

Do we actually understand the impact of renewables on electricity prices? A causal inference approach

Davide Cacciarelli^{1,*}, Pierre Pinson^{1,2,3,4}, Filip Panagiotopoulos⁵, David Dixon⁵, and Lizzie Blaxland⁵

¹Imperial College London, London, UK

²Technical University of Denmark, Kgs. Lyngby, Denmark

³Halfspace, Copenhagen, Denmark

⁴Aarhus University, Aarhus, Denmark

⁵National Energy System Operator, Wokingham, UK

*d.cacciarelli@imperial.ac.uk

ABSTRACT

The energy transition is profoundly reshaping electricity market dynamics. It makes it essential to understand how renewable energy generation actually impact electricity prices, among all other market drivers. These insights are critical to design policies and market interventions that ensure affordable, reliable, and sustainable energy systems. However, identifying causal effects from observational data is a major challenge, requiring innovative causal inference approaches that go beyond conventional regression analysis only. We build upon the state of the art by developing and applying a local partially linear double machine learning approach. Its application yields the first robust causal evidence on the distinct and non-linear effects of wind and solar power generation on UK wholesale electricity prices, revealing key insights that have eluded previous analyses. We find that, over 2018–2024, wind power generation has a U-shaped effect on prices: at low penetration levels, a 1 GWh increase in energy generation reduces prices by up to 7 GBP/MWh, but this effect gets close to none at mid-penetration levels (20–30%) before intensifying again. Solar power places substantial downward pressure on prices at very low penetration levels (up to 9 GBP/MWh per 1 GWh increase in energy generation), though its impact weakens quite rapidly. We also uncover a critical trend where the price-reducing effects of both wind and solar power have become more pronounced over time (from 2018 to 2024), highlighting their growing influence on electricity markets amid rising penetration. Our study provides both novel analysis approaches and actionable insights to guide policy makers in appraising the way renewables impact electricity markets.

Increasing the share of electricity from renewable energy sources, e.g., wind energy and solar energy, is one of the key actions towards climate change mitigation. Projections from the International Energy Agency (IEA) indicate that by 2030, wind and solar power generation assets will jointly contribute to 40% of global electric power generation¹. In the UK, wind energy alone represented 29.4% of total electricity generation in 2023². Solar energy, while currently in a less prevalent position, is increasing its market share rapidly, driven by low technology costs and generous support policies. The integration of these volatile energy sources (also with limited predictability) is profoundly affecting electricity markets, requiring a deeper understanding of their impact on market dynamics. Although the environmental benefits of renewable energy generation, such as reduced CO₂ emissions, are well-documented³, their increasing penetration introduces both opportunities and challenges for the operation of the electricity market⁴. From a market perspective, the inherent lower marginal costs associated with wind and solar power generation are expected to exert downward pressure on wholesale day-ahead prices (also referred to as spot prices) by shifting the equilibrium point through the merit-order effect⁵. However, from an operational point of view, the variability and limited predictability of renewable energy generation induce an additional need for system balancing and ancillary services^{6,7}. The market impact of renewable energy sources, particularly wind energy, has been the subject of extensive research over the past two decades. Early studies were conducted for the Danish electricity market (part of Nord sPool). Although actual wind power generation was found to have limited impact on prices⁸, forecast wind power penetration emerged as the primary determinant, reflecting the information available to market participants at the time of gate closure⁹. Since these pioneering studies, researchers have conducted similar investigations in numerous other electricity markets, such as in Australia^{10–12}, Germany^{13–17}, Ireland^{18,19}, Italy²⁰, Korea²¹, the Netherlands²², the Nordics^{23–25}, Spain^{26–28}, the UK^{29,30}, and the US^{31–34}. More recently, the *cannibalisation* effect, in which the market value of renewable energy declines as their penetration levels increase by displacing more expensive forms of generation, has been investigated for several European markets³⁵. These studies highlight common trends observed across electricity markets globally, such as the price reduction resulting from the merit-order effect and the increased price volatility linked to higher levels of installed wind capacity.

Despite these relevant insights, most of the aforementioned works rely on traditional regression-based approaches, which are often limited in their ability to fully capture the complexity of the impact of renewable energy on electricity markets. Here though, the diagnostic analytics approaches to be developed and the insights to be derived ought to be causal. Specifically, simpler analyses might struggle to unravel the influence of confounding factors, such as demand fluctuations, gas prices, and hourly price profiles. In addition, they often overlook the non-linear dynamics that govern the impact of renewable energy penetration on electricity prices, leading to biased or incomplete conclusions. Indeed, traditional econometric approaches are often constrained by their fully parametric nature, requiring strong assumptions about the functional form of the relationships being modelled. As a result, traditional methods often provide associations, but lack the causal clarity needed to inform policy making effectively. If not revealing true causal effects, consequent market interventions and new policies will be suboptimal at best, and completely inefficient in the worst case.

In contrast, our approach represents a significant methodological and analytical advancement. We introduce a bespoke causal inference framework that generalises double machine learning (DML)³⁶ for use in a non-linear and non-stationary context. Our *local partially linear DML* method enables the identification of non-linear treatment effects while accounting for a wide range of confounding factors. This allows us (i) to isolate the effects of predicted wind and solar power penetration from other influences, (ii) to examine how these impacts vary with predicted penetration levels, and (iii) to track their evolution over time. Focusing on the UK electricity market, which has seen a substantial rise in renewable energy penetration in recent years, our analysis provides critical insights that traditional regression methods cannot offer. We demonstrate that neglecting robust causal inference techniques can lead to inaccurate conclusions about the effect of renewable energy generation on wholesale prices, potentially misguiding analysts and policy makers. While prior studies laid the groundwork for understanding the relationship between renewable energy and electricity prices, our approach offers a new perspective by applying a rigorous analytical and causal lens to uncover the complex, non-linear effects induced by wind and solar energy.

Relationship between renewable energy penetration and electricity prices

We first examine the relationship between predicted renewable energy penetration and wholesale electricity prices in the UK over a period from 2018 to 2024. Figure 1 presents bar plots of the mean day-ahead APX prices in the UK across varying levels of predicted wind and solar power penetration. For wind power penetration (Figure 1a), a clear downward trend emerges, where higher penetration levels are associated with lower electricity prices. This aligns with the well-known merit-order effect, as increased predicted renewable energy generation displaces higher-cost conventional generation, reducing market prices. In contrast, for solar power penetration (Figure 1b), the pattern is more complex. While moderate solar penetration levels are linked to lower prices, an unexpected “bump” in the price distribution emerges at around 4–7% penetration. This counter-intuitive phenomenon may be influenced by external factors such as higher demand during summer months, energy supply disruptions, or other market conditions. Without careful causal analysis, these fluctuations could be mis-attributed solely to solar penetration rather than a broader set of intertwined effects within electricity markets.

To reinforce these initial findings, we performed a regression-based analysis using techniques commonly employed in the literature (using local polynomial (quantile) regression⁹). Figure 2 illustrates the relationship between forecast wind and solar power penetration and spot prices, over the same period. For wind power (Figure 2a), the results corroborate earlier insights, showing that higher predicted wind power penetration generally leads to lower wholesale electricity prices. Specifically, as predicted wind power penetration increases, daily price patterns undergo a noticeable transformation: the pronounced peaks during morning and evening hours at low penetration levels are significantly smoothed. This indicates that wind power forecasts not only reduces average prices but also dampen daily price spikes. Results from the quantile regression model (Figure 2c) provide further insights into the distributional effects of wind power generation. With higher levels of predicted wind penetration, the range between the 10% and 90% quantiles narrows, suggesting reduced price variability. However, the lighter colour shading at high penetration levels indicates limited data in these regions, cautioning against overgeneralising these findings.

