

Hirth, Lion; Khanna, Tarun; Ruhnau, Oliver

Working Paper

The (very) short-term price elasticity of German electricity demand

Suggested Citation: Hirth, Lion; Khanna, Tarun; Ruhnau, Oliver (2022) : The (very) short-term price elasticity of German electricity demand, ZBW - Leibniz Information Centre for Economics, Kiel, Hamburg

This Version is available at:

<https://hdl.handle.net/10419/249570>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The (very) short-term price elasticity of German electricity demand

Lion Hirth, Tarun Khanna, Oliver Ruhnau*

Hertie School, Berlin, Germany

*corresponding author: ruhnau@hertie-school.org

Working Paper

Abstract. Electricity is a peculiar economic good, the most important reason being that it needs to be supplied at the very moment of consumption. As a result, wholesale electricity prices fluctuate widely at hourly or sub-hourly time scales, regularly reaching multiples of their average, and even turn negative. This paper examines whether the demand for electricity responds to such price variations in the very short term. To solve the classical identification problem when estimating a demand curve, we use weather-driven wind energy generation as an instrument. Our robustness checks confirm that wind energy is indeed a strong and valid instrument. Using data from Germany, we estimate that a 1 €/MWh increase in the wholesale electricity price causes the aggregate electricity demand to decline by 67–80 MW or 0.12–0.14%, contradicting the conventional wisdom that electricity demand is highly price-inelastic. These estimates are statistically significant and robust across model specifications, estimators, and sensitivity analyses. At average price and demand, our estimates correspond to a price elasticity of demand of about -0.05 . Comparing situations with high and low wind energy (5–95th percentile), we estimate that prices vary by 26 €/MWh, and the corresponding demand response to wholesale electricity prices is about 2 GW, or 2.6% of peak load. Our analysis suggests that the demand response in Germany can be attributed primarily to industrial consumers.

Keywords: Electricity markets, price elasticity, demand response, instrumental variables, wind energy

Acknowledgements

The authors thank Christian Traxler, Golnoush Soroush, Lars Nolting, and William Zappa, as well as the participants of the 10th INREC conference, the 6th AIEE Energy Symposium, the Hertie Energy Research Seminar, and the Research Meeting at the Florence School of Regulation, for their valuable comments and inspiring discussions.

1 Introduction

The wind and solar challenge. Energy systems worldwide are undergoing a rapid transition toward wind and solar energy. However, further expansion of these variable renewable energy sources and their integration into energy systems and markets is impeded by the fact that they are not always available when electricity is needed. In addition to flexible generation, long-distance transmission, and energy storage, demand-side flexibility is hoped to be part of the solution to such intermittency. Electricity demand that dynamically responds to the availability of cheap renewable electricity would facilitate the transition to a low-carbon power system.

Demand elasticity. While dedicated “demand response programs” exist in many power systems, they are often designed as emergency response mechanisms. During normal operation, the most important mechanism to coordinate supply and demand is represented by wholesale markets for electricity. Wholesale markets clear at high frequency, matching supply and demand at hourly or sub-hourly intervals. Economists refer to consumers adjusting the quantity of power consumed in response to prices as the “price elasticity of demand”; energy system researchers often refer to this as an (implicit) demand response. Conventional wisdom holds that electricity demand is very price-inelastic, which has profound implications for electricity markets. Cramton et al. (2013) summarized this as follows:

One particularly notorious problem of electricity markets is low demand flexibility. Most customers are unaware of the real time prices of electricity, have no reason to respond to them, or cannot respond quickly to them, leading to highly price-inelastic demand. This contributes to blackouts in times of scarcity and to the inability of the market to determine the market-clearing prices needed to attract an efficient level and mix of generation capacity. Moreover, the problems caused by this market failure can result in considerable price volatility and market power that would be insignificant if the demand-side of the market were fully functional.

While it is true for many jurisdictions that most consumers are unaware of hourly variations in real-time prices, some large industrial consumers directly participate in wholesale markets. These consumers may indeed respond to hourly price changes. In this empirical paper, we study how much of such a short-term demand response exists today—that is, the price elasticity of electricity demand in the current power system.

Exogenous variation. During windy or sunny times, the additional low-cost energy supply from wind turbines and solar photovoltaics reduces wholesale electricity prices, which, if the demand is price-elastic, causes an increase in electricity consumption compared to the no-renewables counterfactual. Likewise, a lack of renewable generation may cause a decrease in consumption through higher wholesale prices. We can therefore use the exogenous weather-dependent electricity supply as an instrument for wholesale prices to estimate the responsiveness of aggregate (national) electricity demand price changes. This is the strategy we use to overcome the classical simultaneity problem when estimating demand curves from market outcomes.

Time scales. There is a large body of literature on electricity demand elasticity. The vast majority examined how annual demand responds to changes in retail prices (for a review, see Labandeira et al., 2017). Somewhat confusingly, this line of research identifies a “short-term” response if price changes affect demand in the same year, presumably a consequence of substitutions and geographical shifts in the production of energy-intensive goods (for recent examples, see Cialani and Mortazavi, 2018; Csereklyei, 2020). Accordingly, “long-term” price elasticity refers to responses over multiple years, possibly driven by investment patterns, such as the choice of heating systems in new buildings (Davis, 2021). We are interested in much shorter time scales, such as hours. We address the question: by how much does the aggregate electricity demand respond to (very) short-term price fluctuations, if at all?

Short-term demand response potential. Many publications discuss the technical possibilities and theoretical potential of electricity loads to adjust consumption (for reviews, see Sauer et al., 2019; Kochems, 2020). Energy-intensive industries, such as the chemical, metal, and paper industries, can shut down production processes when high prices turn their margins negative (“load shedding”). Other processes may entail some flexibility to shift load toward hours with low prices while keeping the overall output constant. Such “load shifting” is particularly promising in heating and cooling, exploiting thermal inertia. Bottom-up studies on the theoretical demand response potential are often based on techno-economic models (e.g., Gils, 2014) or choice experiments (e.g., Olsthoorn et al., 2019). While these studies provide detailed insights into how demand *could* respond to prices, it remains unclear how much of this potential is exploited in today’s wholesale electricity markets. Such studies also often fail to convincingly estimate the costs of load shedding and shifting.

Individual consumers. A number of empirical papers provide evidence of the actual demand response of *individual* consumers. Faruqi and Sergici (2013) summarized international evidence on the demand response of households from 34 experimental studies. Other studies have focused on industrial consumers, analyzing their (sub-)hourly consumption patterns (Patrick and Wolak, 2001; Hopper et al., 2006; Zarnikau and Hallett, 2008; Choi et al., 2011). While these studies provide evidence that some consumers do respond to prices, it remains unclear how representative this is for a country and how well this extends beyond the quasi-experimental setting of some of these studies. What ultimately matters for electricity systems and markets is, of course, aggregate demand.

Aggregate demand. This study estimates the price elasticity of aggregate (country-wide) electricity demand—the demand curve—on time scales of hours using wind energy as an instrumental variable. We are aware of only a handful of previous studies that have estimated the elasticity of aggregate demand. The first of these is that of Lijesen (2007), who used data from the Netherlands for 2003 to estimate an exponential demand curve and found a demand elasticity of -0.0014 . A major concern with this paper is endogeneity because the lagged price is used as an instrumental variable. Given the strong serial correlation between high-frequency electricity system data and intertemporal relationships, such as energy storage, thermal inertia in technical systems, and expectations, the exogeneity assumption seems unlikely to hold.

