quantile regression models. By applying this method, we are able to relax the assumption that an average treatment effect is an accurate description of the treatment effect across the unconditional distribution of the response. Moreover, we are able to compare the effect estimated for select points across this distribution. This kind of insight is barred to those who rely on mean regression.

Our results indicate that solar electricity generation has a statistically significant and economically relevant impact on the electricity price that depends on the level of the electricity price itself. At times, subsidies are smaller in magnitude compared to the reduction in spending on electricity from lower electricity prices. However, this is not yet the norm. If the magnitude of the reduction in the cost of electricity exceeds the magnitude of the subsidy, the subsidy is effective - although it changes the allocation of resources. Altogether, our approach contributes in several ways to the existing literature. It exploits the exogeneity of solar insulation. It relies on matching to arrive at causal estimates, and it goes beyond mean regression.

2 Literature

Contributions with regard to the impact of renewables on the electricity price have been made by Joskow (2011) and Borenstein (2012). Würzburg et al. (2013) who scans the existing literature on this question with an eye towards simulations as well as empirical analyses and Hirth (2013) who also provides a summary and comparison of studies on the market value of variable renewable energy sources such as solar and wind power. Morthost et al. (2010) focus in their review on energy price effects based on wind energy.

Altogether, the empirical literature finds a negative effect of renewable electricity production on the electricity price. Gelabert et al. (2011), for example, conclude from their multivariate regression model for the Spanish electricity market that a higher share of energy produced by renewables and cogeneration decrease the electricity price. Cludius et al. (2014) use a similar model for the German market including hourly electricity prices, solar generation, wind electricity generation, and the total energy load as control variables. They also identify a negative price effect due to a higher feed-in of wind and solar energy. In addition to the energy generation, they suggest to also incorporate commodity prices such as coal, gas, and CO_2 prices as explanatory variables when analyzing the electricity price.

Jónsson et al. (2010) point out that renewables might have non-linear effects on the electricity price. We take this finding seriously and interact all control variables with the level of the total load in the system. For analyzing the effect of renewables on the electricity price, empirical analyses often rely on time-series regression. Ketterer (2014), for example, uses a seasonally adjusted GARCH model explaining the daily electricity price by the expected wind generation, the total electricity load and their interaction for the German market. Her model results indicate that an increase in wind electricity generation reduces the electricity price, but at the same time increases price volatility.

Ziel et al. (2015) analyze the effect of wind and solar generation in addition to the total load to explain the hourly electricity price for the German market. Accommodating the fact that electricity demand is different on workdays compared to weekends and holidays, they conclude from their TARCH model that increases in renewable electricity generation reduces the electricity price.

The outline of our paper is the following. Section 2 describes the German electricity market with special respect to the electricity auction at the European Electricity Exchange (EPEX). Section 3 explains the identification strategy, while section 4 presents the matching method as well as the unconditional quantile regression. Sections 5 and 6 inform about the empirical application. In section 7 explains the policy implication of our findings, while section 8 summarizes and concludes.

3 The German electricity market

With an annual gross electricity consumption on the order of 589.8 Terawatt hours (TWh) in 2014, Germany has the largest electricity market in Europe (BMWi, 2015). At the same time, renewable electricity sources such as wind and solar accounted for 161.4 TWhs or 27.4%. This provides us with an opportunity to analyze the impact of solar electricity generation on electricity prices for a large and liquid market with a considerable renewable electricity generation. We contrast the effect of solar electricity generation on the electricity price with the cost of subsidizing solar electricity generation.

Long-term electricity contracts are traded bilaterally and no information regarding prices or volumes are available. Short-term trading, however, is well documented and takes place in the form of day-ahead auctions and intra-day continuous trading at the European Electricity Exchange (EPEX) in Paris, France. In order to participate, a registration at EPEX is necessary.

On any day, the first short-term auction to take place is the day-ahead or spot auction at 12 pm at which hourly contracts for all hours of the following day are traded. It is a sealed double auction in the sense that market participants bid to sell or buy electricity without any information regarding the bids of other market participants. Each bid requires information on whether the market participant intends to buy or sell, as well as the trading volume and the price. Electricity prices are constrained to a range between $-3,000$ and $3,000$ Euro per Megawatt hour (MWh). Orders must be placed before 12 pm, at which time the auction draws to a close and EPEX calculates the equilibrium price and equilibrium volume. Market participants are informed about the outcome of the auction at about 12:40 pm. All transactions clear at the equilibrium price. This equilibrium price is uniform across Germany.

Several theoretical models were established in order to show under which conditions sealed double auctions yield efficient market results. Most analyses focus on auctions that involve a single object, for example, a unique piece of art. In case of electricity, however, a large quantity is traded. Therefore, bidders submit a bid function rather than a single price (Wolfram, 1998).

As shown by Wilson (1985), a sealed double auction leads to efficiency if bidders are risk neutral, independently and identically distributed, symmetric, and the payment is a function of the bids alone. Moreover, the number of buyers and sellers within the auction needs to be sufficiently high.

To close the gap between the abstract mathematical proof of Wilson and the application to the specific auction at EPEX, "sufficiently high" needs to be quantified. Satterthwaite and Williams (1989) argue that buyers in a sealed double auction with only a few bidders have the incentive to understate their actual valuation which leads to inefficient market results. As the number of bidders rises, market pressure forces buyers to reveal their true preferences. By simulating values for double auctions from 2 to 12 bidders they show how quickly incentives to understate vanish. Kagel and Vogt (1993) modify the assumption in a way which leaves sellers with an incentive to bid below cost and buyers again to understate their demand. Nevertheless, the result remains unchanged. By increasing the number of bidders from 2 to 8, the computed values converge to the efficient results.

Daily repetition of auctions can be supportive of collusion compared to one-shot auctions (Rothkopf, 1999; Bower and Bunn, 2001). Still, the participants in the sealed bid auction do not receive any feedback on other bidders behavior (Ockenfels et al., 2008). Moreover, Friedman (1991) points out that strategic behavior in a sealed double auctions requires comprehensive knowledge of the market and its participants. Because of this, he argues that it can be assumed that bidders neglect this option and "play a game against nature". In addition to that, the considerable level of intermittent renewable energy sources increases the complexity of the day-ahead electricity market at EPEX, thereby increasing the risk of loosing out when trying to game the market.

Another important aspect which needs to be addressed is whether bidders are symmetric since electricity markets usually face a certain amount of concentration on the supply side. The incentive to withhold capacity declines with the market size but increases with the market share of the bidder. It is noteworthy that the overall electricity production of the supplier is not equal to their share in the day-ahead market since especially large suppliers sell most of their production in advance (Ockenfels et al., 2008).

Wolfram (1998) for example analyzes daily bids by generating companies in England and Wales from 1992 to 1994 and finds evidence that a larger supplier submits bids with a higher mark-up compared to its smaller competitors. Borenstein et al. (2002) identifies inefficiencies in California's wholesale electricity market where market power was likely to be exercised in peak demand periods between 1998 and 2000. For the German market, Müsgens (2006) finds evidence for higher electricity prices due to market power in some months in the period from 2000 to 2003.

Whereas, at the end of 2003 there were 25 market participants who traded at EPEX, this number has risen to 232 EPEX exchange members at the end of 2015. Therefore, we argue that the number of market participants is sufficiently high so that market participants reveal their true preferences. In this case, the observed market outcomes are not influenced by unobservable and changing behavior.

Figure 1: *Supply and demand curves*

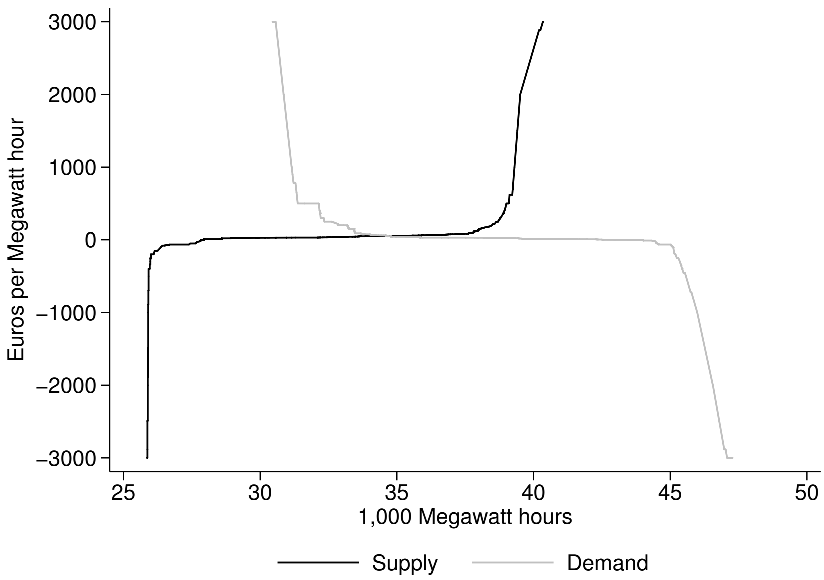

Figure 1 indicates the aggregate supply and demand bids from January 10, 2014 for 12 pm delivery the next day (EPEX 2015). The supply and demand curves found on the electricity market differ strongly compared to other commodities. For one, there are bids to produce electricity at negative prices. These offers are made by inflexible power plants designed for long-term constant output. For these power plants, it is cheaper to produce at negative prices for a limited amount of time in comparison to deviations from their optimal production profile. Negative price offers can also be found on the demand side. At negative prices, for example, large compressed air storage may discharge their load without producing electricity in favor of additional electricity consumption.