The analysis of predicted solar power penetration (Figures 2b and 2d) reveals a more complex relationship. As noted earlier (Figure 1b), higher levels of solar penetration appear to be linked with rising spot prices (which is counter-intuitive). The discrepancy between observed behaviour for wind and solar power also highlights the importance of penetration levels. With wind energy being more pervasive in the UK market, its effect is easier to discern, whereas the impact of solar energy is potentially overshadowed by external factors. These results underscore the limitations of regression-based approaches in disentangling the causal effects of renewable energy penetration. Without accounting for confounding variables, we risk drawing misleading conclusions, such as inferring that higher solar power penetration drives higher wholesale prices. This emphasises the necessity of devising and employing robust causal inference methods to accurately attribute changes in electricity prices to renewable energy penetration.




Figure 1. Price variation across predicted renewable penetration levels. Bar plots show the mean APX prices for wind and solar penetration levels, over 2018-2024, with error bars representing associated 95% confidence intervals. The predicted wind power penetration intervals (a) are defined in 10% increments, while those for predicted solar power penetration (b) are defined in 1% increments, reflecting the lower installed solar power capacity in the UK.

Causal impact of renewables on electricity prices

Understanding the impact of wind and solar energy generation on electricity prices is inherently a causal question: what happens to electricity prices if we predict one more GWh of wind or solar energy? This inquiry goes beyond a simple statistical association; it seeks to uncover the direct, underlying effect of additional renewable energy integration. Historically, experimentation has been central to scientific discovery, providing controlled environments to isolate causes and effects. However, in many real-world environments, such as electricity markets, conducting controlled experiments is neither feasible nor practical. In these cases, deriving causal insights from observational data alone remains one of the most formidable challenges in science and statistics^{37,38}. Given the limitations of traditional regression models, we employ an approach based on DML to move beyond simple associations and uncover the true effects of renewable integration. Specifically, our *local partially linear DML* framework enables us to isolate the non-linear impact of wind and solar power production on wholesale electricity prices. This methodology addresses confounding factors that often distort observational analyses, providing a robust foundation for understanding how renewable integration affects market dynamics. Figure 3 shows the results of our DML framework on the APX prices, for the period 2018-2024 (results for NordPool and intraday prices are included in the supplementary information of the paper, in Figures S1 and S2). The *causal estimate* is the conditional average treatment effect (CATE), which represents the impact of a 1 GWh increase in predicted renewable energy generation on wholesale electricity prices, while controlling for a wide set of confounders. Instead, the *observational mean* reflects the raw association between renewable penetration and electricity prices at the same penetration level, without accounting for confounding factors. The full list of potential confounding variables to consider is given in Table S1 of the supplementary information. It includes variables like varying installed capacity (for wind and solar energy), time of day, time of year, load, gas price, carbon permits, etc.

The causal analysis reveals distinct effects for wind and solar power. Wind power exhibits a U-shaped causal effect on electricity prices (Figure 3a): the price-reducing effect from wind energy generation is strongest at low penetration levels, diminishes at moderate levels, and increases again at higher penetration levels. This pattern likely reflects the interaction between wind energy generation and the slope of the supply curve in the wholesale electricity markets, at various penetration levels. At low wind power penetration, wind power rapidly displaces the most expensive marginal generators, leading to substantial price reductions. As penetration increases, this effect diminishes temporarily, only to re-emerge as wind displaces more entrenched higher-cost generators. The reduced effect in the moderate range (20-30%) could arise because the supply curve flattens as gas providers with similar cost structures enter the market, reducing the sensitivity of prices to additional wind power generation. For solar power (Figure 3b), our causal inference approach uncovers a consistent price-reducing effect across all penetration levels. This finding stands in sharp contrast to the observational mean trends, which suggest a “bump” in electricity prices at moderate solar penetration (around 5%). This bump, observed in simpler analyses, likely arises from un-addressed confounding factors such as demand shifts or major events occurring during periods of high solar