Bönte et al. (2015) estimated demand elasticity at -0.43 , based on ask data from Germany’s leading electricity exchange, EPEX SPOT. However, the demand curve at the power exchange cannot be interpreted as the aggregate (national) demand curve for electricity. They are not identical because, on one hand, the power exchange captures neither volumes traded in bilateral agreements nor electricity produced by integrated generation-retailing companies (in our sample period, EPEX SPOT trading volumes were about half the German electricity demand). On the other hand, traded volume on the exchange includes energy that is traded several times across market stages (financial markets included, the churn rate of German electricity is around 5). Kulakov and Ziel (2019) used the same data and applied a decomposition, yielding an elasticity estimate of -0.0001 . However, the paper does not use instrumental variables, and it is unclear whether endogeneity is adequately addressed. Similarly, Damien et al. (2019) used a hierarchical Bayesian model and data from Texas to regress demand on prices, ignoring endogeneity. They found a positive relationship between demand and prices, which we believe most likely describes the supply curve better than the demand curve: high demand increases prices because of increasing marginal costs.

This leaves us with one paper, that of Knaut and Paulus (2016), which takes an approach similar to ours, using wind energy as an instrument. Using German data from 2015, the authors reported an elasticity of -0.02 to -0.13 . Although we regard their approach as the most reliable one brought forward so far, the paper has never been published in a peer-reviewed journal, and we think it can be

extended in several ways—in particular, by responding to concerns about the validity of the instrument.

Contribution. Our contribution includes five major aspects. First, we develop a theoretical framework for the causal identification of demand response, including an in-depth discussion of how prices impact load, given the institutional setup of the German electricity market. Second, we update and extend the scope of the analysis to five years (2015–2019). This is important for ascertaining the robustness of the analysis and time trends in demand elasticity. Third, we use different model specifications, including a purely linear specification (as in Knaut and Paulus, 2016), a log-linear specification, and specifications with nonparametric elements. This matters because there are many reasons to expect nonlinear relationships in electricity systems. The models are estimated using a 2-stage least squares (2SLS) and a 2-stage generalized additive models (2SGAM). This allows us to study the shape of the demand curve. Fourth, we study the heterogeneity of demand response across time, including the time of the day, day of the week, and season of the year, as well as across different regions, and compared the spatial heterogeneity to the uneven distribution of industrial load across Germany. Finally, we conduct a broad range of sensitivity analyses and robustness checks. Importantly, we perform a long list of tests and analyses concerning challenges to the exogeneity of the instrument, including using alternative and additional instruments and excluding extreme events. Other analyses, such as running the model on first differences, provide insight into the role of lagged effects in the context of the short-term price elasticity of electricity demand.

Findings. We estimate that a 1 €/MWh increase in day-ahead electricity price in the wholesale markets causes the aggregate electricity demand in Germany to decline by 67–80 MW (linear estimates) or 0.12–0.14% (log-linear estimates). At the average price and demand, these estimates correspond to a price elasticity of demand of about -0.05 . Comparing situations with high and low wind energy (5–95th percentile), we estimate that prices vary by 26 €/MWh, and the corresponding demand response is about 2 GW (or 2.6% of peak demand). Our estimates do not vary significantly across years or seasons. We find a strongly nonlinear demand curve during nighttime and a quite linear curve during daytime hours. Our estimates of the regional distribution of demand response match with the regional distribution of energy-intensive industries in Germany, matching our prior assumption that it is the industrial consumers who are in a position to respond to wholesale electricity prices in Germany. The results are statistically significant and robust across a broad range of model specifications, estimators, time periods, alternative instruments, and other sensitivity analyses.

Relevance. Empirical estimates of the price elasticity of electricity demand are relevant for multiple reasons. First and foremost, many theoretical studies assume that demand is price-inelastic, and regulatory and policy recommendations are often made based on this premise. For example, government-mandated capacity mechanisms are often justified by a lack of demand response. Second, demand response is an important parameter in resource adequacy studies and power system planning. Essentially, if demand is reduced as prices peak, less generation capacity is needed for the same level of security of supply. Demand elasticity also matters for operational decisions. Short-term electricity price forecasting, dispatch planning, and grid modeling can all be improved based on more precise information about demand response. Finally, this parameter can be readily used in energy system models for long-term scenarios.

Outline. We continue with introducing our dataset in Section 2 before explaining our identification strategy in Section 3. Section 4 shows the model specifications and estimators, and Section 5 presents the results. In Section 6, we discuss the results, including how they reflect the institutional setup of the German electricity market, and draw conclusions.

2 Data

Variables considered. This section gives an overview of the data used to estimate the price elasticity of electricity demand. The main parameters are the hourly time series of the aggregate electricity demand and the wholesale electricity prices in Germany. To overcome the endogeneity between demand and price, wind energy generation is used as an instrument variable. The analysis controls for possible confounding influences exerted by other variables on the price-demand relationship, including solar electricity generation and temperature, as well as coal, gas, and carbon prices.

Data sources and exclusion of holidays. Hourly time series for German wholesale electricity prices, electricity demand, wind and solar generation, ambient temperature, and other weather variables were obtained from the Open Power System Data platform (Wiese et al., 2019). This platform aggregates data from different primary sources, which are mentioned below. We exclude holidays, bridge days (days between holidays and weekends), and the Christmas vacation period (between Christmas and New Year) from the analysis, as the demand and price patterns on these days are distinct and not representative of the general relationship between prices and demand.

2.1 Electricity demand

Aggregate demand. We use the “Total Load” data series published by the European Network of Transmission System Operators (ENTSO-E) for measuring the aggregate electricity demand in Germany (for details, see Schumacher and Hirth, 2015; Hirth et al., 2018). This is calculated as “Net Generation – Exports + Imports – Absorbed Energy,” where the absorbed energy is the electricity used to charge storage.¹

Completeness of demand data. By definition, the “Total Load” data should include demand that is met by solar and industrial generators on consumption sites (self-generation or behind-the-meter generation). However, we are not entirely sure that this part of the load is completely and correctly estimated. We suspect that the data may be incomplete because of a gap between the annual sum of the “Total Load” time series by ENTSO-E and the annual “Final Consumption” statistics from EUROSTAT, although both follow the same definition for load/consumption. While the ENTSO-E time series are the best source we know for hourly demand data, the annual EUROSTAT data are presumably more complete. We expect the cause of this gap to be related to the load served by behind-the-meter generators, which can be distinguished into two groups: rooftop solar photovoltaics and industrial self-generation, the latter of which often comes with heat cogeneration. We discuss below how uncertainty regarding the self-generation of these two groups could affect our identification strategy. It should be mentioned, however, that the gap between the ENTSO-E and EUROSTAT data decreases over time.²

Small-scale solar. Just like wind generation, solar energy depresses electricity prices and hence, in principle, could be used as an additional instrument. We have decided not to do so in our main specification for the following reason. Much of the solar capacity represents small-scale, decentralized generation. These generators are usually not metered at hourly resolution. To calculate “Total Load,” system operators estimate the hour-by-hour solar generation. Estimation errors that are correlated with prices could induce a relationship between prices and observed demand that is not due to the demand response. This would violate the exclusion restriction, inflating our estimates of demand elasticity. In all our regressions, however, we do control for solar to account for any possible correlation

¹https://transparency.entsoe.eu/content/static_content/Static%20content/knowledge%20base/data-views/load-domain/Data-view%20Total%20Load%20-%20Day%20Ahead%20-%20Actual.html

² From about 50 TWh in 2015 to about 20 TWh in 2019.

with wind generation and demand. In addition, we run a sensitivity test of solar energy as an additional instrument (Section 5.5).