4 Theoretical model

Figure 2 provides for a stylized model of supply and demand. The figure includes two supply functions. One indicates the marginal cost of electricity production when there is no solar electricity generation. The second supply function indicates the marginal cost at times of solar electricity production. At times of no solar generation, the intersection of supply and demand occurs at point G, associated with price E and quantity J. The cost associated with producing the electricity equals the area defined by $\overline{0DGJ}$, the producer surplus is $\overline{DEG}$, and the consumer surplus is $\overline{EFG}$.

Figure 2: *Market equilibrium*

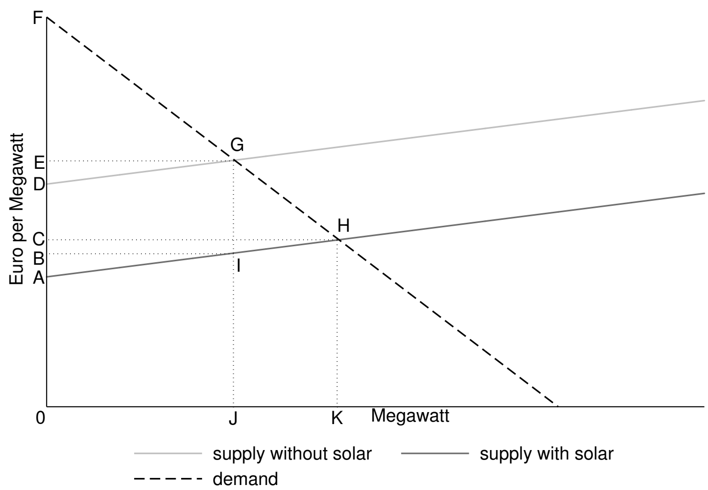

At times of solar electricity generation, the supply function is shifted to the right, because solar generates electricity at zero marginal cost. With a given demand function, the new market equilibrium (point H) clears at a lower price (C) and at a higher quantity (K) compared to before. All else equal, increases in solar electricity generation increase the consumer rent from $\overline{EFG}$ to $\overline{CFH}$. Altogether, the rightward shift in the supply function lowers the price from G to I, while the additional demand increases the price.

Figure 3: *Market equilibrium holding demand constant*

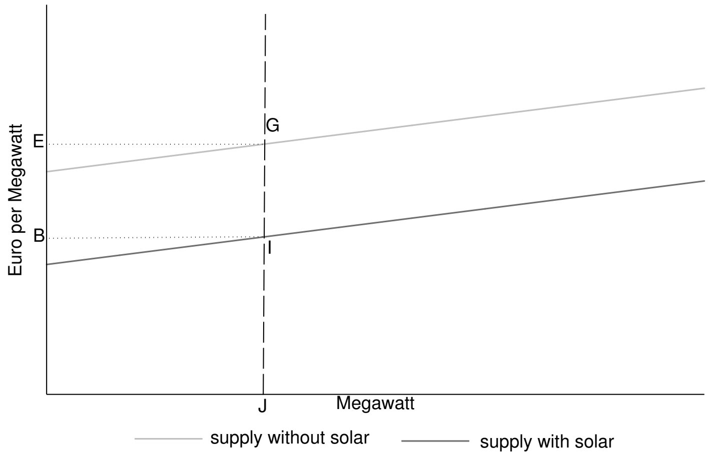

We are, however, interested in isolating the effect of solar electricity generation on the electricity price holding all else constant. To this end, we compare market outcomes with similar characteristics including the equilibrium quantity. In the context of a stylized supply and demand diagram, this is represented by a vertical demand function that fixes the equilibrium quantity at point J.

5 Methods

We pursue estimating the impact of solar electricity production on the electricity price by combining the data with coarsened exact matching and unconditional quantile regression. The matching process guarantees that there are no systematic differences between treatment and control groups with the exception of the treatment itself. Unconditional quantile regression is applied to test whether the impact of the treatment depends on the level of the electricity price. The intuition is that at times of high electricity prices the effect of solar electricity production has a higher impact compared to times of low electricity prices.[1]

5.1 Matching method

The ultimate aim of econometric analyzes is to identify causal effects. Nevertheless, causality itself is a theoretical concept independent of statistical methods (Ho et al., 2007). In this paper, we refer to the causality identified employing contrived experiments (Splawa-Neyman et al., 1990; Fisher, 1935; Cox, 1958) compared to causality from predictions (Granger, 1969; Sims, 1972). For overviews on causality, we refer the reader to Cartwright (2004) or Heckman (2005, 2008).

In the most simplest of cases, an experiment consists of observation units i randomly assigned to two groups in addition to a binary treatment. For the treatment group, the binary treatment indicator is $T_i = 1$, while the indicator is $T_i = 0$ for the observation units in the control group. In case that there is a causal effect of solar electricity generation on the electricity price on the day-ahead market, it is impossible to contrive any kind of randomized controlled experiment capable of revealing a the true effect. Randomizing solar insulation is impossible.

But it is possible to approximate a randomized controlled experiment very closely. To this end, we identify observation units from the control group that are similar with regard to all relevant characteristics with the exception of the treatment itself: solar electricity generation. Ideally, we would want to observe the same observational units in the counterfactual situation of employment so that it would be guaranteed that the unobservable characteristics would be identical, too. Observing the counterfactual is, of course, impossible and poses "the fundamental problem of causal inference" (Holland, 1986). We employ coarsened exact matching described by Ho et al. (2007); Iacus et al. (2011b,a) because it is superior to other means of matching such as propensity score matching or Mahalanobis matching.

[1] Similar descriptions of the methods applied in this paper can be found in Ritter (2016a) and Ritter (2016b).

The higher the similarity after matching, the closer the quasi-experiment approximates a carefully randomized controlled experiment. In perfectly balanced groups, a t-test suffices to ascertain possible differences in the outcome. Should imbalances remain after matching, a model needs to control for the differences. Altogether, successful matching considerably reduces the influence of the specification of the model on the results. (Iacus et al., 2011a) refer to this as model dependence.

Let T_i inform about the treatment with

$$T_i = \begin{cases} 1 & \text{for treated } i \\ 0 & \text{for untreated } i \ , \end{cases} \tag{1}$$

and let the sample of size n be a drawn from a population N with $n \leq N$. Each n belongs either to the treatment group with the potential outcome $Y_i(1)$ or the control group with the potential outcome $Y_i(0)$, thus

$$Y_i = T_i Y_i(1) + (1 - T_i) Y_i(0) \ . \tag{2}$$

One important difference between the treatment and the control group is that all observations that qualify as $Y_i(1)$ are known, while the observations that have similar characteristics to the ones in the treatment group need to be identified using the matching procedure.

In addition to the response Y and the treatment variable T, the k-dimensional matrix $\boldsymbol{X} = (X_1, X_2, \ldots, X_k) = [X_{ij}, i = 1, \ldots, n; j = 1, \ldots, k]$ holds any remaining controls. The treated observational units are indexed by $\mathcal{T} = \{i : T_i = 1\}$, and the number of treated units is $n_T = \#\mathcal{T}$. Accordingly, the control units are indexed following $\mathcal{C} = \{i : T_i = 0\}$, while their number is $n_C = \#\mathcal{C}$. Thus, the sample size is $n = n_T + n_C$. The aim of the matching process is to match a treated unit $i \in \mathcal{T}$ with covariates $\boldsymbol{X}_i$ to a control unit $l \in C$ with similar covariates $\boldsymbol{X}_l$, so that the dissimilarity d is minimized to a degree that it is of little consequence for the analysis or $d(X_i, X_l) \simeq 0$. In case of exact matching, one aims for a dissimilarity on the order of $d(X_i, X_l) = 0$.

Coarsened exact-matching involves three steps. First, continuous variables in $\boldsymbol{X}$ are coarsened into categories $H(\boldsymbol{X})$. Second, exact matching is performed on the categories by sorting the observations into a number of strata $s \in S$ equaling the number of unique values of $H(\boldsymbol{X})$. Third, all strata s are discarded which are only populated by untreated observations. Moreover, the strata s are kept in which exist observations from the treatment and the control group.