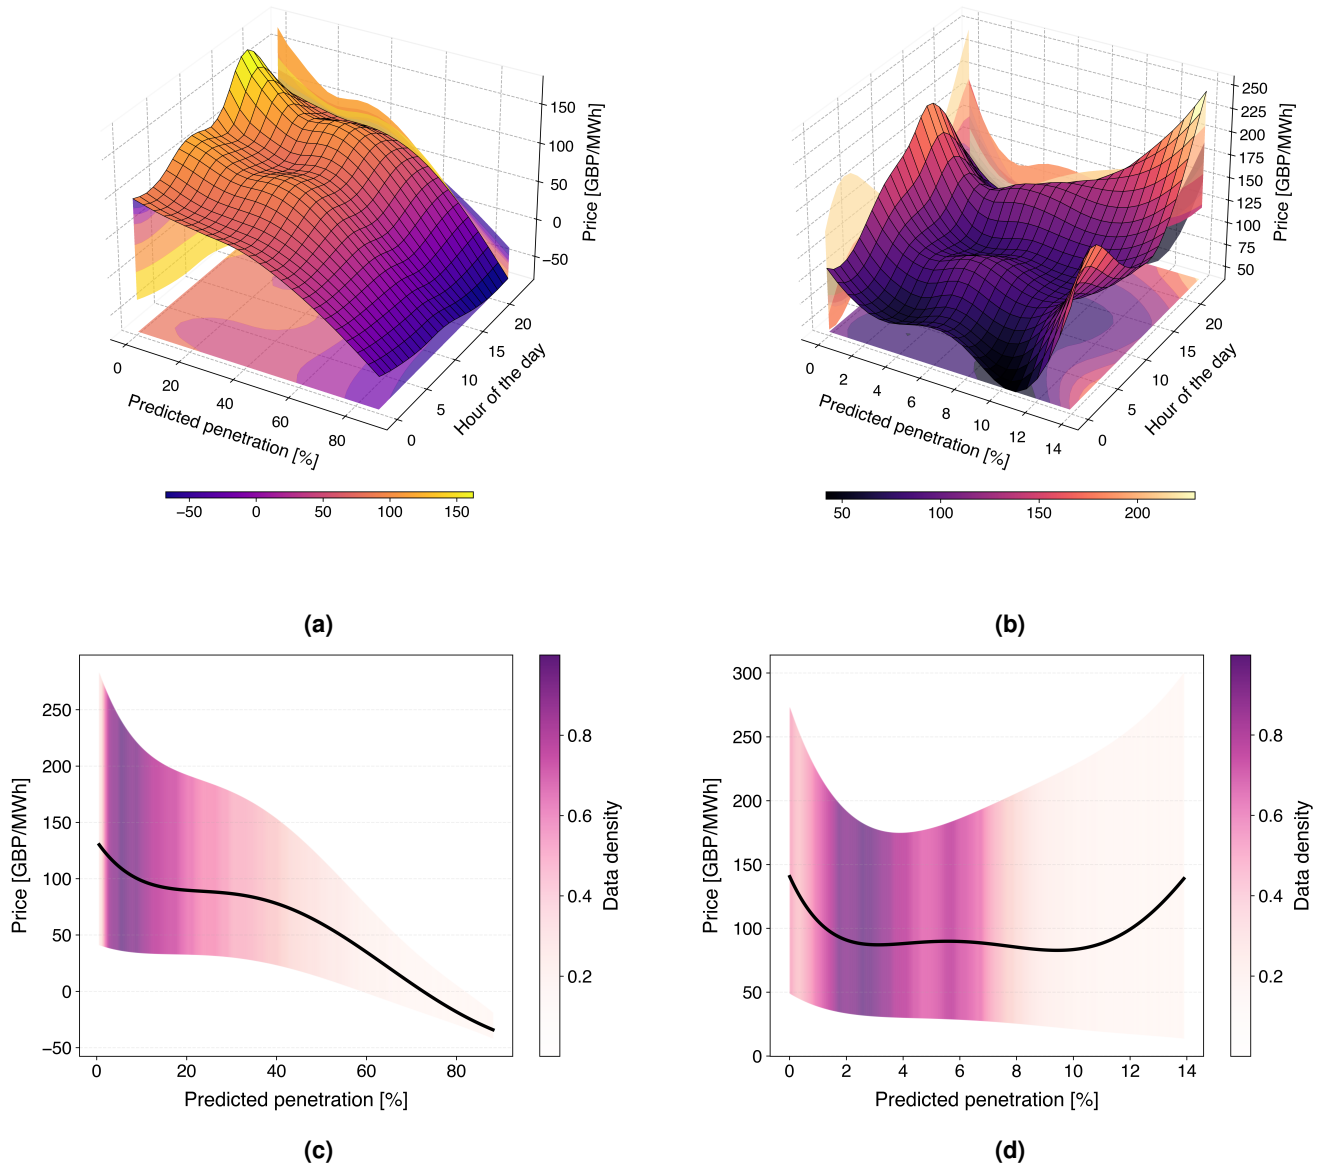


Figure 2. Regression analysis of predicted renewable penetration and APX wholesale prices. The top row shows the mean spot price modelled as a function of the hour of the day, as well as predicted (a) wind power penetration and (b) solar power penetration. The bottom row presents the results of quantile regression models, with the solid line representing the mean spot price, and the shaded region representing central 80%-coverage intervals, where the colour intensity indicates the data density (in terms of penetration data). These models capture the distributional effects of predicted (c) wind power penetration and (d) solar power penetration, on spot prices.

output. The causal analysis effectively isolates the intrinsic merit-order effect of solar power, demonstrating its role in reducing electricity prices even when confounders might obscure this relationship. Finally, we hypothesise that the effect of solar power on electricity prices might also exhibit a U-shaped pattern, similar to wind, since the effect fundamentally comes from the overall shape of the supply curve in the market. However, this potential trend is not fully observable in our analysis due to the limited attainable penetration of solar power, which reflects the reduced installed solar capacity in the UK. As solar power penetration increases in the future, it is plausible that a more pronounced U-shaped relationship could emerge, driven by dynamics analogous to those seen for wind power. Comparable patterns and causal relationships have been observed while analysing the effects on the NordPool (Figure S1) and intraday (Figure S2) electricity prices. In the supplementary information of the paper (Figure S3), we explored an alternative approach by using predicted renewable penetration not only as a contextual

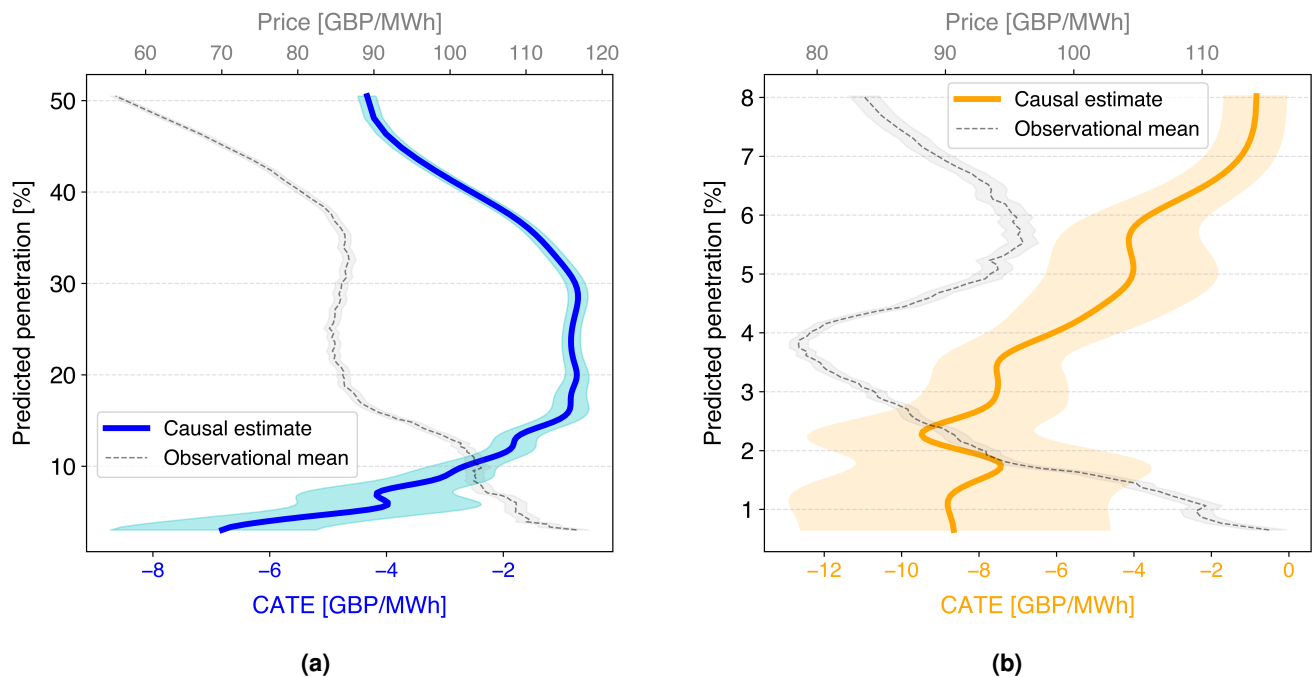


Figure 3. Causal impact of predicted renewable power production on electricity prices. Comparison of observational mean and causal effects of wind (a) and solar (b) power production on wholesale electricity prices. The solid lines represent the non-linear CATE estimates derived using our local partially linear DML framework, capturing the price impact (in GBP/MWh) of a 1 GWh increase in renewable energy generation. The dashed lines show the observational mean trends as a function of renewable penetration levels, illustrating the differences between raw associations and the true causal effects. Shaded areas denote 80% confidence intervals.

variable but also as the main input to the model. This analysis provides a complementary view by estimating the causal effect of increasing renewable penetration (%) directly, rather than focusing on energy generation (GWh).

These findings have significant implications for analysts, policy makers and market participants. The U-shaped effect observed in Figure 3a underscores the importance of considering how renewable energy influences market structures across varying levels of market penetration. For solar power, the discrepancy between observational and causal trends underscores the risks of relying on simpler regression methods, which may lead to disproportionately erroneous conclusions about the economic impacts of renewables. By providing an accurate assessment of the true effects of renewable generation, our approach supports better-informed market designs and policy interventions aimed at maximising the benefits from renewable energy integration.

Evolution of causal effects over time

The impact of renewables on electricity prices is not static. It evolves alongside market conditions, technological advancements, and changes in renewable penetration levels driven by additional installed generation capacity. To investigate these dynamics, we analyse the non-linear causal effect over time using a sliding window approach spanning two financial years. This approach provides a temporally resolved view of how the price effects of wind and solar energy have changed over the period 2018-2024, while ensuring a sufficient number of observations to apply our DML framework. The findings in Figure 4 reveal a growing impact of renewable energy on electricity prices over the years. In particular, the magnitude of wind's price-reducing effect (Figure 4a) has significantly increased in the last four years, reflecting its expanding role in the electricity market and the growing relevance of the merit-order effect. This trend underscores the profound influence of wind power on market outcomes, highlighting its policy implications as governments and regulators aim to integrate higher levels of renewable energy into the grid. The evolution of the effect of solar power on electricity prices (Figure 4b) also appears to be increasing, indicating that the impact of solar energy may become substantial in a near future with the further deployment of solar power generation capacity in the UK.




Figure 4. Temporal evolution of the causal impact of renewable power production on electricity prices. non-linear CATE estimates of (a) wind and (b) solar power production over time, obtained using a sliding window spanning two financial years, with each window containing approximately 35,000 observations. Shaded regions represent 80% confidence intervals.