Industrial self-generation. Industrial self-generation could be positively correlated with prices if industrial consumers decided to shut down their power plants at low prices and buy from wholesale markets instead. If this was not properly accounted for in the total load data, such a dispatch decision on self-generation would introduce an upward bias in our demand-response estimates. Because the gap between the ENTSO-E and EUROSTAT data decreases over time, while our estimates remain consistent over the years (see Section 5.2), we are confident that the demand responsiveness that we are measuring is not entirely due to industrial self-generation. In addition, we conducted the following analysis: the German regulator (Bundesnetzagentur) maintains a public list of power plants containing 1,980 units. A manual review resulted in 366 units that could potentially be industrial behind-the-meter plants, according to the reported names of the operators. Of those, 28 have a nameplate capacity of 100 MW or larger, meaning they should report hour-by-hour generation on the ENTSO-E Transparency Platform. We were able to identify such time series for nine units. On the aggregate of these nine time series, we ran the same 2SLS model as the one we present and apply below for our main analysis. In doing so, we found that a 1 €/MWh price increase leads to a 2.6 MW increase in production, with a confidence interval of 2.0 MW to 3.3 MW. Scaling the estimated price response of our sample by the total annual output of industrial self-generation in Germany³ yields a price response of the German-wide industrial self-generation of around 8 MW per €/MWh, just above a tenth of the price elasticity we estimate in our main analysis. In other words, this assessment suggests that 90% of our main estimate of demand response is actually an adjustment of consumption processes, while 10% reflects industrial behind-the-meter generators.

Temporal patterns in aggregate demand. Figure 1 illustrates the temporal patterns of the aggregate German electricity demand used in our analysis. There are two distinct diurnal peaks in electricity consumption, one in the morning at 11 a.m. and the other in the evening at 6 p.m., and demand is significantly lower on weekends compared to weekdays, reflecting lower commercial and industrial demand. Electricity consumption is also higher during the winter months compared to other times of the year. The aggregate demand for Germany was relatively stable across the years considered in our analysis. There is only a marginal underlying positive trend from 2015 to 2018 and a decline in 2019. A linear least squares model with demand as the dependent variable shows that the hourly, monthly, weekday, annual, and time trends together explain 83% of the variation in hourly aggregate demand in Germany.

³ The nine plants have a joint annual output of about 8 TWh compared to the total industrial self-generation of about 25 TWh (according to the German Association of Energy and Water Industries).




Figure 1: Boxplots of the electricity demand by year, month, weekday, and hour⁴

2.2 Electricity prices

Day-ahead wholesale prices. We use data from the German day-ahead auction at the EPEX power exchange to measure electricity prices. In Germany, the EPEX day-ahead auction is the primary wholesale market for the physical delivery of electricity. The auction clears supply and demand for each hour of the next day separately, yielding electricity prices at an hourly resolution. Although not all electricity consumed is traded in the day-ahead market, the day-ahead prices are the relevant benchmark for bilateral trades and self-generation because they reflect the opportunity costs of not trading in the power exchange. Furthermore, our analysis based on day-ahead prices will capture demand response to market prices at other lead times, such as intraday and balancing prices, to the extent that these prices are correlated with the day-ahead prices.

Price exposure. Most electricity consumers in Germany are not directly exposed to variations in day-ahead market prices. Instead, retail suppliers of electricity purchase electricity in the wholesale markets for them. These consumers can freely choose between retail suppliers and tariffs that reflect wholesale prices in day-ahead markets to varying degrees: from not at all (constant price during a year) to real-time pricing (prices change hourly). In Germany, constant prices are the most common option, and only a few large industrial consumers and some household consumers opt for real-time pricing. In addition, some large electricity consumers from energy-intensive industries (such as metal and chemicals) also procure electricity directly from the wholesale market and are fully exposed to variations in the day-ahead wholesale price of electricity. Therefore, our estimates of the elasticity of electricity demand to wholesale market prices must only reflect the response of consumers who are exposed to such price variations, although we cannot identify the role of individual consumers or types of consumers.

Temporal patterns in electricity prices. The nature of the time-series data of day-ahead spot prices in electricity markets has been previously studied (Knittel and Roberts, 2005; Kosater and Mosler, 2006; Huisman et al., 2007; Liebl, 2013). Usually, time series of electricity spot prices are assumed (i) to have deterministic annual variations and also monthly, weekly, and hourly patterns, (ii) to show price-dependent volatilities, and (iii) to be stationary (after controlling for seasonal patterns). These temporal trends are clearly visible in our data as well (Figure 2). The average spot price of electricity changes over time in our data, with prices falling slightly in 2016 from 2015 before increasing to a peak

⁴ The black lines in the middle of the boxes indicate the median; the boxes extend from the first to the third quartile (inter-quartile range), and the whiskers include the 5–95% confidence interval of the observations. Observations outside of this confidence interval are depicted as black dots, and the white points represent the mean of the distribution.

in 2018 and sharply falling in the first quarter of 2019. Besides this trend over the years, there are the expected intra-annual patterns. Electricity prices change over the course of the year, depending on the month. Prices are higher during the weekdays and lower over the weekends, reflecting lower demand from industrial and commercial activity. Over the course of a day, there are two distinct peaks: a morning and an evening peak around 8 a.m. and 7 p.m., respectively. Fitting a nonparametric generalized additive model (GAM) to electricity prices reveals that the nonlinear coal and gas prices and the seasonal time patterns together explain about 48% of the variation in the hourly spot price of electricity. The residuals spot prices, once the deterministic seasonal and time trends are accounted for, seem stochastic.




Figure 2: Boxplots of the electricity price by year, month, weekday, and hour⁵

Correlation with fuel prices. The monthly averages of wholesale electricity prices reflect the underlying cost of electricity generation, closely following the prices of natural gas and coal (Figure 3).⁶ We therefore include monthly fuel prices converted to EUR/MWh as control variables in our models. In some of our model specifications, we also control for a nonparametric time trend (see Section 4).

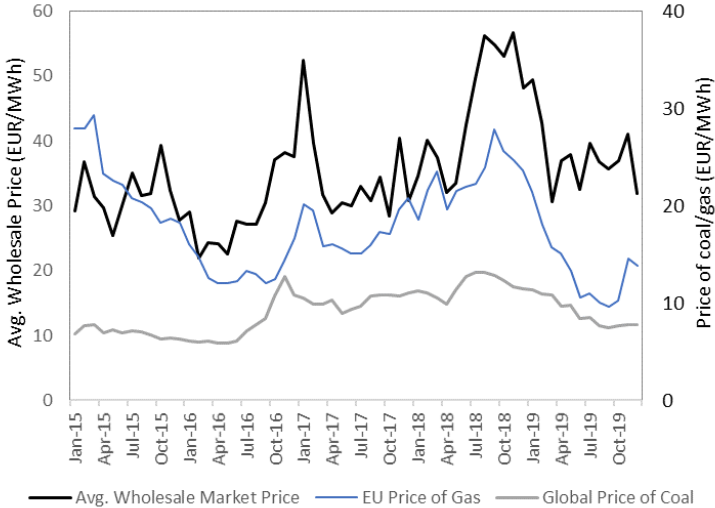


Figure 3: Day-ahead prices in wholesale markets in Germany vs. global prices of coal and gas

⁵ Note that some extreme prices are not visible because they exceed the limits of the ordinate. See Footnote 3 for an explanation of the boxplot elements.

⁶ We use the data series “Global price of Coal, Australia, U.S. Dollars per Metric Ton, Monthly, Not Seasonally Adjusted” and “Global price of Natural gas, EU, U.S. Dollars per Million Metric British Thermal Unit, Monthly, Not Seasonally Adjusted” published by the International Monetary Fund for coal and gas prices, respectively.