There are two choices with respect to strata s which are only populated by treated observations. On the one hand, one can keep these observations and use extrapolated values of the control units. Another option is to discard these observations to preserve a tight combination of treated with control observations that qualify as counterfactuals. **?** argues that "in many observational studies the sample sizes are sufficiently large that sampling variances of estimators will be small, the sensitivity of estimators to biases is the dominant source of uncertainty." Following Rubin's argument, we remove observations from the treatment group for which there are no counterfactuals available.

For successful matching, the imbalance between the treatment and the control group needs to be smaller after matching compared to before. To measure the imbalance, Iacus et al. (2011b) suggest to use

$$\mathcal{L}_1(f, g; H) = \frac{1}{2} \sum_{l_1 \dots l_k \in H(\boldsymbol{X})} \left| f_{l_1 \dots l_k} - g_{l_1 \dots l_k} \right| , \tag{3}$$

with f^m and g^m being frequencies for the for the uncoarsened variables in the treatment and the control group. $H(\boldsymbol{X}) = H(X_1) \times \cdots \times H(X_k)$ describes the categories into which the variables have been coarsened. While $\mathcal{L}$ has no interpretation, a smaller imbalance after matching m with $\mathcal{L}(f^m, g^m) \leq \mathcal{L}(f, g)$ indicates a successful matching.

Iacus et al. (2011b) compare coarsened exact matching to propensity score matching and Mahalanobis matching. They present evidence that coarsened exact matching is superior to both. In addition to the empirical findings, Iacus et al. (2011b) present a theoretical argument against propensity and for coarsened exact matching. Turning to Figure 4, the reader is presented with the distributions of two variables. The monomodal distribution is for the variable from the treatment and the biomodal distribution is for the control group. In both cases the expected value is 1. Propensity score matching focuses on the mean of variables. Because the mean value of both distributions is the same, no observations would be dropped. Coarsened exact matching is not limited to the mean but rather leads to very similar distributions in the control compared to the treatment group after matching as would be the case if stratified random-sampling had taken place.

Figure 4: *Mono- and bimodally distributed variables with identical mean*

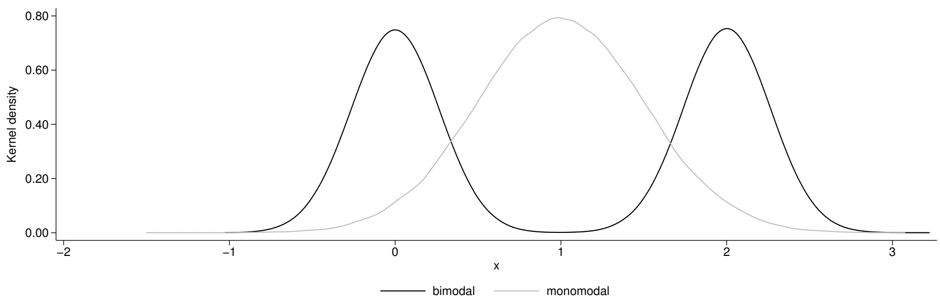

5.2 Conditional quantile regression

Diminishing returns are an established concept in economics, electricity production is no exemption in the sense that the production cost for each Kilowatt-hour increases with the level of electricity production. Thus, at times of high electricity demand the price of electricity is high, especially when renewable electricity production is low and conventional electricity generation capacities become increasingly scarce. In such a situation, additional electricity production from renewable sources with zero marginal cost such as solar and wind should have a larger impact compared to times in which prices are low. While mean regression allows for the estimation of an average effect, quantile regression suggested by Koenker and Bassett (1978) allows to estimate the effect of any control at any point in the conditional distribution of a continuous response. Thus, it is possible to find out whether the impact of a control depends on the level of the response itself - in this case whether the impact of renewables on the electricity price depends on the level of the electricity price.

Determining the quantile of a variable, say the median, can be achieved in two ways. First, one can sort the variable of interest in an ascending or a descending order and count until ones finds the value for which half of the observations are above and half are below. Second, one can employ an analytical approach that minimizes the absolute sum of residuals. Koenker and Hallock (2001) provide for an excellent overview article which exemplifies how to determine the median by solving

$$y_{\text{median}} = \min_{\xi \in \mathbb{R}} \sum_{i=1}^{n} \rho_\tau (y_i - \xi) \ , \tag{4}$$

where ρ_τ is the absolute value function. In case of the median, ρ_τ is symmetrical to ensure that there are as many observations below and above. Any other quantile can be found by tilting the absolute value function.

So far, it was described how to arrive at the unconditional median of a variable. A simple thought experiment by Koenker and Hallock (2001) shows how to transfer the train of through used to derive the conditional expectation $E(Y|X)$ from the unconditional mean $E(Y)$ to quantile regression. Given a random sample $[y_1, y_2, \ldots, y_n]$ the sample mean is calculated following

$$E(Y) = \min_{\mu \in \mathbb{R}} \sum_{i=1}^{n} (y_i - \mu)^2 \ . \tag{5}$$

By substituting μ with a parametric function $\mu(x_i, \beta)$, the conditional expectation function is found by solving

$$E(Y|X) = \min_{\mu \in \mathbb{R}} \sum_{i=1}^{n} (y_i - \mu(x_i, \beta))^2 \ . \tag{6}$$

The same train of thought applied to quantile regression yields

$$E(Y_\tau|X) = \min_{\xi \in \mathbb{R}} \sum_{i=1}^{n} \rho_\tau (y_i - \xi(x_i, \beta)) \ , \tag{7}$$

where τ indicates the quantile of interest.

5.3 Unconditional quantile regression

Let $dq_\tau(T)/dT$ describe the direct effect of the treatment t on the τth quantile q_τ of the marginal distribution $F_Y(y)$, where $T = 1$ if the observation is from the treatment group and $T = 0$ otherwise. The unconditional quantile estimator suggested by (Firpo et al., 2009) for the treatment is

$$dq_\tau(T)/dT = \frac{E[Y > q_\tau|T = 1] - E[Y > q_\tau|T = 0]}{f_Y(q_\tau)} \ . \tag{8}$$

It is interpreted as the impact of the treatment on the τth unconditional quantile of the response, for example, the impact of a unit of solar electricity production in the 10th quantile of the unconditional distribution of the electricity price. The more frequently applied conditional quantile regression as suggested by Koenker and Bassett (1978) described above, estimates the impact of the treatment on the conditional distribution.

$$\beta_\tau = F_Y^{-1}(\tau|T = 1) - F_Y^{-1}(\tau|T = 0) \ . \tag{9}$$

In case of conditional quantile regression, adding additional controls changes the distribution of the response analyzed, while in case of unconditional quantile regression, the distribution of the response analyzed remains the same. Moreover, the interpretation of any unconditional quantile is straightforward and can easily be calculated. This, however, does not apply to the conditional quantiles which depend on the control variables included.

The estimator suggested by Firpo et al. (2009) relies on the concept of the influence function $\mathrm{IF}(Y; v, F_Y)$ which describes the impact of an observation on a distributional statistic $v(F_Y)$. The τth quantile of the influence function IF is known to be

$$\mathrm{IF}(Y; q_\tau, F_Y) = (\tau - \mathbb{1}\{Y \leq q_\tau\})/f_Y(q_\tau) \ . \tag{10}$$

By adding $v(F_Y)$ to the influence function, Firpo et al. (2009) generate what they refer to as the re-centered influence function (RIF), which has the desirable property that its expected value equals $E(\text{RIF}) = v(F_Y)$. Thus, the RIF of the τth quantile is

$$\text{RIF}(Y; q_\tau, F_Y) = q_\tau + \text{IF}(Y; q_\tau, F_Y) \ . \tag{11}$$

The unconditional quantile estimator by Firpo et al. (2009) can be estimated using ordinary least squares transforming the observed response variable according to equation 11 in a three step procedure estimating $\hat{q}_\tau$, $f_Y(q_\tau)$, and the average marginal effect $dE[Y > q_\tau|T]/dT$. First, estimate the sample quantile q_τ according to Koenker and Bassett (1978):

$$\hat{q}_\tau = \underset{q}{\text{argmin}} = \sum_{i=1}^{N} (\tau - \mathbb{1}\{Y_i - q \leq 0\}) \cdot (Y_i - q) \ . \tag{12}$$

Second, estimate the density $\hat{f}_Y(\cdot)$ using the kernel density estimator

$$\hat{f}_Y(\hat{q}_\tau) = \frac{1}{N \cdot b} \cdot \sum_{i=1}^{N} \mathcal{K}_Y \left(\frac{Y_i - \hat{q}_\tau}{b} \right) \tag{13}$$

with $\mathcal{K}_Y(z)$ being a kernel function with a positive bandwidth b. Third, estimate the average marginal effect $dE[Y > q_\tau|T]/dT$ using OLS.