Discussion and conclusions

We bridge a critical gap in understanding the causal impact of renewable energy penetration on electricity prices, offering insights that extend beyond regression-based associations only. By applying a robust causal inference framework, we demonstrated how solar power consistently reduces electricity prices, underscoring its fundamental merit-order effect. Similarly, the U-shaped causal effect of wind power highlights the non-linear impact of wind integration, linked to the slope of the overall market supply curve at the equilibrium point, at varying penetration levels. Our analysis further revealed the temporal trends of these causal effects, demonstrating how the influence of renewables on electricity prices has steadily intensified over the period from 2018 to 2024. These findings carry important implications for energy policy makers and market participants. Traditional analyses that rely solely on observational data risk misinterpreting the economic effects of renewables, potentially mis-informing future policy interventions in electricity markets. By contrast, our causal framework provides evidence-based insights that more accurately reflect the true economic benefits of renewables. This evidence underscores the importance of embedding causal methods into the decision-making processes that shape future energy systems.

The implications of our findings extend beyond the UK market only. As many electricity markets transition to higher shares of renewables, the lessons from our analysis become increasingly relevant. The methodology we employ can be adapted to other contexts, providing a versatile tool for understanding how renewable integration impacts market outcomes under diverse conditions. Furthermore, the effects we observe are likely to amplify as renewable energy penetration increases globally, suggesting that the price-reducing effect of renewables will become even more pronounced in markets with higher targets for wind and solar energy. This study also highlights areas for future research. A key question involves similarly assessing the impact of renewables on balancing costs and the growing need for reserves to address variability, uncertainty, and grid constraints. Additionally, exploring interactions between renewable energy penetration and emerging technologies such as long-duration storage and demand-side management, or evaluating the broader social benefits of renewables such as emissions reductions and energy equity through a causal lens would provide a more comprehensive picture of their societal value.

In conclusion, our findings emphasise the importance of grounding energy policy and market interventions in robust causal evidence. As the deployment of wind and solar power generation capacities is poised to strengthen in the future, adopting analysis methodologies able to accommodate the complexities of modern electricity systems will be essential for achieving affordable, sustainable, and reliable power systems. By advancing the causal understanding of renewable energy impact on electricity markets, we contribute to a foundation for smarter policies and a smoother transition to a decarbonised future.

Methods

This study employed different layers of analysis to explore the impact of predicted wind and solar power penetration on electricity prices. To provide a baseline using methodologies commonly employed in the literature, we first analysed the relationship between predicted renewable penetration and spot prices using regression analysis. Then, we developed a robust causal inference approach to estimate the evolving and non-linear impact of renewables while controlling for a large set of confounding variables.

Data

The dataset for this study comprises electricity market data, wind and solar generation forecasts, and various other market-relevant variables for the UK, spanning the period from 2018 to 2024. A comprehensive list of data sources, detailed descriptions of the variables, and the preprocessing steps used to derive additional variables are provided in the supplementary information of the paper (Table S1).

Regression-based analysis

Mean smoothing. Given the well-established non-linear nature of the relationship between electricity prices, predicted renewable energy penetration, and time of day, we employed locally weighted polynomial regression models⁹. These models assume that although the overall relationship is non-linear, it can be locally approximated by a series of linear models. So, by choosing a set of m fitting points, we can locally estimate m linear models representing wholesale electricity prices as a function of predicted wind or solar penetration and time of the day. For this study, we choose 576 fitting points arranged in a 24×24 grid. At each fitting point \mathbf{x}_u , given that we have n observations (\mathbf{z}_i, p_i) available, the price p_i is modelled as

$$p_i = \beta_u^\top \mathbf{z}_i + \varepsilon_i, \quad i = 1, \dots, n, \quad (1)$$

where β_u represents the vector of regression coefficients tailored specific to the fitting point \mathbf{x}_u , \mathbf{z}_i comprises a vector of input features, and ε_i denotes the error term (centred, and with finite variance). While with a higher number of fitting points we could adequately approximate non-linear functions with simple first-order models, using polynomial features enables the modelling of more complex relationships. If we let r_i and h_i represent the normalised renewable power level (wind or solar power) and the hour of the day, we can obtain a second-order model by including quadratic and interaction terms, after expanding each data point $\mathbf{x}_i = [r_i \ h_i]^\top$ to $\mathbf{z}_i = [1 \ r_i \ h_i \ r_i^2 \ r_i h_i \ h_i^2]^\top$. Locally weighted polynomial regression then estimates the model coefficients β_u by weighting observations based on their proximity to the fitting point (using a distance $d(\mathbf{x}_u, \mathbf{x}_i)$). This is achieved through a kernel function that assigns higher weights to nearby observations, reducing influence as the distance increases. Here, we used the tri-cube weight function given by

$$w_{u,i} = \begin{cases} (1 - d(\mathbf{x}_u, \mathbf{x}_i))^3, & \text{if } d(\mathbf{x}_u, \mathbf{x}_i) < 1 \\ 0, & \text{otherwise} \end{cases}, \quad (2)$$

where d is chosen as the normalised Euclidian distance between an observation \mathbf{x}_i and the fitting point at hand \mathbf{x}_u , i.e., $d(\mathbf{x}_u, \mathbf{x}_i) = \|\mathbf{x}_i - \mathbf{x}_u\|_2 / h_u$. An adaptive bandwidth h_u , for each fitting point \mathbf{x}_u , can be used to ensure that observations beyond a certain distance from the current point exert no influence on parameter estimation. We calculated h_u using the 30th percentile of distances from all data points to a given fitting point, accommodating to local data density and variability. At each fitting point, the estimated parameter vector $\hat{\beta}_u$ is obtained by minimising the weighted sum of squared residuals, as in

$$\hat{\beta}_u = (\mathbf{Z}^\top \mathbf{W}_u \mathbf{Z})^{-1} \mathbf{Z}^\top \mathbf{W}_u \mathbf{p}, \quad (3)$$

where \mathbf{W}_u is a diagonal matrix containing the weights $w_{u,i}$, \mathbf{Z} is the design matrix expanded to the quadratic model form (i.e., whose rows are all input feature observations \mathbf{z}_i^\top), and \mathbf{p} is the vector of observed response values p_i . This methodology effectively captures local variations by emphasising the influence of nearby observations, thus providing a detailed and smooth representation of the relationship between wind power production, time of day, and wholesale prices.

Quantile modelling. Understanding the relationship between predicted renewable penetration, time of the day, and mean wholesale prices provides interesting insights on the behaviour of the market. However, to cope with the inherent uncertainty of renewable generation, it is equally important to describe the distribution of data around the mean trend. Using quantile regression, we can get insights into the conditional distribution of the price, indicating the value below which a certain proportion of observations fall for a given quantile with nominal level q . For example, if $q = 0.9$, the quantile regression line will show the value below which 90% of the observed data points lie. Locally weighted polynomial regression can be extended to model quantiles by minimising a cost function that differently weights prediction errors depending on whether they

fall below or above the specified quantile. This approach captures the asymmetric risk associated with underestimations or overestimations. The model parameters for a specific quantile q are determined by

$$\hat{\beta}_{u,q} = \underset{\beta}{\operatorname{argmin}} \sum_{i=1}^n w_{u,i} \rho_q(p_i - \mathbf{z}_i^\top \beta), \quad (4)$$

where p_i represents the observed price, \mathbf{z}_i is the i -th observation expanded to the model form and w_i is its weight (computed using (2)), β is the vector of coefficients to be estimated, and ρ_q is the quantile loss function defined as $\rho_q(\varepsilon) = \varepsilon(q - \mathbf{1}\{\varepsilon < 0\})$, with $\mathbf{1}\{\varepsilon < 0\}$ being an indicator function that takes the value 1 if the residual ε is negative and 0 otherwise, and q represents the nominal level of the quantile of interest (e.g., 0.1 for the 10% quantile and 0.9 for the 90% quantile). Unlike the case for mean modelling, there is no closed-form solution for the parameter vector $\hat{\beta}_{u,q}$, which is estimated at each fitting point using an iterative gradient-based optimisation method to minimise the quantile loss function in (4). This methodology provides a detailed representation of the behaviour of prices, beyond average effects, for varying levels of predicted wind power production or penetration.