2.3 Other data series

Wind and solar generation. Our time series for wind and solar energy originate from the ETNSO-E Transparency Platform. They represent actual electricity generation, which has two implications. First, renewable energy is traded in the day-ahead market based on forecasts such that it is the predicted rather than the actual electricity generation that impacts day-ahead prices. We argue, however, that the actual electricity generation can still be used as an explanatory variable for day-ahead prices, given the high correlation with predicted values. Second, the actual generation may differ from the potential generation due to curtailment at negative electricity prices or grid congestion. What curtailment means for our identification strategy is discussed in Section 3.2. Furthermore, part of the solar generation is not individually metered at hourly resolution, which may introduce estimation errors. Therefore, we decided to use wind energy only as an instrument and to include solar energy only as a control variable. The temporal patterns of wind energy are displayed in Figure 4.




Figure 4: Boxplots of the wind energy generation by year, month, weekday, and hour⁷

Further control variables. In addition to fuel prices, we use the daily price of European Union Allowances (EUAs) for carbon emissions issued via the European Union Emissions Trading Scheme, which we retrieved from the Intercontinental Exchange (ICE).⁸ Furthermore, we use hourly time series for ambient temperature from renewables.ninja.⁹ They provide population-weighted averages of spatial temperature data across Germany. For the linear models, we derived the hourly heating and cooling degrees from the ambient temperature using 15°C as a threshold.

3 Causal inference

Price elasticity of demand. We are interested in estimating the price elasticity of electricity—that is, the causal effect of the wholesale electricity price from the day-ahead market (treatment) on the aggregate electricity demand (outcome) at very short time scales, such as hours. We do this by using hourly data on the day-ahead price of electricity in the wholesale markets and the hourly aggregate demand for electricity in Germany.

⁷ Note that some extreme prices are not visible because they exceed the limits of the ordinate. See Footnote 3 for an explanation of the boxplot elements.

⁸ <https://www.theice.com/>

⁹ <https://www.renewables.ninja/>

3.1 How prices impact demand

Direct procurement and real-time pricing. As discussed in Section 2.2., there are two ways in which consumers can be exposed to variations in wholesale electricity prices in Germany. On the one hand, some large electricity consumers procure electricity directly from the wholesale market (the direct causal path between wholesale prices and demand in Figure 5). On the other hand, some consumers have real-time tariffs that reflect the wholesale price variation (the indirect causal path between prices and demand in Figure 5). Using hourly day-ahead prices in the wholesale markets enables us to estimate the aggregated response of consumers with direct procurement and with real-time pricing.

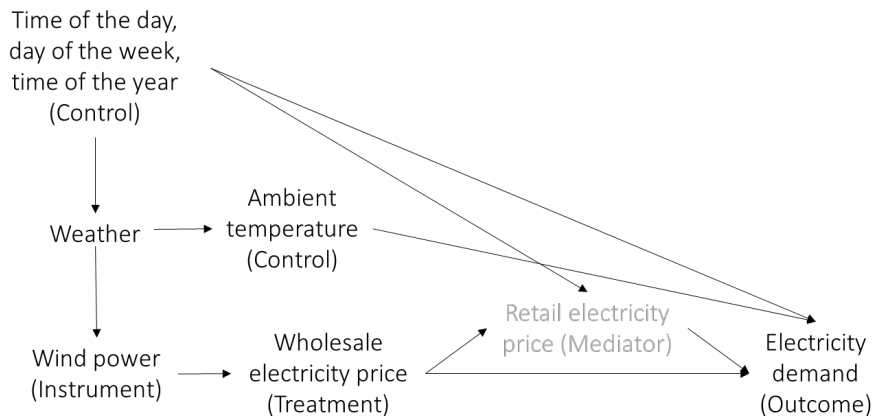


Figure 5: Main causal relationships

Time-of-use pricing. Some retail tariffs feature deterministic time-of-use pricing in the sense that the price of electricity changes in predetermined ways to reflect recurring diurnal trends in the wholesale electricity markets.¹⁰ For example, consumers may pay higher prices for the electricity consumed during the day than at night. Such tariffs do not reflect the stochastic part of the variation in wholesale prices caused by, for example, the stochasticity of wind speed. Although electricity demand may respond to such time-of-use prices, our instrumental regression using wind energy as an instrument does not capture such a response.

Additional price components. For all electricity consumers, the price of energy is only one part of what they pay as the electricity price. What comes on top are network fees, taxes, and statutory levies (renewable energy financing costs, metering costs, retail margins, etc.). We refer to these as additional price components. Typically, these price components are independent of fluctuations in the wholesale price of electricity. This has two implications for our research. First, while fluctuations in wholesale electricity prices may be fully passed through to the consumer in absolute terms, the relative variation in prices will be smaller. As our model is based on prices in absolute terms, this should not affect our estimates. Second, these price components may incentivize certain behaviors. For example, industrial network fees often comprise hefty capacity payments that are calculated based on consumption during the hour with the highest demand during the year. Therefore, the marginal cost of consumption during such hours, on top of the wholesale price and other taxes and levies, is very high. To minimize such payments, industrial consumers may try to avoid consumption peaks, even if the wholesale market prices are low. The restriction that capacity payments and other regulated price components impose on how electricity demand can respond to wholesale prices will be reflected in our estimates.

¹⁰ In Germany, time-of-use pricing is rare and mostly applies to night storage heating. We can think of time-of-use pricings as a causal path between the time variables and retail electricity prices (Figure 5).

Put differently, our estimates of the price elasticity of demand are conditional on the current regulation of additional price components in Germany.

3.2 Identification strategy

Instrument variable. There is a long tradition in economics of applying instrument variables for estimating supply and demand curves and the price elasticity of demand for various products. Because the supply and demand curves shift over time, the observed data on quantities and prices reflect a set of equilibrium points on both curves. Due to the simultaneity problem, ordinary least squares cannot be used to determine either the supply or the demand curve (Angrist and Krueger, 2001; Wooldridge, 2013). To isolate the demand curve and estimate the price elasticity of demand from time series data, we use the quantum of hourly wind energy generation as an instrument for electricity prices. To be a valid instrument, wind energy must fulfill three conditions: it must be exogenous and relevant and fulfill the exclusion restriction.

Exogeneity. To fulfill the exogeneity condition, wind generation must be unaffected by electricity prices and demand. The potential wind power supply is driven by weather conditions—more precisely, wind speed, which is certainly unaffected by electricity markets. Indeed, because of its natural stochasticity, weather variables have a long history of being used as instruments to estimate the demand elasticity of other goods, such as oil and fish (e.g., Wright, 1928; Graddy, 2006). We can also assume that wind generation is strictly exogenous in the sense that current wind generation is unaffected by past electricity market outcomes, which is important for causal inference with time series.

Wind generation vs. speed. Instead of using wind speed directly, we include wind generation in the main model specification. This is because the relationship between wind speed and wind generation is nonlinear due to the underlying S-shaped power curves of single wind turbines. Furthermore, wind turbines shut down at very high wind speeds, which may compromise the condition that the relationship between the instrument and the endogenous variable must be monotonous. However, we use wind speed as an instrument for sensitivity analyses, and we exploit the nonlinear relationship between wind speed and wind energy to perform a zero-first-stage test, finding supporting evidence for the validity of our instrument (see Section 5.6 below).

Curtailement. In contrast to the purely physically determined *potential* wind power supply, the *actual* supply may be affected by the decision not to produce, even though the wind is blowing. This may be due to two reasons: negative wholesale electricity prices or congested networks. The marginal cost of generating wind energy is virtually zero, and many wind generators receive, on top of wholesale revenues, additional income for each MWh produced (e.g., in the form of contracts-for-difference or tradable green certificates). Because of the opportunity costs related to this additional income, they will stop producing only if the wholesale price turns steeply negative. This is the case during only a few hours of the year. Transmission network congestion mostly occurs when both wind speeds and electricity demand are high. Hence, conditional on high wind speeds, the wind power supply may be somewhat lower when demand is high. In such situations, day-ahead prices may be relatively high because of the high demand,¹¹ which could cause us to *underestimate* demand elasticity. Nevertheless, we conduct two sensitivity analyses to address endogeneity concerns: one in which we exclude hours with negative prices from the dataset and one using wind *speed* instead of wind *generation* as an

¹¹ Note that day-ahead prices are determined without accounting for grid congestion such that the day-ahead price will not be affected by the lower wind supply due to curtailment.

instrument. As our estimates are robust in both cases (see Sections 5.5 and 5.7 below), we conclude that the endogeneity of wind energy is not an issue.