6 Data

The data underlying the analysis originates from four sources and covers the period from 01.01.2011 - 31.03.2014. For all variables with the exception of the coal price, the gas price, and the exchange rate, the data is made available under the transparency regulations of the European Union. The hourly electricity spot market price and volume are from the EPEX website. The production of solar and wind energy as well as the out of service capacity are published on the European Energy Exchange (EEX) transparency website. Coal prices, gas prices, and CO_2 prices were obtained using DataStream, while the exchange rate was obtained via the German central bank's website. Information obtained from DataStream is proprietary. However, we are happy to provide the full data set underlying the analysis to any requesting party proving having access to DataStream.

6.1 Descriptive statistics

Descriptive statistics by treatment group for the full data set before matching are provided in Table 1. The data indicates that the average electricity price is higher for treated observations. There is also a considerable difference in the average total load, with approximately $58,300$ MW for the treated and $49,800$ MW for the non-treated. Other control variables include the amount of solar and wind generation, the total system load, the generation capacity out of service, in addition to the prices for inputs in electricity generation, such as gas and coal. We also control for the price of carbon certificates. Electricity producers have to surrender a number certificates depending on the amount of their carbon emissions. Thus, the price for certificates informs the electricity price. We also test the hypothesis that the mean values of the observable characteristics are identical comparing the treatment to the control group. The third column of Table 1 presents the test results that all reject the hypothesis of equality at any conventional level of significance.

Table 1: *Descriptive statistics for full sample (N=12,384)*

Variables	**control group** (N=5,548)	**treatment group** (N=6,836)	**t-tests** (p-values)
electricity price (Euro / MWh)	30.434 (16.586)	39.817 (15.572)	0.000
solar (1,000 MWh)	0.000 (0.000)	5.473 (5.770)	0.000
wind (1,000 MWh)	6.658 (5.078)	5.781 (4.872)	0.000
out of service (1,000 MW)	12.641 (4.668)	13.958 (5.149)	0.000
co2 (Euro / ton)	4.973 (1.177)	4.693 (1.153)	0.000
load (1,000 MWh)	49.777 (7.929)	58.256 (7.808)	0.000
coal (Euro / ton)	59.061 (4.466)	58.139 (4.406)	0.000
gas (Euro / MWh)	25.200 (1.229)	25.070 (1.171)	0.000

Standard deviation in parentheses.

Table 2: *Probit results explaining assignment to treatment group*

Variables	**Coefficients** N=12,384
wind (MW)	−0.019*** (0.003)
load (MW)	0.139*** (0.002)
CO2 price (Euro / t CO2)	−0.900*** (0.013)
coal price (Euro / ton)	0.040*** (0.005)
gas price (Euro / MWh)	−0.243*** (0.017)
out of service (MW)	0.190*** (0.005)
intercept	−5.607*** (0.313)

Standard errors in parentheses. ***, **, * indicates statistical significance at the 0.1%, 1%, and 5% level.

To further test for systematic differences between the treatment and the control group, we estimated a probit model. The results in Table 2 provide for overwhelming evidence for systematic differences between the two groups. All control variables used in the analysis in addition to the treatment are significant at the 0.1% level. This violates the principle that there should be no systematic differences between control and treatment groups with the exception of the treatment itself.

6.2 Identification strategy

This paper exploits the natural experiment character of solar electricity generation, which takes place only during the light hours of the day. Besides, at night time, the equipment is identical to day time. Often this also true for the fuel prices, and the wind electricity generation. We place hours with solar electricity generation into the treatment group, while the hours without serve as the control group. Using coarsened exact matching (Ho et al., 2007; Iacus et al., 2011a,b; King and Nielsen, 2015), we find similar observations regarding the wind electricity production, the capacity out of service, the price of carbon emissions, the total load, as well as for the prices of coal, and gas.

Table 3: *Descriptive statistics after matching (N=1,450)*

Variables	**control group** (N=725)	**treatment group** (N=725)	**t-tests** (p-values)
electricity price (Euro / MW)	40.138 (18.630)	36.842 (18.378)	0.001
solar (MW)	0.000 (0.000)	2.337 (3.814)	0.000
wind (MW)	5.836 (5.031)	5.809 (5.090)	0.916
out of service (MW)	13.167 (5.188)	13.221 (5.157)	0.843
co2 (Euro / ton)	4.834 (1.109)	4.835 (1.109)	0.993
load (MW)	54.074 (10.436)	54.276 (10.382)	0.712
coal (Euro / ton)	58.762 (4.477)	58.761 (4.473)	0.998
gas (Euro / Mwh)	25.096 (1.249)	25.098 (1.246)	0.974

Standard deviation in parentheses.

Turning to Table 3, the reader learns about the data set by treatment group after the matching process. The data set consists of two groups of 725 observations each. According to Table 3, for example, the average electricity price is higher for the control group compared to the treatment group. In the full data set described in Table 1, however, the price was higher for the treatment group compared to the control group.

Table 4: *Probit results explaining assignment to treatment group*

Variables	**after matching** N=1,450
wind (MW)	−0.001 (0.007)
load (MW)	0.005 (0.005)
CO2 price (Euro / t CO2)	0.000 (0.033)
coal price (Euro / ton)	0.004 (0.012)
gas price (Euro / MWh)	−0.017 (0.042)
out of service (MW)	0.010 (0.012)
intercept	−0.205 (0.741)

Standard errors in parentheses. ***, **, * indicates statistical significance at the 0.1%, 1%, and 5% level.

Further, the control group is characterized by the absence of solar production while the average level in the treatment group was $2,337$ MW. In both groups, wind generation was a little over $5,800$ Megawatt, while about $13,000$ MW of capacity were out of service for scheduled maintenance. The average carbon price was 4.834 Euro per ton in both sub samples, while the coal price was 58.76 Euro per ton. The average price for gas in the sample was 25.10 Euro per MW.

To test the success of the randomization, a probit model was estimated. It confirms that by employing coarsened exact matching, we managed to distill two groups of observations for which no significant differences can be discerned (see Table 4). The matched data set approximates a genuine experiment as closely as possible. Thus, removing the observations that cause the systematic differences between the treatment and the control group contributes in a considerable way to identify causal effects.

The similarity of the individual variables achieved by coarsened exact matching extends beyond the mean. This technique accounts for higher moments, and, thus, is superior to matching techniques that do not. For

Figure 5: *Density distribution of wind production by treatment*

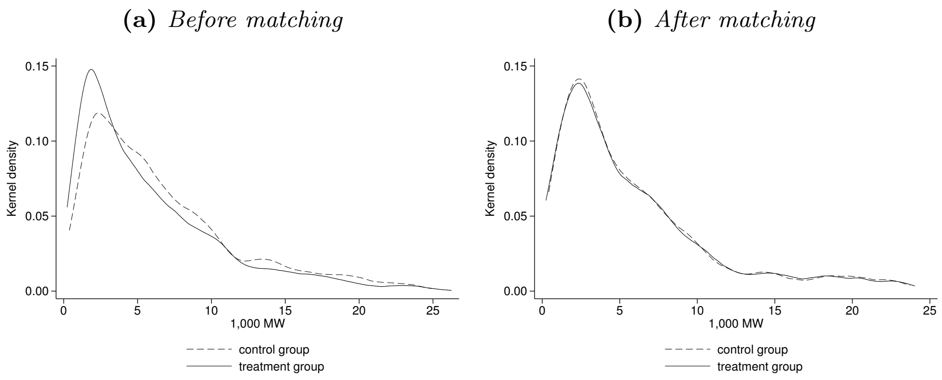

example, Figure 5 indicates the density distribution of wind electricity production. The left panel depicts the distribution before the matching process, while the right panel plots the distribution after the matching process. Before the matching, the treatment group includes more observations with wind production levels below 5,000 Megawatt hours (MWh) compared to the control group. Turning to the right panel in Figure 5, the reader learns that the matching process greatly increases the similarity in the density distributions.