Causal inference using double machine learning (DML)

Traditional regression models may fall short in disentangling causation from confounding influences, especially when the data involve complex, non-linear interactions or high-dimensional confounders like fuel prices, demand fluctuations, and seasonal patterns. The DML framework³⁶ leverages machine learning algorithms to flexibly model the relationships between confounders, treatments, and outcomes. This approach enables the estimation of causal effects in the presence of complex, high-dimensional data while mitigating the spurious effects introduced by confounders. Within the partially linear DML framework³⁹, we assume that the response p (here, the electricity price) is a function of the treatment t (e.g., predicted renewable energy penetration) and other confounding variables \mathbf{a} , as in

$$p_i = \beta t_i + f(\mathbf{a}_i) + \varepsilon_i, \quad i = 1, \dots, n, \quad (5)$$

where β represents the (linear) effect of t on the outcome p , f is a potentially non-linear function of the confounders \mathbf{a} , and ε is an error term (centred, and with finite variance). Similarly, we assume that the treatment variable itself can be modelled as a function of the confounders, or a subset of them, i.e.,

$$t_i = g(\mathbf{a}_i) + \eta_i, \quad i = 1, \dots, n, \quad (6)$$

where g represents the modelled relationship between t and \mathbf{a} , and η is an error term (centred, and with finite variance). The DML framework is designed to isolate the causal effect of the treatment variable β by adjusting for confounding influences in two stages. First, we estimate the nuisance parameters using two machine learning models to approximate the functions f and g . Then, we use the trained models to “remove” the influence of confounders from both the treatment and the response variables. Then, it regresses the residualised response $\tilde{p}_i = p_i - f(\mathbf{a}_i)$, $i = 1, \dots, n$, on the residualised treatment $\tilde{t}_i = t_i - g(\mathbf{a}_i)$, $i = 1, \dots, n$, to estimate β . Here, β coincides with the *average treatment effect* (ATE), or average causal effect, which measures the overall impact of a treatment on the outcome across the entire population. Formally, for a continuous treatment variable t , the ATE is defined as

$$\text{ATE} = \mathbb{E}[p \mid do(t+1)] - \mathbb{E}[p \mid do(t)], \quad (7)$$

where $p|do(t)$ represent the potential outcome with treatment t and $p|do(t+1)$ the potential outcome after a unit increase in t . ATE can be interpreted as the expected increase in the response p , resulting from a unit increase in the treatment t . The relationship is hence linear and not conditional on any contextual variable.

Local partially linear DML

While the partially linear DML framework provides reliable causal estimates under the assumption of linear treatment effects, it may not fully capture the complexities of the electricity market, where the effect of renewables may vary based on the penetration level. To address this issue, we extend the standard DML approach by introducing a local partially linear DML framework based on a boxcar kernel. Rather than estimating a global ATE across the entire dataset, we assume that the relationship between treatment (renewable energy generation) and outcome (electricity prices) can be locally approximated using a partially linear model within small, predefined subsets of the data. This enables us to estimate the so-called *conditional average treatment effect* (CATE), which can vary depending on the segment of the data under consideration. Within this context, the boxcar kernel acts as a window function to isolate a subset of the data around each point of interest (e.g., for a given level of

penetration, time). The idea is similar in essence to the locally weighted polynomial regression models used previously. For a conditioning (and possibly multidimensional) variable \mathbf{x} and a fitting point \mathbf{x}_u , the boxcar kernel is defined as

$$K_u(\mathbf{x}) = \begin{cases} 1, & \text{if } \|\mathbf{x} - \mathbf{x}_u\|_2 \leq \frac{h}{2}, \\ 0, & \text{otherwise} \end{cases}, \quad (8)$$

where h is the width of the window. The kernel takes the value 1 within the bandwidth parameter h around \mathbf{x}_u and 0 otherwise. This kernel is applied to partition the data into smaller, rectangular regions where the partially linear model is assumed to hold, yielding a subset $\mathcal{W}(\mathbf{x}_u) = \{(\mathbf{x}_i, \mathbf{a}_i, t_i, p_i), K_u(\mathbf{x}_i) = 1\}$. For each segment of the data within the window, we estimate the CATE by applying the partially linear DML method. Formally, at each fitting point \mathbf{x}_u (i.e., a given predicted wind penetration level), we define the local CATE as

$$\text{CATE}(\mathbf{x}_u) = \mathbb{E}[p \mid do(t+1), \mathbf{x} \in \mathcal{W}(\mathbf{x}_u)] - \mathbb{E}[p \mid do(t), \mathbf{x} \in \mathcal{W}(\mathbf{x}_u)], \quad (9)$$

where $\mathcal{W}(\mathbf{x}_u)$ is the subset of data within the window centred around the fitting point \mathbf{x}_u , and the expectation is taken within this window. We refer to $\mathcal{W}(\mathbf{x}_u)$ as the contextual subset of data. The procedure follows two main steps, similar to the standard DML, but with the additional local partitioning. For each contextual subset $\mathcal{W}(\mathbf{x}_u)$, we use machine learning models to estimate the nuisance models f and g , and residualise both the treatment variable t and the outcome p . This removes the confounding effects, leaving us with the contextual residualised variables \tilde{t}_u and \tilde{p}_u , with vectors of values $\tilde{\mathbf{t}}_u$ and $\tilde{\mathbf{p}}_u$ that corresponds to the contextual subset $\mathcal{W}(\mathbf{x}_u)$. The variables used for residualising wind power production, solar power production, and spot prices are listed in Table S1 of the supplementary information. Then, at each window we estimate the CATE with

$$\hat{\beta}_u = \text{OLS}(\tilde{\mathbf{p}}_u, \tilde{\mathbf{t}}_u \mid \mathbf{x} \in \mathcal{W}(\mathbf{x}_u)) = (\tilde{\mathbf{t}}_u^\top \tilde{\mathbf{t}}_u)^{-1} \tilde{\mathbf{t}}_u^\top \tilde{\mathbf{p}}_u. \quad (10)$$

We use a bootstrap procedure for each contextual subset $\mathcal{W}(\mathbf{x}_u)$ to ensure robust results and to capture the distributional uncertainty of our estimates. Details of the estimation procedure are provided in Algorithms 1 and 2 in the supplementary information for this paper. Additionally, a Gaussian filter was applied to reduce the noise in the CATE estimates and to provide a smoother representation of the non-linear effects. The individual CATE estimates prior to smoothing are presented in the supplementary information (Figures S4 and S5).

Data availability

The data utilised in this study is available at <https://github.com/dcacciarelli/market-impact-renewables/tree/main/data>.

Code availability

The Python scripts to replicate all the results presented in the paper are available at <https://github.com/dcacciarelli/market-impact-renewables>.

Author contributions statement

D.C. developed the methodology, conducted the analysis, and drafted the manuscript. P.P. provided continuous supervision in methodology development, experimental design, result analysis, while also contributing to the writing of the manuscript. F.P. and D.D. supported data sourcing and preprocessing and contributed to the result analysis. L.B. conceptualised the initial idea and provided feedback throughout the project. All authors reviewed the manuscript and accepted its submission for potential publication.

Competing interests

The authors declare no competing interests.

Supplementary information

Supplementary notes

Partially linear DML

The DML framework³⁶ is a powerful approach that integrates machine learning with econometric techniques to provide unbiased causal estimates even when the data is characterised by nonlinear relationships among numerous confounders. In the context of DML, one of the most commonly encountered setup is the partially linear one³⁹.