Relevance. To fulfill the relevance condition, wind power must be correlated with electricity prices. In economic terms, increased wind generation is a supply shock that affects the electricity price by shifting the electricity supply curve outward. An increase in wind energy generation, *ceteris paribus*, leads to a decrease in electricity prices, while a decrease in wind energy generation causes electricity prices to increase. Indeed, wind energy has a high explanatory power for price variations that cannot be explained by other covariates (see Section 5.1 below). In statistical terms, the partial R^2 of wind energy in the first stage of our two-stage least squares regression is high. Hence, wind energy is a relevant instrument for the wholesale electricity price.

Exclusion restriction. To fulfill the exclusion restriction, wind power must not affect the electricity demand other than through the wholesale electricity price. The wind energy supply is naturally stochastic because of the underlying weather patterns. Nevertheless, there are three broad concerns about the exclusion restriction: concomitant seasonality in wind speeds and electricity consumption, correlation between wind generation and temperature, and the direct impact of winds on electricity demand through heating demand or interruptions. These are discussed in the following paragraphs.

Seasonality. Both wind speeds and electricity demand feature diurnal and annual seasonality. The root causes for these are different: winds are driven by atmospheric physics, while demand seasonality reflects patterns of human life and business activity (of course, at a fundamental level, they are both caused by the Earth's rotation). Nevertheless, this correlation could bias our estimate of the price elasticity of demand. To correct for this, we control for time of day and time of year in our regressions, using hourly and monthly dummy variables, as well as nonlinear time trends in electricity demand and prices.

Temperature. Beyond seasonality, electricity demand is impacted by weather, particularly ambient temperature, because electricity is used for heating and cooling. Hence, the degree to which wind speeds and temperature are physically correlated may bias our estimate of the price elasticity of demand. To address this issue, we control for the ambient air temperature using heating and cooling degrees or nonparametric specifications.

Wind sensitivity of heating and cooling. We can think of two ways through which wind speeds may directly impact electricity demand (other than through prices): wind-sensitive demand for space heating and cooling and interrupting electricity-consuming services. Wind speeds affect the apparent temperatures and heat losses of buildings, implying that heating and cooling demand may be a function of not only temperature but also wind speed directly (Huang et al., 1990; Sholahudin and Han, 2016). Thus, wind speed could directly affect electricity demand through heating and cooling, and we would expect the sign of the effect to depend on the season: wind would increase heating and electricity demand during winter and decrease cooling and electricity demand during summer. Hence, these wind effects could lead to an overestimation of the price elasticity of demand during winter because demand increases in times of high wind speeds are related not only to low prices but also to the direct positive link between wind and demand. The potential bias would be reversed during summer (and probably smaller because there is less electric cooling than heating in Germany). To control for this causal path, in a sensitivity analysis, we run separate regressions using observations from only one season (see Section 5.2 below). Because our estimates remain unchanged, we conclude that our findings on price elasticity are not substantially biased by direct wind-heating or wind-cooling effects.

Wind-related demand disruption. Storms may lead to interruptions of services, such as the operation of electric railways, and hence lower electricity consumption. This lower demand in times of high wind supply and, hence, low electricity prices would lead us to *underestimate* the price elasticity of demand; therefore, this is not a challenge to our identification strategy. However, to explore this possible effect, we conduct another sensitivity analysis excluding hours with very high wind speeds from our dataset (see Section 5.7 below). As our estimates remain virtually unchanged, we conclude that our exclusion restriction holds.

4 Model specification and estimation approach

Overview. This section presents our five main model specifications (Table 1). All specifications comprise two equations: one for price as a function of supply and one for demand as a function of price. For all models, we use wind as a supply-side instrumental variable to address the simultaneity problem of price and demand. The models differ in their assumptions of the linearity of relationships. The first two specifications assume a linear and a log-linear relationship between demand and price, with linear control variables (Section 4.1). The other specifications assume a linear, a log-linear, and a nonparametric demand–price relationship, while including nonparametric smooth functions for control variables (Section 4.2).

Table 1: Overview of our five main model specifications

Model	1	2	3	4	5
Model class	Parametric models		Nonparametric models		
Specification of controls	Linear		Linear/Nonparametric		
Estimator	Two-stage least squares (2SLS)		Two-stage generalized additive model (2SGAM)		
Demand–price relationship	Linear	Log-linear	Linear	Log-linear	Nonparametric

4.1 Parametric models

Linear specification. For the purely linear specification, we define the simultaneous equation model as

$$Price_t = \alpha_0 + \alpha_1 I_t + \alpha_2 \mathbf{C}_t + \alpha_3 \mathbf{D}_t + v_t \quad (1)$$

$$Demand_t = \beta_0 + \beta_1 Price_t + \beta_2 \mathbf{C}_t + \beta_3 \mathbf{D}_t + u_t \quad (2)$$

where

$Price_t$	Wholesale price of electricity at hour t
$Demand_t$	Electricity demand at hour t
β_1	Price elasticity of demand
I_t	Instrument (wind energy)
\mathbf{C}_t	Controls (solar power, heating and cooling degrees, and EUAs)
\mathbf{D}_t	Dummy controls (hour of the day, weekday, month of the year, and year)
α, β	Modeled linear coefficients
u, v	Error terms

and bold letters indicate vectors of multiple variables or coefficients. Note that some of the control and dummy variables (e.g., hour of the day) are necessary to fulfill the exclusion restriction because both wind energy and electricity demand feature diurnal patterns (see Section 3.2). Other variables are included simply to improve the model fit (e.g., weekdays); they explain demand patterns but are

uncorrelated to wind speeds. The second model specification uses the logarithm of demand in the second stage. The coefficient of price (β_1) can then be interpreted as the *relative* change in demand per absolute change in price. Taking the logarithm of price to estimate the dimensionless elasticity parameter with the typical log-log specification is not possible because hourly electricity prices become zero and even negative.

2SLS estimator. Because of the simultaneity of Eqs. (1) and (2), the key regressor (price) is correlated with the error term of Eq. (2), implying that ordinary least squares estimates are biased. Hence, we use the 2SLS model with Eq. (1) as the first stage and Eq. (2) as the second stage. In the first stage, the price is regressed on our exogenous instrument (wind energy). In the second stage, the price is replaced by the predicted price based on Eq. (1).

2SLS standard errors. Our input data feature strong serial correlation, and so do the error terms v_t and u_t in Eqs. (1) and (2). To account for this, we calculate heteroscedasticity and autocorrelation (HAC) robust standard errors of the estimates. More precisely, we calculate kernel standard errors using the Bartlett kernel with automatic bandwidth selection as implemented in the *linearmodels* Python package. Related to this, Section 5.4 presents a sensitivity regression based on first differences, which significantly reduces serial correlation in the residuals.