Figure 6: *Density distribution of total load by treatment*

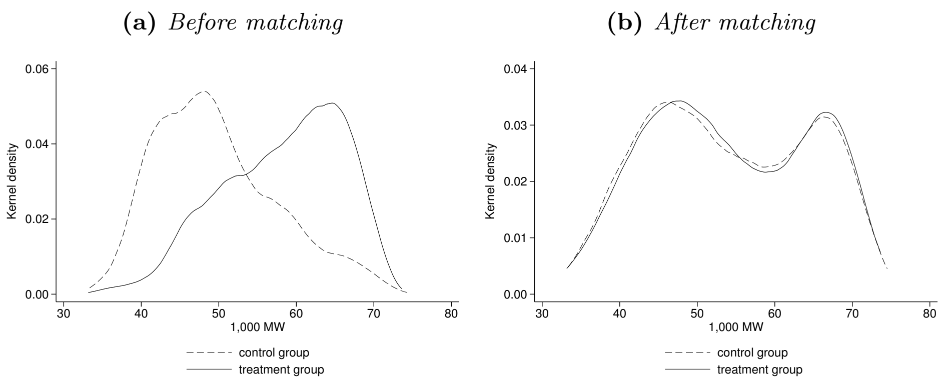

The improvement in the similarity with respect to the total load is even greater. This is indicated in Figure 6. Before the matching (6a), the highest density in the total load is found between $60,000$ and $70,000$ MW in the treatment group, while the highest density in the control group occurs between 40,000 and 50,000 MW. This pronounced difference in the total load indicates that the intersections of supply and demand are considerably higher for the treatment group compared to the control group before matching. After the matching process, the total load in both groups is virtually identical for many points in their distributions.

Before the matching, there already existed a high degree of similarity in the distributions of the CO_2 price (see Figure 7a). Nevertheless, the probit model found that the CO_2 price was a significant determinant of treatment status (see Table 2). After the matching, however, both groups share virtually identical density distributions.

Figure 7: *Density distribution of CO2 prices by treatment*

(a) *Before matching* **(b)** *After matching*

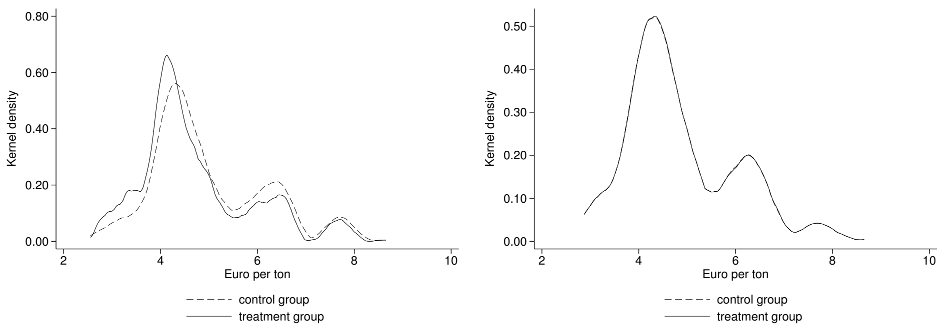

Regarding the price of coal (Figure 8), there were more observations with prices below 55 Euro per ton and between 59 and 62 Euro per ton in the treatment group. Between 56 and 59 Euros per ton and between 62 and 67 Euros per ton there were more observations for the other group. After the matching, the distributions of coal prices are again virtually identical.

Table 2 also indicated significant differences between the treatment and the control group with respect to the gas price. According to Figure 9a, these differences occur in the upper part of the distributions. After the matching process, the distributions are again very similar.

Figure 8: *Density distribution of coal prices by treatment*

(a) *Before matching* **(b)** *After matching*

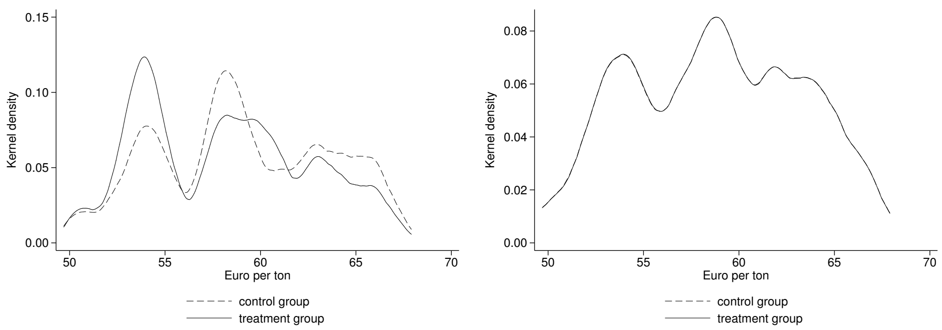

Figure 9: *Density distribution of gas prices by treatment*

(a) *Before matching* **(b)** *After matching*

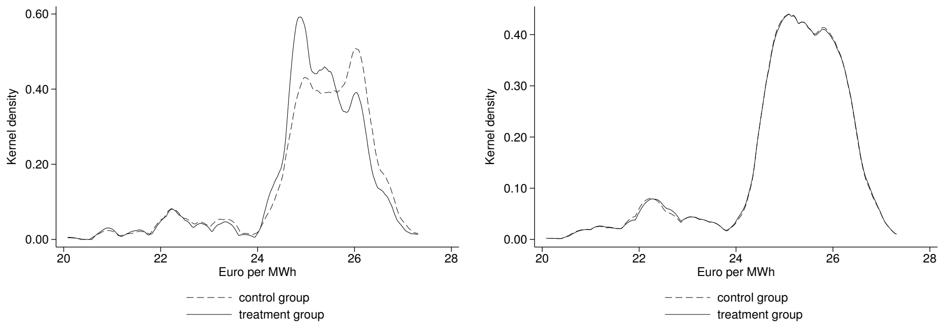

Turning to the electricity generation capacity out of service, the reader learns that before the matching the control group had many more observations below about 15,000 MW. At the same time, the control group had less observations for capacity out of service above 15,000 MW. After the matching, the distributions are again near identical.

Matching also had an influence on the dependent variable, the electricity price, although it was not part of the matching process. Figure 11 shows the distributions of the electricity price before and after the matching. For example, the lowest price found in the treatment and control group is approximately −50 Euro after matching. Before, the data set contained electricity prices as low as −200 Euro per Megawatt hour. Therefore, the analysis at

Figure 10: *Density distribution of out of service capacity by treatment*

(a) *Before matching* **(b)** *After matching*

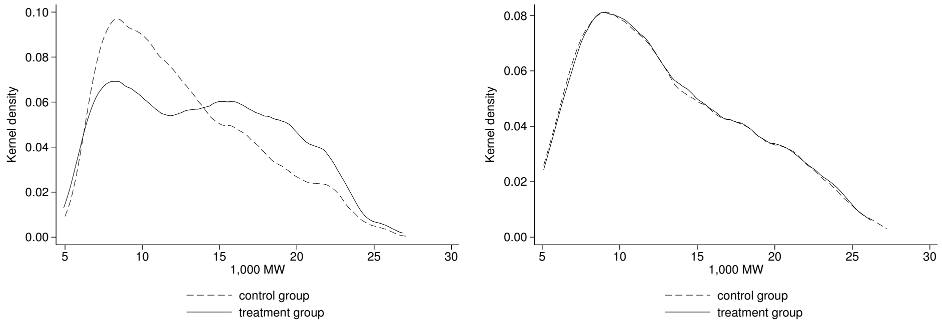

hand cannot make statements about prices lower than −50 Euro. However, there are only few observations on negative electricity prices.

Figure 11: *Electricity price before and after matching*

(a) *before* **(b)** *after*

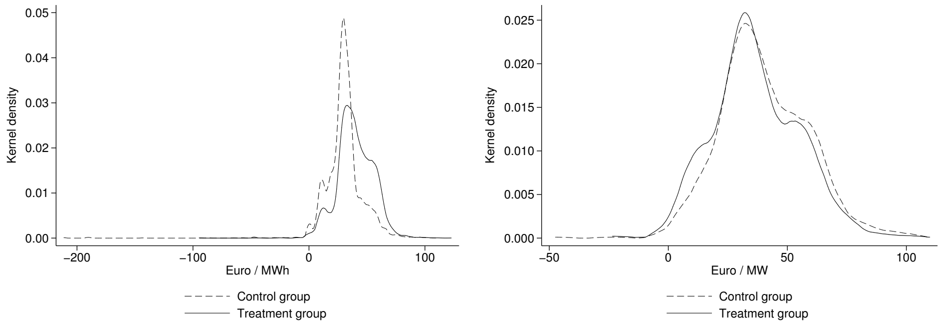

Scrutinizing Figures 5 through 10, the reader learns that the range of observations is the same before and after the matching. Therefore, the analysis at hand applies to all states of the control variables with the exception of the price. Comparing the descriptive statistics before and after matching on solar electricity production (Tables 1 and 3), one might arrive at the conclusion that the data on solar electricity generation is quite different before and after the matching. Figure 12 provides for strong evidence against this notion. In fact, the range of the observations is similar. Moreover, the shape of the distribution is also similar, with some deviation in the lower part of the distribution. Therefore, the analysis is also applicable to all states of solar electricity generation observed in the observation period. Altogether, there is strong evidence that our analysis has external validity.