Algorithm 1 Partially linear DML

- 1: **Input:** dataset $D = \{(\mathbf{a}_i, t_i, p_i)\}_{i=1}^n$, number of folds K
- 2: **Output:** linear ATE estimate $\hat{\beta}$
- 3: Randomly partition the dataset into K folds. ▷ Split the data
- 4: **for** each fold $k = 1, \dots, K$ **do** ▷ Train predictive models
- 5: Train a model \hat{f}_{-k} using $K - 1$ folds to predict the response p from the confounders \mathbf{a} .
- 6: Train a model \hat{g}_{-k} using $K - 1$ folds to predict the treatment t from the confounders \mathbf{a} .
- 7: Use \hat{f}_{-k} and \hat{g}_{-k} to predict \mathbf{p} and \mathbf{t} in the held-out fold k .
- 8: Compute residuals ▷ Orthogonalisation
$$\tilde{\mathbf{p}}_k = \mathbf{p}_k - \hat{f}_{-k}(\mathbf{A}_k)$$
$$\tilde{\mathbf{t}}_k = \mathbf{t}_k - \hat{g}_{-k}(\mathbf{A}_k)$$
- 9: Regress $\tilde{\mathbf{p}}_k$ on $\tilde{\mathbf{t}}_k$ using OLS ▷ Effect estimation
$$\hat{\beta}_k = \text{OLS}(\tilde{\mathbf{p}}_k, \tilde{\mathbf{t}}_k) = (\tilde{\mathbf{t}}_k^\top \tilde{\mathbf{t}}_k)^{-1} \tilde{\mathbf{t}}_k^\top \tilde{\mathbf{p}}_k$$
- 10: **end for**
- 11: Compute the final estimate as the mean:

$$\hat{\beta} = \frac{1}{K} \sum_{k=1}^K \hat{\beta}_k$$

The main steps of DML-based effect estimation for the partially linear case are reported in Algorithm 1, where \mathbf{a}_i represents the i th vector of confounding variables, t_i is the value of the treatment for the i th observation, and p_i is the corresponding price value. The function f represents the modelled relationship between p and \mathbf{a} , while g represents the modelled relationship between t and \mathbf{a} . It should be noted that the models trained in steps 5 and 6 of Algorithm 1 can be any machine learning model, allowing for flexibility in capturing complex relationships between variables. In our implementation, we employed LightGBM regression models⁴⁰.

Locally partially linear DML

Our *locally partially linear DML* approach makes use of a boxcar kernel to estimate local effects (CATEs), instead of a global ATE. The key steps of the estimation procedure are described in Algorithm 2, where \mathbf{x} is the conditioning variable used to contextualise the estimation of the causal effect. In our implementation, we used kernels with a size h of 10,000 observations, and a step size s of 1,000 observations. The smoothing parameter of the Gaussian filter was set to 1.5 times the standard deviation of the mean effect estimates. The number of bootstrap replications within each boxcar kernel was set to 100.

Algorithm 2 Locally partially linear DML

- 1: **Input:** dataset $D = \{(\mathbf{x}_i, \mathbf{a}_i, t_i, y_i)\}_{i=1}^n$, kernel size h , step size s , number of bootstrap replications B , number of folds K
 - 2: **Output:** nonlinear CATE estimate
 - 3: **for** each fitting point \mathbf{x}_u **do**
 - 4: Select subset $\mathcal{W}(\mathbf{x}_u) = \{(\mathbf{x}_i, \mathbf{a}_i, t_i, p_i), \|\mathbf{x} - \mathbf{x}_u\|_2 \leq \frac{h}{2}\}$ ▷ Contextual subset of the data
 - 5: **for** each bootstrap iteration $b = 1, \dots, B$ **do**
 - 6: Draw a resampled dataset $\mathcal{W}^b(\mathbf{x}_u)$ from $\mathcal{W}(\mathbf{x}_u)$ with replacement ▷ Bootstrap resampling
 - 7: Randomly partition the dataset into K folds. ▷ Split the data
 - 8: **for** each fold $k = 1, \dots, K$ **do** ▷ Train predictive models
 - 9: Train a model \hat{f}_{-k}^b using $K - 1$ folds to predict the response p from the confounders \mathbf{a} .
 - 10: Train a model \hat{g}_{-k}^b using $K - 1$ folds to predict the treatment t from the confounders \mathbf{a} .
 - 11: Use \hat{f}_{-k}^b and \hat{g}_{-k}^b to predict \mathbf{p} and \mathbf{t} in the held-out fold k .
 - 12: Compute residuals ▷ Orthogonalisation
 - 13: Regress $\tilde{\mathbf{p}}_k^b$ on $\tilde{\mathbf{t}}_k^b$ using OLS ▷ Effect estimation
 - 14:
$$\hat{\beta}_k^b(\mathbf{x}_u) = \text{OLS}(\tilde{\mathbf{p}}_k^b, \tilde{\mathbf{t}}_k^b) = (\tilde{\mathbf{t}}_k^{b\top} \tilde{\mathbf{t}}_k^b)^{-1} \tilde{\mathbf{t}}_k^{b\top} \tilde{\mathbf{p}}_k^b$$
 - 15: **end for**
 - 16: Compute the mean estimate for the current bootstrap sample $\mathcal{W}^b(\mathbf{x}_u)$ as:
 - 17:
$$\hat{\beta}^b(\mathbf{x}_u) = \frac{1}{K} \sum_{k=1}^K \hat{\beta}_k^b(\mathbf{x}_u)$$
 - 18: **end for**
 - 19: Compute the mean estimate for the current kernel $\mathcal{W}(\mathbf{x}_u)$ as:
 - 20:
$$\hat{\beta}(\mathbf{x}_u) = \frac{1}{B} \sum_{b=1}^B \hat{\beta}^b(\mathbf{x}_u)$$
 - 21: Aggregate $\hat{\beta}(\mathbf{x}_u)$ across all kernel centres \mathbf{x}_u to construct the nonlinear CATE estimate
 - 22: Smooth the aggregated CATE estimates to reduce noise using a Gaussian filter ▷ Optional
-

Supplementary figures

Results from NordPool and intraday markets

NordPool prices. To establish the robustness of our locally partially linear DML framework, we extend the results presented in the paper to the NordPool (day-ahead) and the intraday prices. Here, we estimate the causal effects of wind and solar generation on electricity prices, highlighting differences between observed trends and true causal impacts. Figure S1 compares observational mean trends with smoothed nonlinear CATEs for wind and solar power on electricity prices. These results confirm the behaviour observed in the case of APX prices, where both wind and solar power generation exert a significant downward pressure on the electricity prices.