4.2 Nonparametric models

Nonparametric specifications. A priori, there is no reason to believe that the relationship between demand and price should be linear. For example, consumers may respond only to very high prices when they decide to ramp down their industrial production processes. To allow for the determination of such possible nonlinear effects of predictors on the response variable in Eqs. (1) and (2), we use a GAM:

$$Price_t = \alpha_0 + s_{\alpha,1}(I_t) + \mathbf{s}_{\alpha,2}(\mathbf{C}_t^s) + \boldsymbol{\alpha}_3 \mathbf{D}_t + v_t \quad (3)$$

$$Demand_t = \beta_0 + \beta_1 Price_t + \mathbf{s}_{\beta,2}(\mathbf{C}_t^s) + \boldsymbol{\beta}_3 \mathbf{D}_t + u_t \quad (4)$$

where

$Price_t$	Wholesale price of electricity at hour t
$Demand_t$	Electricity demand at hour t
β_1	Price elasticity of demand
I_t	Instrument (wind energy)
\mathbf{C}_t^s	Nonparametric controls (solar power, temperature, EUA and time)
\mathbf{D}_t	Dummies (hour of the day, weekday, month of the year, and year)
$s(\cdot)$	Nonparametric smooth functions
α, β	Modeled linear coefficients
u, v	Error terms

The terms $s_1(x_1) \dots s_p(x_p)$ denote smooth, nonparametric functions, which, depending on the underlying patterns in the data, may be nonlinear. The GAM structure allows for the parametric and nonparametric terms to be combined additively, allowing for the usual interpretation of the parametric terms. The expected value of the dependent variable is $g(E(Y)) = \alpha + s_1(x_1) + \dots + s_p(x_p)$, where Y is the dependent variable, $E(Y)$ denotes the expected value of the dependent variable, and $g(Y)$ denotes a link function (for example, a linear or logarithmic function), which links the expected value to the predictor variables x_1, \dots, x_p . In addition, we run two more model specifications with nonparametric dummies that differ in the demand–price relationship: a log-linear

model and one with a nonparametric demand response, obtained by replacing $\beta_1 Price_t$ in Eq. (4) with $s(Price_t)$.

2SGAM. As in the case of the linear specification, the presence of a simultaneous relationship between demand and price can lead to inconsistent estimates of Eq. (4). To address this issue, we use the 2SGAM approach laid out by Marra and Radice (2011). For the endogenous variable, the price, consistent estimates of α , and $s(\cdot)$ are obtained by fitting the GAM corresponding to Eq. (3). Then, we estimate the value of the residuals $\hat{v}_t = Price_t - (\hat{\alpha}_0 + \hat{s}(I_t) + \hat{s}(C_t^s) + \hat{\alpha}_3 D_t)$, which captures the endogenous variation in the price variable. Next, a GAM defined by equation Eq. (4) that includes an additional term $s(\hat{v}_t)$ is calculated to flexibly account for the endogenous variation in the price variable. In doing so, the linear/nonlinear effects of price can be estimated consistently. Note that $s(\hat{v}_t)$ will contain a mixture of effects, which makes it not interpretable. However, this is not problematic, since all that is required is to account for the presence of endogeneity (see Marra and Radice, 2011). In practice, the 2SGAM estimator can be implemented using GAMs represented via any penalized regression spline approach. The 2SGAM approach described in this section is fitted using the *gam* function in the R package *mgcv*. The *mgcv* package fits the 2SGAM estimator using penalized likelihood, which can be maximized by penalized iteratively reweighted least squares (P-IRLS) (Marra and Radice, 2011; Zanin et al., 2015; Wood, 2017).

2SGAM confidence intervals. Confidence intervals for the components of a GAM can be constructed using Bayesian confidence intervals (Wood, 2006; Marra and Wood, 2012; Gu, 2013). Because the second stage of 2SGAM cannot account for the additional source of variability introduced via the residuals calculated in the first step, the intervals for the components in the second-step model will be too narrow, leading to poor coverage probabilities. This can be rectified via posterior simulation (Zanin et al., 2015). Intuitively, samples from the posterior distribution of each first-step model are used to obtain samples from the posterior of the quantities of interest v_t . Then, given that the N_b vector replicates for each v_t , N_d random draws from the N_b posterior distributions of the second-stage model are used to construct approximate intervals for the smooth functions. Marra and Radice (2011) suggested that $N_b = 25$ and $N_d = 100$ yield good coverage probabilities. These settings were used to construct 95% intervals for the function estimates obtained using the 2SGAM.

5 Results

Outline. In this section, we first present the main results on demand elasticity for all five model specifications (Section 5.1). We then conduct a series of sensitivity and robustness analyses: we split the dataset into seasons, day/night, etc., report regional differences, analyze lagged effects, test additional instruments and zero first stage, and exclude extreme weather (Sections 5.2–5.7). Finally, we check the plausibility of the control and dummy variables in our main model specifications (Section 5.8).

5.1 Demand elasticity

First stage. The results of the first-stage regressions are reported in Table 2. The first stage is independent of whether the demand–price relationship is assumed to be linear, log-linear, or nonparametric, but it does differ between the 2SLS and GAM estimation. As expected, we find that wind energy has a significant negative effect on prices in both model specifications (at the 0.001% level). Furthermore, in the linear model, the partial R^2 of wind energy and the F-statistic of the excluded instrument are very high. The F-statistic exceeds the critical values 13.91 and 16.38 for weak instrument tests based on max. 5% 2SLS bias and max. 10% 2SLS size, respectively (Stock and Yogo,

2005). The first-stage partial F statistic calculated using bootstrapped standard errors is calculated as 14,669, much higher than even more stringent criteria (Lal et al., 2021). Therefore, we conclude that wind energy generation is a strong instrument. The results from the nonparametric models are similar to those of the linear specification. The relationship between price and wind generation is largely linear, with a slightly concave shape at higher shares (Figure 6). The relationship between price and other coefficients turns out to be nonlinear, as detailed below.

Table 2: First-stage regression of the endogenous price variable on the instrument wind power and controls

	2SLS	GAM
Adjusted R ²	0.76	0.79
Partial R ² of wind energy	0.46	-
Partial F-statistic	930	-
Wind energy (GW)	-0.94 *** [-1.01, -0.88]	Spline***
Solar energy (GW)	-1.12 *** [-1.19, -1.05]	Spline***
Heating degrees (°C)	0.36 *** [0.22, 0.49]	-
Cooling degrees (°C)	0.45 *** [0.33, 0.58]	-
Temperature (°C)	-	Spline***
EUAs (€/t)	1.08 *** [0.84, 1.32]	Spline***
Coal price (€/MWh)	1.67 *** [1.33, 2.01]	Spline***
Gas price (€/MWh)	0.35 *** [0.12, 0.57]	Spline***
Hour	Dummies	Dummies
Weekday	Dummies	Dummies
Month	Dummies	Dummies
Year	Dummies	Dummies
Time	-	Spline***

95% confidence intervals are reported in brackets; significance levels: *** 0.001 ** 0.01 * 0.05. The significance of the dummy variables can be found in Figure A1 in the Appendix. The partial F-statistic is the F-statistic of the excluded instrument (wind energy).




Figure 6: Estimated nonparametric wind–price relationship from the GAM. The Y-axis shows the change in price at different levels of wind generation while holding all the other variables constant.