Figure 12: *Solar electricity generation before and after the matching for the treatment group*

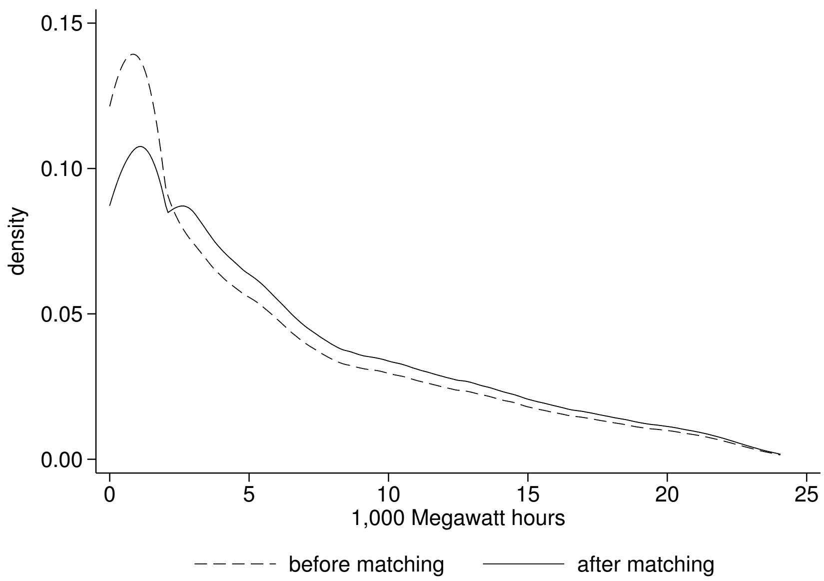

Having shown that the matching has lead to very similar observations comparing the treatment and the control group, there remains the question whether more observations could have been preserved in the interest of smaller standard errors in the following regressions. When applying coarsened exact matching, one has to chose amongst 4 different algorithms that coarsen the continuous variables into categorical ones. This choice might influence the number of matches found. Table 5 indicates the size of the post matching data set depending on the algorithm chosen.

Table 5: *Impact of the algorithm choice on the number of observations*

Algorithm	Observations	Systematic differences
Sturge's rule	$1,450$	no
Scott's rule	678	no
Friedman-Diaconis rule	540	no
Shimazaki-Shinomoto's rule	$1,210$	no

The column "systematic differences" indicates whether a probit model found significant differences between the treatment and the control group.

While all algorithms lead to data sets devoid of systematic differences between the control and the treatment group, choosing Sturge's rule preserves the highest number of observations. Therefore, we apply Sturge's rule in the matching process.

7 Empirical Application

We estimate the following model:

$$y = \beta \times \text{treatment} + \delta \times \text{load} + \gamma \times \text{treatment} \times \text{load} + X\phi + \epsilon \ . \tag{14}$$

In the context of this paper, the treatment variable refers to the level of solar electricity production. It appears in its interacted form with the total level of load. The other control variables accounted for in Table 3, such as the prices of coal and gas, are included in the matrix X. ϵ indicates an error term. The decision to interact the control variables is straightforward: The supply follows a step function with cheap electricity production provided by renewables at zero marginal costs and slightly higher costs for nuclear and lignite power plants. Compared to lignite and nuclear, hard coal operated plants face higher costs. Gas turbines are even more expensive.

Since the market clearing electricity price is determined by the costs of the last unit, the hard coal price only affects the market clearing given that the last unit is provided by hard coal. The same logic applies to gas. Thus, the interaction is necessary to guarantee that the model only ascribes an effect to the fuel used to produce the last unit.

Table 6 holds the estimation results. The estimated model includes the variables discussed in the data section and additional interaction terms of the overall load with a) the level of the treatment, b) the level of wind electricity generation and, c) the electricity generation capacity out of service. The interaction term consisting of the level of the treatment and the overall load was included to test whether the impact of solar electricity generation on the electricity price depends on the level of the overall load. The remaining two interactions were included to treat wind electricity generation and the loss of production capacity the same as solar electricity generation.

The estimation results in Table 6 are from ordinary least squares and unconditional quantile regression. The reader learns that the coefficients measuring the treatment, i.e. solar electricity generation, vary significantly across the conditional quantiles. For example, in the 10th percentile, the coefficient for the treatment is -10.046 compared to 1.266 for the mean. However, to gauge the impact of the treatment on the electricity price, the marginal effect as a combination of the level of the treatment and the interaction of the treatment with the overall load has to be calculated.

Table 6: *Mean and unconditional quantile regression results (N=1,450)*

Variables	**Mean**	**Percentile of price**				
	Mean	10th	25th	50th	75th	90th
treatment	1.266**	−10.046***	−3.539***	1.754***	5.646***	6.200***
	(0.401)	(1.791)	(0.703)	(0.519)	(0.957)	(1.028)
load	2.872***	7.742***	9.491***	5.765***	−6.344***	−5.661***
	(0.518)	(1.867)	(0.958)	(0.730)	(1.339)	(1.768)
treatment × load	−0.050***	0.147***	0.042**	−0.059***	−0.130***	−0.134***
	(0.007)	(0.029)	(0.014)	(0.010)	(0.020)	(0.021)
wind	−0.934***	−12.753***	−4.477***	0.988*	5.949***	5.529***
	(0.224)	(0.895)	(0.427)	(0.398)	(0.397)	(0.567)
wind × load	−0.010**	0.186***	0.057***	−0.039***	−0.136***	−0.120***
	(0.004)	(0.015)	(0.007)	(0.007)	(0.008)	(0.012)
out of service	0.507	−4.473***	−3.603***	−0.870*	5.200***	4.113***
	(0.274)	(1.147)	(0.516)	(0.418)	(0.629)	(0.768)
out of service × load	0.004	0.121***	0.087***	0.029***	−0.094***	−0.073***
	(0.005)	(0.020)	(0.009)	(0.009)	(0.014)	(0.017)
co2 price	6.943***	25.035***	14.296***	5.318**	−3.046	−6.724
	(1.140)	(5.015)	(2.281)	(1.704)	(2.470)	(3.700)
co2 price × load	−0.101***	−0.413***	−0.242***	−0.089***	0.083	0.133
	(0.020)	(0.081)	(0.036)	(0.028)	(0.049)	(0.075)
coal	−0.084	−2.310	−1.066	−1.725***	−1.248	1.298
	(0.374)	(1.805)	(0.727)	(0.525)	(0.867)	(1.185)
coal price × load	−0.001	0.048	0.023*	0.032***	0.022	−0.029
	(0.007)	(0.029)	(0.012)	(0.010)	(0.018)	(0.025)
gas price	3.450*	18.832**	21.913***	13.299***	−15.172***	−19.101***
	(1.502)	(6.689)	(3.085)	(2.258)	(3.513)	(4.463)
gas price × load	−0.024	−0.372***	−0.384***	−0.223***	0.349***	0.387***
	(0.026)	(0.109)	(0.050)	(0.039)	(0.068)	(0.089)
intercept	−169.421***	−406.208***	−532.024***	−309.546***	302.348***	338.231***
	(29.948)	(113.758)	(59.426)	(44.035)	(70.227)	(90.851)

Standard errors in parentheses. ***, **, * indicates statistical significance at the 0.1%, 1%, and 5% level.

The marginal effect of the treatment at varying points in the unconditional distribution of the load are indicated in Table 7. Depending on the overall load, the impact of $1,000$ additional Megawatt hours (MWh) of solar electricity generation on the electricity price estimated using ordinary least squares is between -0.754 Euros per Megawatt hour at the 10th percentile of total load and -2.129 at the 90th percentile of total load. While the mean regression estimates vary with the total load, the unconditional quantile regression estimates additionally vary with respect to the level of the electricity price.

For example, at the 10th percentile in the unconditional distribution of the price, the impact of $1,000$ additional MWh of solar electricity is -4.053 Euro per MWh when the overall load is 40.706 MW, -2.231 Euro per MWh at 53,082 MW, and 0.024 Euro per MWh at 68,402 MW. As expected, additional solar electricity generation has a negative impact on the electricity price. The magnitude of the effect increases with the load and the level of the price. The most negative effect with −4.053 Euro per MWh for each 1,000 MWh increase in solar electricity generation indicated at in Table 7 occurs at the 10th percentile in the unconditional distribution of the price and at an overall load on the order of 68,402 MW.