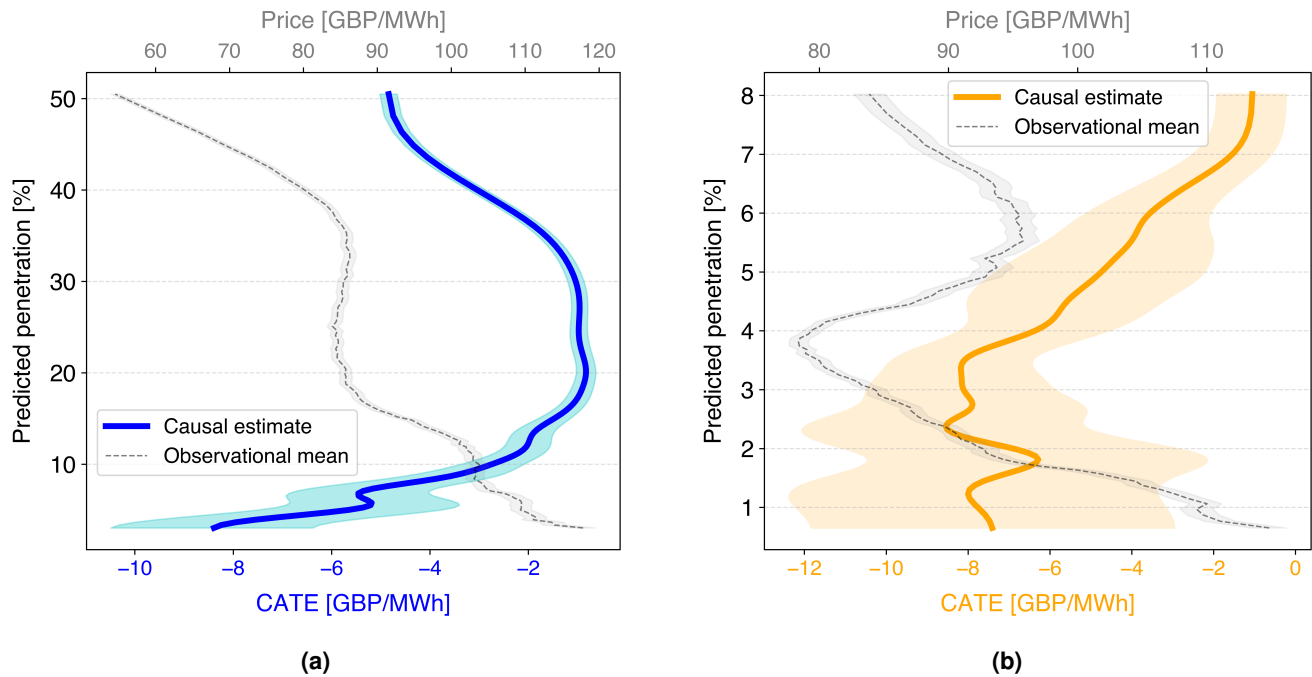


Figure S1. Results from the NordPool day-ahead market. Comparison of observational mean and causal effects of wind (a) and solar (b) power production on wholesale electricity prices. The solid lines represent the nonlinear CATE estimates derived using our locally partially linear DML framework, capturing the price impact [GBP/MWh] of a 1 GWh increase in renewable energy generation. The dashed lines show the observational mean trends as a function of renewable penetration levels, illustrating the differences between raw associations and the true causal effects. Shaded areas denote 80% confidence intervals obtained with bootstrap procedures.

Intraday prices. We further extend the analysis to the intraday market, as shown in Figure S2, where the smoothed CATE estimates illustrate the impacts of predicted renewable power generation on intraday electricity prices. Here, it is possible to notice how while wind power seems to have a quite comparable behaviour, solar power seems to have a much reduced effect on intraday prices, compared to the day-ahead markets.




Figure S2. Causal impact of predicted renewable power production on electricity prices (intraday market).

Comparison of observational mean and causal effects of wind (a) and solar (b) power production on wholesale electricity prices. The solid lines represent the nonlinear CATE estimates derived using our locally partially linear DML framework, capturing the price impact (GBP/MWh) of a 1 GWh increase in renewable energy generation. The dashed lines show the observational mean trends as a function of renewable penetration levels, illustrating the differences between raw associations and the true causal effects. Shaded areas denote 80% confidence intervals obtained with bootstrap procedures.

Using penetration as input

Here, we attempt to estimate the CATE using predicted renewable penetration as the input variable. Previously, predicted penetration was used as a contextual variable within the boxcar kernel, slicing the data into subsets to estimate localised effects. In that framework, the CATE represented the price impact of adding 1 extra GWh of renewable energy generation to the system. Here, we present an alternative analysis where the CATE reflects the price impact of increasing renewable penetration by 1 percentage point (%). Figure S3 shows the results of applying our DML framework with predicted renewable penetration serving both as the input variable for the final model and as a contextual variable for the boxcar kernel. This approach introduces potential challenges, particularly regarding additional confounding factors. Normalising the predicted renewable production by the total estimated load to compute predicted penetration may introduce dependencies that are not fully accounted for by the standard confounders used in residualising production. This added confounding underscores the need for caution when interpreting these results. Unlike penetration, renewable production forecasts are primarily driven by well-understood factors such as weather conditions and seasonal patterns, making them less susceptible to hidden confounders. To mitigate these issues, we applied residualisation to predicted penetration using the same confounders employed for renewable production—namely, wind (or solar) capacity, daylight hours, hour of the day, and month of the year. Additionally, we included estimated load and predicted renewable production (wind or solar) as confounders to address dependencies introduced by the penetration calculation.




Figure S3. Effect of the predicted penetrations. Nonlinear CATE estimates derived using our locally partially linear DML framework, capturing the price impact [GBP/MWh] of a 1% increase in renewable penetration.

In Figure S3, the smoothing factor of the Gaussian filter was set equal to the standard deviation of the mean CATE. For both wind and solar, we observe increased variability in the estimates. For wind, this variability is particularly pronounced below 10% penetration. For solar, the variability remains consistently higher across the entire range of predicted penetration. This heightened variability is likely due to the smaller effect size, making the estimates more susceptible to confounding factors introduced during the normalisation process. Nevertheless, the mean causal estimates remain negative throughout the range, reaffirming the merit order effect, whereby increased renewable penetration reduces electricity prices.

Individual effect estimates prior to the smoothing

We now provide a visualisation for the individual CATE estimates obtained from each bootstrap iteration of our causal inference framework. These estimates represent the causal impact, in GBP/MWh, of a 1 GWh increase in renewable energy generation. Each data point corresponds to the output of our locally partially linear DML method prior to any smoothing. Figure S4 illustrates the distribution of individual CATE estimates for wind and solar generation. Meanwhile, Figure S5 shows the temporal evolution of these estimates across all windows, highlighting variability over time.






Figure S4. Individual CATE estimates prior to smoothing. Individual CATE estimates of the effect of wind (a) and solar (b) power production on electricity prices.



(a)



(b)

Figure S5. Temporal evolution of the individual CATE estimates of wind (a) and solar (b) power production on electricity prices.

Supplementary tables

Table S1 provides a comprehensive overview of the variables employed in our DML framework, emphasising their data sources and roles across different models. The dataset integrates observations from the EnAppSys platform⁴¹, direct inputs from the National Energy System Operator (NESO), and publicly available data from the NESO Data Portal⁴².

Variable	Description	Source	Price (wind model)	Wind forecast	Price (solar model)	Solar forecast
Date	Date of the observation (datetime index). One observation for each 30-minute settlement period.	EnAppSys	×	×	×	×
Year	Year of observation.	Derived	✓	×	✓	×
Month	Month of the year (1-12).	Derived	✓	✓	✓	✓
Day	Day of the week (0-6, Monday to Sunday).	Derived	✓	×	✓	×
Hour	Hour of the day (0-23).	Derived	✓	✓	✓	✓
Daylight hours	Number of hours of daylight based on geographical location. It has been computed using the Python Astral package ⁴³ .	Derived	✓	✓	✓	✓
APX price	Day-ahead electricity price on the APX exchange [GBP/MWh].	EnAppSys	×	×	×	×
NordPool price	Day-ahead electricity price on the NordPool exchange [GBP/MWh].	EnAppSys	×	×	×	×
Intraday price	Intraday or within-day (MIDP) electricity price [GBP/MWh].	EnAppSys	×	×	×	×
Actual load	Initial transmission system demand outturn [MW].	EnAppSys	×	×	×	×
Estimated load	Estimated electricity load, generated from actual demand with noise [MW].	Derived	✓	×	✓	×
Gas price	National Balancing Point (NBP) price for natural gas [GBP/MWh].	EnAppSys	✓	×	✓	×
Carbon permits	MID CO ₂ prices from the EU Emissions Trading Scheme [GBP/tCO ₂ e].	EnAppSys	✓	×	✓	×
Wind capacity	Installed wind capacity [MW].	NESO	×	✓	×	×
Wind forecast	Predicted wind production [MW].	NESO	×	×	✓	×
Predicted wind penetration	Obtained as: wind forecast / estimated load * 100 [%].	Derived	×	×	×	×
Solar capacity	Installed solar capacity [MW].	NESO	×	×	×	✓
Solar forecast	Predicted solar production [MW].	NESO	✓	×	×	×
Predicted solar penetration	Obtained as: solar forecast / estimated load * 100 [%].	Derived	×	×	×	×

Supplementary Table S1. Dataset Description with variables used in the DML framework. The check marks (✓) indicate whether the variable in the row was used to residualise the variable in the column. For example, if a variable is marked with a ✓ under “Price (wind model)”, it means this row’s variable was used as a confounder in residualising the electricity price in the wind model.