Second stage. The results of the second stage are presented in Table 3. For all model specifications, we find the coefficient of price to be statistically significant at the 0.001 level of significance. For a 1 €/MWh increase in the day-ahead electricity price, the electricity demand is estimated to decrease by about 67–80 MW (linear estimates) or 0.12–0.14% (log-linear estimates). The nonparametric specifications yield about 10% smaller estimates than the linear specifications. The confidence intervals reported for the 2SLS model in Table 3 were calculated using HAC standard errors to correct for serial correlation in errors (see Appendix B for treatment of serial correlation). We also use bootstrap methods and calculate the Anderson-Rubin confidence intervals for the IV coefficient, as it is robust to arbitrarily weak instruments (Lal et al., 2021). Both the bootstrapped and the Anderson-Rubin confidence intervals do not diverge much from our 2SLS estimates. The 95% confidence interval obtained using bootstrap methods [-83.98, -75.10] is smaller than the one obtained using HAC standard errors for the linear specification of 2SLS. Further, the Anderson-Rubin confidence [-84.02, -75.14] is also smaller than the confidence interval that we report for the 2SLS model. We conclude that our 2SLS estimates are robust to the choice of estimation method and biases arising from presence of a weak instrument. Table 3 also reports Bayesian confidence intervals for 2SGAM calculated using posterior simulation (see section 4.2)

Table 3: Second-stage regression of demand and log(demand) on price and other covariates

	Linear		Log-linear ^a	
	2SLS	GAM	2SLS	GAM
Adjusted R ²	0.89	0.94	0.90	0.94
Price (€/MWh)	-79.6 *** [-91.3, -67.8]	-67.3 *** [-72.2, -62.7]	-0.14 *** [-0.16, -0.12]	-0.12 *** [-0.13, -0.11]
Solar energy (GW)	-125.3 *** [-153.6, -97.1]	Spline	-0.14 *** [-0.19, -0.10]	Spline
Heating degrees (°C)	310.9 *** [279.8, 342.0]	-	0.55 *** [0.49, 0.61]	-
Cooling degrees (°C)	149.9 *** [113.0, 186.8]	-	0.32 *** [0.25, 0.38]	-
Temperature (°C)	-	Spline ***	-	Spline ***
EUAs (€/t)	98.1 *** [37.5, 158.7]	Spline ***	0.19 *** [0.06, 0.31]	Spline ***
Coal price (€/MWh)	299.8 *** [221.0, 378,7]	Spline ***	0.53 *** [0.37, 0.68]	Spline ***
Gas price (€/MWh)	14.1 [-39.4, 67.7]	Spline ***	0.03 [-0.08, 0.14]	Spline ***
Hour	Dummies	Dummies	Dummies	Dummies
Weekday	Dummies	Dummies	Dummies	Dummies
Month	Dummies	Dummies	Dummies	Dummies
Year	Dummies	Dummies	Dummies	Dummies
Time	-	Spline ***	-	Spline ***

^a All estimated parameters of the log-linear model are reported as percentages.

95% confidence intervals are reported in brackets; significance levels: *** 0.001 ** 0.01 * 0.05.

The significance of the dummy variables can be found in Figure A2 in the Appendix.

Nonparametric demand curve. In the fifth model specification, we allow for a nonparametric relationship between price and demand. Here, we find a statistically significant and generally negative relationship between price and demand (Figure 7). Furthermore, the estimated relationship is also significantly nonlinear; the estimated degrees of freedom (EDF) for the price variable are equal to 8.758. At high and low prices, the relationship appears to be largely linear and negative, although

demand seems largely insensitive to price changes between 20 and 50 €/MWh. However, Figure 7 also reveals a wide variation in the values of demand around the estimated relationship.




Figure 7: Estimated nonlinear demand curve in the GAM

5.2 Temporal differences

Approach. So far, we have presented results based on the entire 5-year dataset. In the following paragraphs, we run the same model specifications on the temporal subsets of the data. These subsets separate the data along four dimensions: into different years, different seasons, weekdays and weekend, and day and night. There are two motivations for doing so. First, we want to explore the robustness of our estimates against changes in the data sample, in part because previous studies used a single year of data. Second, we want to investigate whether the estimated elasticity varies over time and whether consumers respond differently to price changes during the day/night or on weekends. This helps building hypotheses which types of consumers are responding to prices.

Linear estimates. Figure 8 shows the linear price coefficients for the different subsets. Across years, the 2SLS estimates are very stable, with a decrease in demand of 61–82 MW per 1 €/MWh increase in price. The results from the GAM show somewhat higher variation but remain highly significant in all years. We cannot identify a consistent year-to-year trend in the price elasticity of demand. While the seasonal variation in the GAM estimates is difficult to explain, it is reassuring to see that the 2SLS estimates are fairly stable across seasons. This is an important finding because it suggests that heating and cooling are not challenges to our identification (see Section 3.2). Finally, we find that the linear estimates for demand response from both the 2SLS model and GAM are larger during weekdays and daytime hours (8 a.m. to 8 p.m.) than during weekends and nighttime hours. This seems plausible, as it correlates with (industrial) demand and economic activity in general.

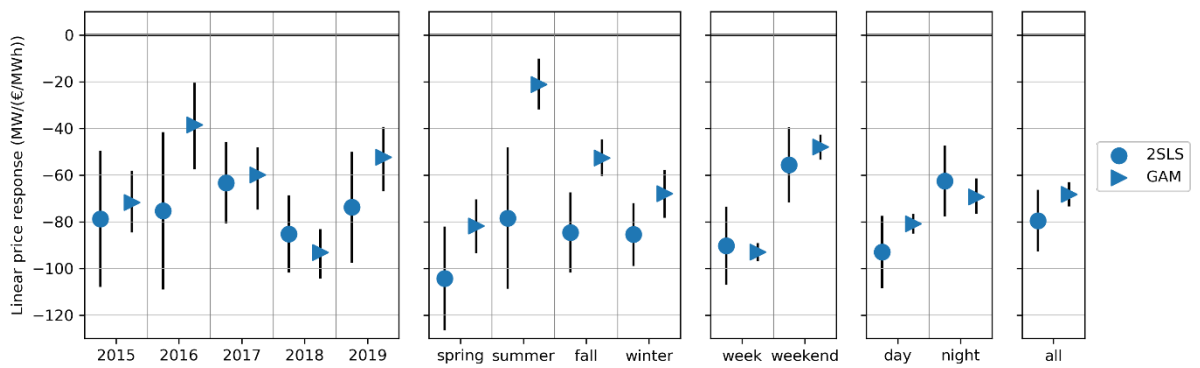


Figure 8: Estimates of linear demand elasticity based on the dataset being split into years, seasons, weekdays, weekends, and daytime and nighttime hours compared to the estimate based on the entire dataset (“all”). The whiskers indicate the 95% confidence interval.

Nonparametric estimates. The log-linear estimates are very similar to the linear specification and highly robust for both estimators (see Figure A3 in the Appendix). The nonparametric estimates show distinct patterns for day and night: during daytime hours, the demand curve is almost linear, but during nighttime hours, it shows a distinct level around the mean price (Figure 9). Similarly, we observe an almost linear curve for weekdays but a nonlinear shape for weekends. The nonlinearity of the demand curve estimated for the total data (Figure 7) seems to be driven by the nonlinearities during nighttime hours. This is confirmed by model diagnostics (see Figures A4-6 in the Appendix).

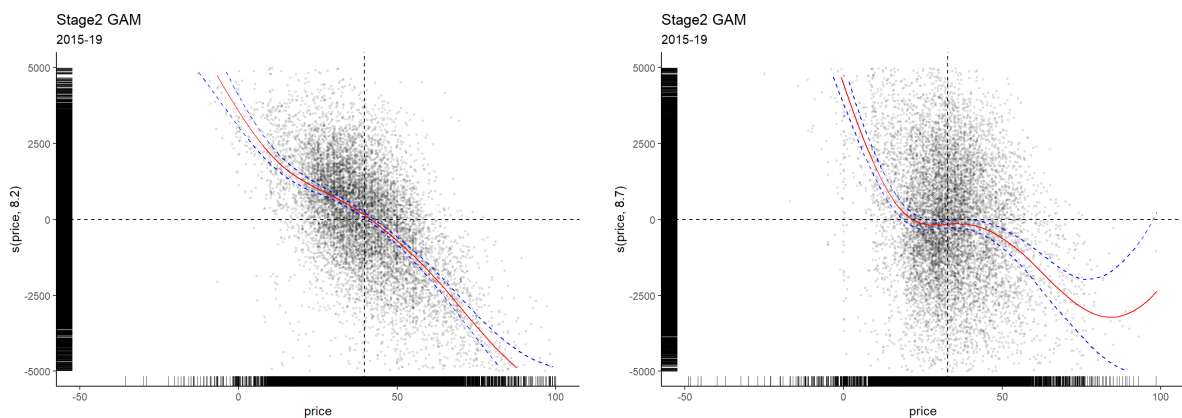


Figure 9: Estimates of the nonlinear demand curve based on the dataset being split into daytime (left) and nighttime hours (right)

5.3 Regional differences

Approach. The previous results are based on the aggregated electricity demand for Germany as a whole. Germany is served by four transmission system operators (TSOs) that report load data for their respective service areas: 50Hertz, Amprion, TennetT, and TransnetBW. We ran the linear and log-linear specification with regional demand as the dependent variable in the second stage to test whether regional elasticities vary (since Germany has one national wholesale electricity price, the first stage is identical across regions).