Table 7: *Marginal effects in Euro per 1,000 MWh (N=1,450)*

Percentile of load	**Mean**	**Percentile of price**				
(in 1,000 MW)		10th	25th	50th	75th	90th
10th	−0.754***	−4.053***	−1.836***	−0.648***	0.341*	0.746***
40.706	(0.114)	(0.631)	(0.187)	(0.133)	(0.163)	(0.179)
25th	−0.992***	−3.348***	−1.636***	−0.931***	−0.283**	0.104
45.495	(0.087)	(0.501)	(0.144)	(0.105)	(0.094)	(0.101)
50th	−1.369***	−2.231***	−1.319***	−1.379***	−1.272***	−0.912***
53.082	(0.065)	(0.309)	(0.124)	(0.102)	(0.135)	(0.134)
75th	−1.919***	−0.600***	−0.855***	−2.033***	−2.717***	−2.397***
64.165	(0.104)	(0.207)	(0.212)	(0.177)	(0.339)	(0.351)
90th	−2.129***	0.024	−0.678**	−2.283***	−3.269***	−2.965***
68.402	(0.130)	(0.278)	(0.261)	(0.215)	(0.421)	(0.439)

Standard errors in parentheses. ***, **, * indicates statistical significance at the 0.1%, 1%, and 5% level.

A previous analysis of Cludius et al. (2014) estimated the average price-reducing effect of 1,000 MWh additional solar electricity of −1,14 Euro per MWh from 2010-2012 for the German market. Würzburg et al. (2013) find a similar effect of −1,26 Euro per MWh in this respective period. Both estimates are similar to our result for the median total load of 53,082 MW in the mean regression.

Comparing the mean estimates with those obtained by unconditional quantile regression reveals that there is a large degree of heterogeneity regarding the impact of solar electricity generation on the electricity price. For example, at a total load of 64,165 MW, the mean estimate is −1.919 indicating that for each additional 1,000 MWhs of solar electricity generation the price of electricity decreases by almost 2 Euro per MWh. Unconditional quantile regression, however, implies that at the same level of load, the effect of solar electricity generation varies between −0.855 and −2.717. Thus, the difference in the point estimates across the conditional distribution of the price indicated in Table 7 is on the order of 1.862 Euro per MWh. The highest estimate is $\left(\frac{-0.855}{-1.919}\right) - 1 = 0.554$ times smaller compared to the mean estimate, while the lowest estimate is $\left(\frac{-2.717}{-1.919}\right) - 1 = 0.416$ times larger in magnitude.

Figure 13 compares the marginal effect for the mean compared and the median estimates (Panel 13a) and the comparison for the 10th compared to the 90th percentile of the conditional distribution of the price (Panel 13b). In Figure 13a, there is significant overlap in the 95% confidence intervals. In fact, the interval for the median estimate nearly covers the entire interval of the mean estimate. In both cases, the marginal effect of solar electricity on the electricity price is negative and the magnitude of the effect increases with increases in the percentile of the total load. In Figure 13b, the marginal effect of solar generation is shown for times for which the electricity price is low (10th percentile) compared to the time for which the electricity price is high (90th percentile).

Figure 13: *Marginal effect of solar electricity production by percentiles*

(a) *Mean vs. 50th percentile* **(b)** *10th vs. 90th percentile*

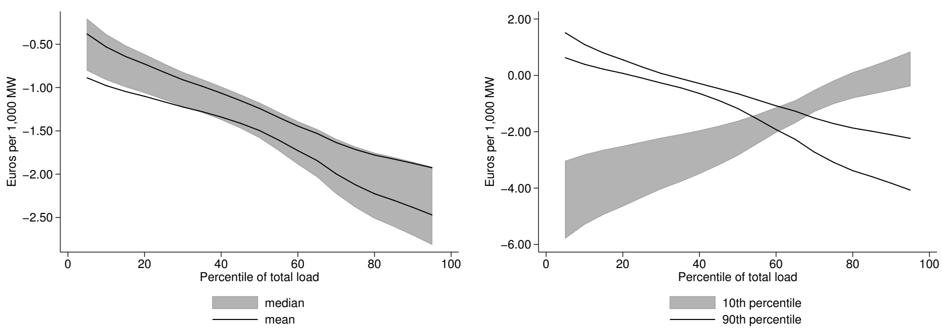

At low electricity prices, the effect decreases with increases in the total load holding all other controls constant. However, Table 7 shows that from the 50th percentile of the electricity price onwards, the marginal effect increases in magnitude with increases in the percentile of the load. In a few cases, the marginal effect is positive, which is counter intuitive. This positive effect occurs when the electricity price and the total load are simultaneously at the low ends of their respective distributions. Incidentally, the mean regression results are also driven upwards by these observations, but mean regression techniques do not allow for the identification of this effect.

We now test whether the estimated coefficients are significantly different from one another across the unconditional distribution of the price and the total load. Table 8 provides for these test results. The first two columns indicate the percentiles for which the difference in marginal effects are calculated. For example, the first row indicates the difference in the marginal effect at the 10th percentile and the 25th percentile in the unconditional distribution of the electricity price. The other columns indicate the total load for which the marginal effects are calculated. At a median load, the difference in the point estimate for the 10th and the 75th percentile is 0.891 Euro per MWh. With a standard error on the order of 0.382, the difference is statistically significant at the 0.1% level.

Table 8: *Differences in marginal effects*

Percentile of price		**Percentile of load**				
i	j	10th	25th	50th	75th	90th
10	25	−2.069**	−1.603*	−0.866*	0.212	0.624
		(0.720)	(0.566)	(0.349)	(0.290)	(0.387)
10	50	−3.276***	−2.317***	−0.797*	1.423***	2.272***
		(0.710)	(0.564)	(0.361)	(0.305)	(0.391)
10	75	−4.304***	−2.983***	−0.891***	2.166***	3.334***
		(0.728)	(0.568)	(0.382)	(0.464)	(0.588)
10	90	−4.716***	−3.377***	−1.255**	1.845***	3.030***
		(0.706)	(0.551)	(0.369)	(0.451)	(0.572)
25	50	−1.207***	−0.713**	0.069	1.212***	1.649***
		(0.242)	(0.186)	(0.169)	(0.302)	(0.370)
25	75	−2.235***	−1.379***	−0.025	1.954***	2.711***
		(0.316)	(0.224)	(0.240)	(0.509)	(0.629)
25	90	−2.647***	−1.773***	−0.389	1.633***	2.406***
		(0.324)	(0.212)	(0.217)	(0.514)	(0.643)
50	75	−1.028***	−0.666***	−0.094	0.742	1.062
		(0.241)	(0.161)	(0.217)	(0.485)	(0.598)
50	90	−1.440***	−1.060***	−0.458*	0.421	0.757
		(0.265)	(0.164)	(0.205)	(0.497)	(0.619)
75	90	−0.413*	−0.394***	−0.364*	−0.321	−0.305
		(0.199)	(0.113)	(0.160)	(0.408)	(0.509)

Standard errors in parentheses. ***, **, * indicates statistical significance at the 0.1%, 1%, and 5% level.

In fact, the test results in Table 8 reveal statistically significant differences in the marginal effects holding the level of total load constant. The highest differences can be found at high total loads for the upper percentiles of the price. In the 10th percentile of total load, the highest difference in marginal effects occurs between the 10th and 90th percentile of the price with a magnitude of -4.716 Euro per MWh for an increase in solar production of 1,000 MWh. Altogether, the tests allow for the conclusion that there is significant heterogeneity in the marginal effects at the same level of total load. Thus, the application of unconditional quantile regression has unearthed insights into the effect of the treatment otherwise overlooked by mean regression.

8 Return on renewable energy subsidies

The subsidization of renewable energies is a means in the pursuit of a battery of policy targets. First amongst these targets is the comprehensive reduction of carbon emissions. Another goal is the reduction of the dependence on energy imports. A third target is the development of cost-effective renewable energy technologies by way of demand-pull strategies.

While the reduction in carbon emissions by substituting conventional electricity generation with renewable energy sources conveys benefits, their quantification is beyond the scope of this paper. The same is true with regard to any technological progress that resulted from the spending on renewable energy technologies. It is, however, possible to compare the subsidies for solar electricity with the impact of solar electricity generation on the electricity price. In 2014, the most recent year in the data set, the remuneration for solar electricity was 13.68 Euro cents per kilowatt hour or $136,680$ Euros per $1,000$ Megawatt hours.

Table 9: *Total cost reduction in Euro for a 1,000 MWh increase in solar electricity generation*

Load	**Price**					
	Mean	10th	25th	50th	75th	90th
10th (40.706)	-30,710	-164,987	-74,750	-26,387	13,881	30,353
25th (45.495)	-45,138	-152,322	-74,430	-42,351	-12,885	4,730
50th (53.082)	-72,651	-118,440	-69,997	-73,179	-67,519	-48,436
75th (64.165)	-123,117	-38,478	-54,863	-130,429	-174,307	-153,834
90th (68.402)	-145,631	1,646	-46,363	-156,146	-223,591	-202,823

Total load in 1,000 MWh. All prices in Euros of 2010.