References

1. IEA. Net zero roadmap: A global pathway to keep the 1.5 °c goal in reach – 2023 update. Technical Report, International Energy Agency (2023).
2. National Grid. How much of the uk’s energy is renewable (2023). Accessed: 2023-04-12.
3. IPCC. Climate change synthesis report 2023 – summary for policymakers. In Lee, H. & Romero, J. (eds.) *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, 1–34 (IPCC, Geneva, Switzerland, 2023).
4. Morales, J. M., Conejo, A. J., Madsen, H., Pinson, P. & Zugno, M. *Integrating renewables in electricity markets: operational problems*, vol. 205 (Springer Science & Business Media, 2013).
5. Antweiler, W. & Muesgens, F. On the long-term merit order effect of renewable energies. *Energy Econ.* **99**, 105275 (2021).
6. Miettinen, J. & Holttinen, H. Impacts of wind power forecast errors on the real-time balancing need: a Nordic case study. *IET Renew. Power Gener.* **13**, 227–233 (2019).
7. Goodarzi, S., Perera, H. N. & Bunn, D. The impact of renewable energy forecast errors on imbalance volumes and electricity spot prices. *Energy Policy* **134**, 110827 (2019).
8. Morthorst, P. E. Wind power and the conditions at a liberalized power market. *Wind. Energy* **6**, 297–308 (2003).
9. Jónsson, T., Pinson, P. & Madsen, H. On the market impact of wind energy forecasts. *Energy Econ.* **32**, 313–320 (2010).
10. Cutler, N. J., Boerema, N. D., MacGill, I. F. & Outhred, H. R. High penetration wind generation impacts on spot prices in the Australian national electricity market. *Energy Policy* **39**, 5939–5949 (2011).
11. Forrest, S. & MacGill, I. Assessing the impact of wind generation on wholesale prices and generator dispatch in the Australian National Electricity Market. *Energy Policy* **59**, 120–132 (2013).
12. Bell, W. P., Wild, P., Foster, J. & Hewson, M. Revitalising the wind power induced merit order effect to reduce wholesale and retail electricity prices in Australia. *Energy Econ.* **67**, 224–241 (2017).
13. Würzburg, K., Labandeira, X. & Linares, P. Renewable generation and electricity prices: Taking stock and new evidence for Germany and Austria. *Energy Econ.* **40**, S159–S171 (2013).
14. Ketterer, J. C. The impact of wind power generation on the electricity price in Germany. *Energy economics* **44**, 270–280 (2014).
15. Gürtler, M. & Paulsen, T. The effect of wind and solar power forecasts on day-ahead and intraday electricity prices in Germany. *Energy Econ.* **75**, 150–162 (2018).
16. Do, L. P. C., Lyócsa, Š. & Molnár, P. Impact of wind and solar production on electricity prices: Quantile regression approach. *J. Oper. Res. Soc.* **70**, 1752–1768 (2019).
17. Keeley, A. R., Matsumoto, K., Tanaka, K., Sugiawan, Y. & Managi, S. The impact of renewable energy generation on the spot market price in Germany: ex-post analysis using boosting method. *The Energy J.* **42**, 1–22 (2021).
18. O’Mahoney, A. & Denny, E. The merit order effect of wind generation on the irish electricity market. *MPRA Pap.* (2011).
19. O’Flaherty, M., Riordan, N., O’Neill, N. & Ahern, C. A quantitative analysis of the impact of wind energy penetration on electricity prices in Ireland. *Energy Procedia* **58**, 103–110 (2014).
20. Clò, S., Cataldi, A. & Zoppoli, P. The merit-order effect in the Italian power market: The impact of solar and wind generation on national wholesale electricity prices. *Energy Policy* **77**, 79–88 (2015).
21. Shcherbakova, A., Kleit, A., Blumsack, S., Cho, J. & Lee, W. Effect of increased wind penetration on system prices in Korea’s electricity markets. *Wind. Energy* **17**, 1469–1482 (2014).
22. Nieuwenhout, F. & Brand, A. The impact of wind power on day-ahead electricity prices in the Netherlands. In *2011 8th International Conference on the European Energy Market (EEM)*, 226–230 (IEEE, 2011).
23. Hu, X., Jaraitė, J. & Kažukauskas, A. The effects of wind power on electricity markets: A case study of the Swedish intraday market. *Energy Econ.* **96**, 105159 (2021).
24. Spodniak, P., Ollikka, K. & Honkapuro, S. The impact of wind power and electricity demand on the relevance of different short-term electricity markets: The Nordic case. *Appl. energy* **283**, 116063 (2021).
25. Tselika, K., Tselika, M. & Demetriades, E. Quantifying the short-term asymmetric effects of renewable energy on the electricity merit-order curve. *Energy Econ.* **132**, 107471 (2024).

26. Gelabert, L., Labandeira, X. & Linares, P. An ex-post analysis of the effect of renewables and cogeneration on Spanish electricity prices. *Energy economics* **33**, S59–S65 (2011).
27. Cruz, A., Muñoz, A., Zamora, J. L. & Espínola, R. The effect of wind generation and weekday on Spanish electricity spot price forecasting. *Electr. Power Syst. Res.* **81**, 1924–1935 (2011).
28. Azofra, D., Jiménez, E., Martínez, E., Blanco, J. & Saenz-Díez, J. Wind power merit-order and feed-in-tariffs effect: A variability analysis of the Spanish electricity market. *Energy Convers. Manag.* **83**, 19–27 (2014).
29. Barthelmie, R., Murray, F. & Pryor, S. The economic benefit of short-term forecasting for wind energy in the UK electricity market. *Energy Policy* **36**, 1687–1696 (2008).
30. Joos, M. & Staffell, I. Short-term integration costs of variable renewable energy: Wind curtailment and balancing in Britain and Germany. *Renew. Sustain. Energy Rev.* **86**, 45–65 (2018).
31. Woo, C.-K., Horowitz, I., Moore, J. & Pacheco, A. The impact of wind generation on the electricity spot-market price level and variance: The Texas experience. *Energy Policy* **39**, 3939–3944 (2011).
32. Woo, C.-K. *et al.* Merit-order effects of renewable energy and price divergence in California’s day-ahead and real-time electricity markets. *Energy Policy* **92**, 299–312 (2016).
33. Martinez-Anido, C. B., Brinkman, G. & Hodge, B.-M. The impact of wind power on electricity prices. *Renew. energy* **94**, 474–487 (2016).
34. Prol, J. L., Steininger, K. W. & Zilberman, D. The cannibalization effect of wind and solar in the California wholesale electricity market. *Energy Econ.* **85**, 104552 (2020).
35. Stiewe, C., Xu, A. L., Eicke, A. & Hirth, L. Cross-border cannibalization: Spillover effects of wind and solar energy on interconnected European electricity markets. *arXiv preprint arXiv:2405.17166* (2024).
36. Chernozhukov, V. *et al.* Double/debiased machine learning for treatment and structural parameters (2018).
37. Pearl, J. *Causality* (Cambridge university press, 2009).
38. Peters, J., Janzing, D. & Schölkopf, B. *Elements of Causal Inference: Foundations and Learning Algorithms* (MIT Press, 2017).
39. Fuhr, J., Berens, P. & Papies, D. Estimating causal effects with double machine learning—a method evaluation. *arXiv preprint arXiv:2403.14385* (2024).
40. Ke, G. *et al.* LightGBM: A highly efficient gradient boosting decision tree. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 3149–3157 (2017).
41. EnAppSys. EnAppSys market data platform. <https://www.enappsys.com/> (2024).
42. NESO data portal. <https://www.neso.energy/data-portal> (2024).
43. Kennedy, S. Astral: a python package for calculating the times of various aspects of the sun and moon. <https://sffjunkie.github.io/astral/> (2023). Version 3.3.2, Python library for astronomical calculations.