Table 9 indicates that the treatment effect exceeds the 2014 costs of subsidized solar electricity in 8 out of the 30 indicated combinations of load and price. While the total cost reduction relying on mean regression indicates that the cost reduction exceeds the subsidies only for the 90th percentile of load, the unconditional quantile estimates provide for a more nuanced interpretation. For example, at low electricity prices and low total loads the cost reduction exceeds the subsidies. The same occurs for combinations of high prices and high loads. Moreover, there exist some combinations of price and load for which effect and subsidy are nearly balanced. Altogether, unconditional quantile regression uncovers significant effects in cases in which mean regression is unable to reject the null. Policy makers gather a more nuanced understanding about the circumstances in which renewable subsidies are cost effective.

9 Conclusion

In this analysis, we addressed the causal effect of solar electricity production on day-ahead electricity price at the electricity exchange EPEX. The identification of causal effects rests on two pillars. First, we exploit the exogeneity of solar radiation that occurs only during the light hours of the day.

The amount of solar electricity output varies with the exogenous variation in solar radiation and, thus, is also exogenous to market participants.

However, electricity production and consumption vary by time of day. Therefore, the second pillar of our identification strategy consists of employing coarsened-exact matching. The intuition of this matching method is to "coarsen" continuous variables into categorical variables with sufficiently small categories in order to be of little consequence for the analysis. After the coarsening, exact matching is carried out on all variables relevant for the analysis. After the matching, the coarsened variables are replaced with their original values to preserve all information.

The application of coarsened exact matching allows us to identify counterfactual observations with near identical observations. Altogether, the combination of the exogeneity of solar radiation and near identical counterfactuals approximates an experimental setting with successful stratified random allocation of observation units into treatment and control groups. We indicate comprehensive evidence that coarsened exact matching removed the systematic differences in the observable characteristics between the treatment and control group that existed before the matching. Moreover, we indicate that although observations for which no counterfactual existed were removed from the data set, the range of the observed characteristics were not influenced by the matching. Thus, the external validity of our analysis is not affected in any way.

We estimated a model that interacts the treatment in form of solar generation with the overall load. This takes into account that the marginal costs of electricity production increase considerably with increases in total load. Combining data provided by the electricity exchange EPEX with unconditional quantile regression, we find that the effect of the treatment depends on the level of total load and on the level of the dependent variable itself. At a given level of total load, the effect of the treatment is significantly different for different levels of the electricity price.

For example, at the median total load the average treatment effect of an additional $1,000$ Megawatt hours of solar electricity generation estimated using mean regression is -1.369. Applying unconditional quantile regression, the effect of the same treatment varies between -2.231 and -0.912 depending on the level of the electricity price. At the 90th percentile of the total load, the mean estimate is -2.129, while the unconditional quantile estimates vary between -3.269 and 0.024. This kind of insight is barred to those who rely on mean regression techniques.

Because solar electricity generation is subsidized in the form of feed-in tariffs, we compared the cost of the subsidy to the effect of the corresponding solar electricity production on the electricity costs. Depending on the level of the electricity price and the total load, the maximal cost reduction is about $223,591$ Euro compared to a level of subsidies on the order of $136,680$ Euro for 1,000 Megawatt of solar electricity production. However, the effect of the subsidies exceed their costs only for certain combinations of electricity prices and total loads.

References

BMWi (2015). *Zeitreihen zur Entwicklung der erneuerbaren Energien in Deutschland*. Bundesministerium für Wirtschaft und Energie.

Borenstein, S. (2012). The private and public economics of renewable electricity generation. *Journal of Economic Perspectives 26*(1), 67–92.

Borenstein, S., J. B. Bushnell, and F. A. Wolak (2002). Measuring market inefficiencies in california's restructured wholesale electricity market. *American Economic Review 92*(5), 1376–1405.

Bower, J. and D. Bunn (2001). Experimental analysis of the efficiency of uniform-price versus discriminatory auctions in the england and wales electricity market. *Journal of economic dynamics and control 25*(3), 561–592.

Cartwright, N. (2004). Causation: One word, many things. *Philosophy of Science 71*(5), 805–820.

Cludius, J., H. Hermann, F. C. Matthes, and V. Graichen (2014). The merit order effect of wind and photovoltaic electricity generation in Germany 2008–2016: Estimation and distributional implications. *Energy Economics 44*, 302–313.

Cox, D. R. (1958). Planning of experiments.

Firpo, S., N. M. Fortin, and T. Lemieux (2009). Unconditional quantile regressions. *Econometrica 77*(3), 953–973.

Fisher, R. A. (1935). The design ofexperiments. *Oliver and Boyd, Edinburgh*.

Friedman, D. (1991). A simple testable model of double auction markets. *Journal of Economic Behavior & Organization 15*(1), 47–70.

Gelabert, L., X. Labandeira, and P. Linares (2011). An ex-post analysis of the effect of renewables and cogeneration on Spanish electricity prices. *Energy Economics 33*(1), 59–65.

Granger, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, 424–438.

Heckman, J. J. (2005). The scientific model of causality. *Sociological methodology 35*(1), 1–97.

Heckman, J. J. (2008). Econometric causality. *International statistical review 76*(1), 1–27.

Hirth, L. (2013). The market value of variable renewables: The effect of solar wind power variability on their relative price. *Energy Economics 38*, 218 – 236.

Ho, D. E., K. Imai, G. King, and E. A. Stuart (2007). Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis 15*(3), 199–236.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American statistical Association 81*(396), 945–960.

Iacus, S. M., G. King, and G. Porro (2011a). Causal inference without balance checking: Coarsened exact matching. *Political analysis*, mpr013.

Iacus, S. M., G. King, and G. Porro (2011b). Multivariate matching methods that are monotonic imbalance bounding. *Journal of the American Statistical Association 106*(493), 345–361.

Jónsson, T., P. Pinson, and H. Madsen (2010). On the market impact of wind energy forecasts. *Energy Economics 32*(2), 313–320.

Joskow, P. L. (2011). Comparing the costs of intermittent and dispatchable electricity generating technologies. *The American Economic Review 101*(3), 238–241.

Kagel, J. H. and W. Vogt (1993). Buyers bid double auctions: Preliminary experimental results. *The Double Auction Market: Institutions, Theories, and Evidence*, 285–305.

Ketterer, J. C. (2014). The impact of wind power generation on the electricity price in Germany. *Energy Economics 44*, 270–280.

King, G. and R. Nielsen (2015). Why propensity scores should not be used for matching. *Working Paper*.

Koenker, R. and G. Bassett (1978). Regression quantiles. *Econometrica 46*(1), 33–50.

Koenker, R. and K. Hallock (2001). Quantile regression: An introduction. *Journal of Economic Perspectives 15*(4), 43–56.

Morthost, P., S. Ray, J. Munksgaard, and A. Sinner (2010). *Wind energy and electricity prices. Exploring the merit order effect.* Belgium: European Wind Energy Association.

Müsgens, F. (2006). Quantifying market power in the german wholesale electricity market using a dynamic multi-regional dispatch model. *The Journal of Industrial Economics 54*(4), 471–498.

Ockenfels, A., V. Grimm, and G. Zoettl (2008). Strommarktdesign–preisbildungsmechanismus im auktionsverfahren für stromstundenkontrakte an der eex. *Gutachten im Auftrag der European Energy Exchange AG zur Vorlage an die Sächsische Börsenaufsicht 11.*

Ritter, N. (2016a). Does involuntary unemployment cause deterioration of mental or physical health? *unpublished.*

Ritter, N. (2016b). Experimental evidence on the effect of wage increases on total revenues, shifts worked, and work effort. *unpublished.*

Rothkopf, M. H. (1999). Daily repetition: A neglected factor in the analysis of electricity auctions. *The Electricity Journal 12*(3), 60–70.

Satterthwaite, M. A. and S. R. Williams (1989). The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. *The Review of Economic Studies 56*(4), 477–498.

Sims, C. A. (1972). Money, income, and causality. *The American economic review 62*(4), 540–552.

Splawa-Neyman, J., D. Dabrowska, and T. Speed (1990). On the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science 5*(4), 465–472.

Wilson, R. (1985). Incentive efficiency of double auctions. *Econometrica 53*(5), 1101–1115.

Wolfram, C. D. (1998). Strategic bidding in a multiunit auction: An empirical analysis of bids to supply electricity in england and wales. *The Rand Journal of Economics*, 703–725.

Würzburg, K., X. Labandeira, and P. Linares (2013). Renewable generation and electricity prices: Taking stock and new evidence for Germany and Austria. *Energy Economics 40*, 159–171.

Ziel, F., R. Steinert, and S. Husmann (2015). Efficient modeling and forecasting of electricity spot prices. *Energy Economics 47*, 98–111.