Results. Figure 10 displays the results for the linear (left) and log-linear (right) 2SLS models. The linear specification reveals how the absolute per-€/MWh demand elasticity is distributed to different

regions. While demand in the areas of 50Hertz and TransnetBW is virtually inelastic, roughly two thirds of the national elasticity occur in the Amprion and one third in the TenneT area. The log-linear model yields similar results: the estimate for relative change in demand is close to zero for 50Hertz and TransnetBW, similar to the national estimate for TenneT, and about twice as large for Amprion. This is revealing because the Amprion region is home to most of Germany’s heavy industry. This finding matches well with our expectation that it is the industrial consumers who are able to respond to variations in electricity prices in the wholesale markets.




Figure 10: Estimates of linear (left) and log-linear (right) demand elasticity obtained using the 2SLS model based on demand data from single TSOs compared to the estimates based on the aggregated German demand (“DE”). The whiskers indicate the 95% confidence interval.

5.4 Lagged effects

Lagged effects. A change in the electricity price at one hour may affect demand not only during that same hour but also in subsequent hours. Moreover, because prices can be reasonably well predicted at short horizons, the (anticipated) electricity price in one hour may also affect demand in the preceding hours. Thinking of electricity as a heterogeneous good across time, one can interpret such a response as intertemporal cross-price elasticity (Hirth et al., 2016; Mier and Weissbart, 2020).

Load shifting vs. inertia. The sign of such intertemporal cross-price elasticity of electricity demand is ambiguous a priori: if load is shifted away from hours with high prices toward preceding or subsequent hours with lower prices, the intertemporal cross-price elasticity may be positive (demand *increases* with higher preceding and subsequent prices). By contrast, if consumption processes are subject to inertia, which means that load is adjusted only if prices are low for several hours in a row, the intertemporal cross-price elasticity may be negative (demand *decreases* with higher preceding and subsequent prices).

Interpretation of our main estimates. Our main estimates of demand elasticity partly reflect such intertemporal cross-price elasticity of demand, in addition to response to prices in any particular hour, because prices are serially correlated. Put differently, our main model estimates the demand response to an increase in the day-ahead price, given that day-ahead prices with temporal proximity will also be higher due to serial correlation.

First differences. As an alternative to our main model specification, this section explores an estimation of the price elasticity of demand based on first differences. The time-related dummy variables remain as level variables. Taking the first differences allows us to exclude the effect of lagged prices and, hence, identify the response of the demand in one hour to a change in price in the same hour.

Results and interpretation. The 2SLS results based on the first-differenced time series are displayed in Figure 11. The magnitudes of both the linear and loglinear estimates of demand response increase substantially. These results suggest that on an average the demand response to change in prices in the same hour is even larger than our main estimates, which partly include intertemporal cross-price elasticities. It also suggests that in our dataset intertemporal cross-price elasticity is positive, i.e., load shifting (positive cross-price elasticity) is larger than load inertia (negative cross-price elasticity).




Figure 11: Estimates of linear (left) and log-linear (right) demand elasticity obtained using the 2SLS model based on first differences compared to the main estimates based on levels. The whiskers indicate the 95% confidence interval.

We also note that the serial correlation in the residuals of the first differences model is lower than in the model with level variables. Additionally, the model fit for the first stage also decreases, although the F-statistic is still above the acceptable threshold (see Appendix B).

5.5 Other instruments

Other instruments. To further assess the robustness of our estimates, we repeated our analysis with different instruments (using the non-differentiated time series), using wind speed and solar generation as instruments.

Wind speed. First, we use wind *speed* as an instrument instead of wind *energy*. We do this to address endogeneity concerns: wind energy, in contrast to wind speed, may be affected by economic or grid curtailment (see Section 3.2). Wind speed, a purely meteorological variable, is clearly exogenous. We calculate wind speed as the population-weighted average wind speed across Germany. Substituting the instrument hardly changes our estimates (see Figure 12). We take this as an indication that the potential endogeneity of wind energy generation is not an issue.

Solar energy. As a second robustness check, we use solar energy as a second instrument (in addition to wind energy). We refrained from using solar energy as an instrument in our main specification because of concerns that solar energy might be inaccurately estimated, violating the exclusion restriction (see Section 2.1). Figure 12 suggests, however, that including solar energy as an additional instrument does not substantially change our estimate. If anything, including solar as a second instrument reduces the size of the estimate, in contrast to our concern that it might inflate it. This suggests that solar energy is also a valid instrument.

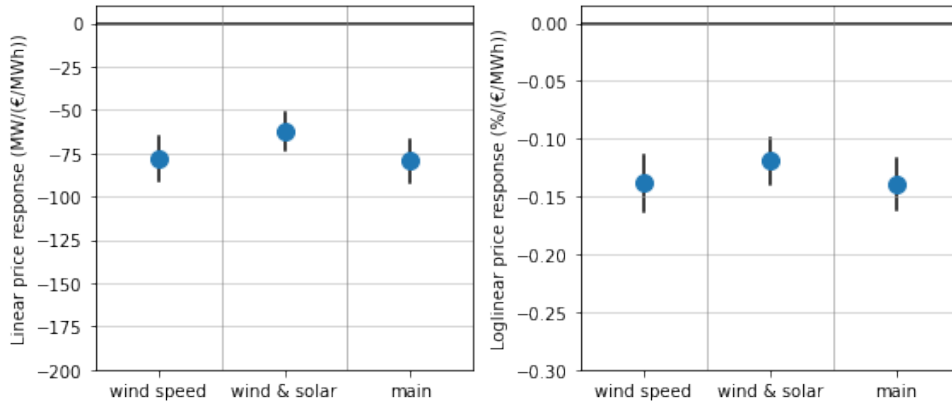


Figure 12: Estimates of linear (left) and log-linear (right) demand elasticity obtained via the 2SLS model using wind speed and wind and solar energy generation as instruments compared to the main estimates obtained using wind energy as an instrument. The whiskers indicate the 95% confidence interval.

5.6 Zero-first-stage test

National power curve. As explained in Section 3.2, we expect the relationship between wind speed and wind energy to be nonlinear because of the S-shaped power curve of wind turbines and the shutdown of turbines at very high wind speeds. This is also evident at the national level (Figure 13).




Figure 13: Wind energy in Germany compared to wind speed. For every observation (hour), the x-axis shows the 90% quantile of all locations sorted by wind speed. This is because wind turbines tend to be located at windy locations. The solid black line indicates the local average wind energy per 0.1 m/s wind speed. The dashed black line indicates the threshold below which observations are considered for the zero-first-stage test.

Test setting. We make use of this nonlinearity to check the validity of our instrument in a zero-first-stage test (van Kippersluis and Rietveld, 2018; Lal et al., 2021). In the first step, we limit our dataset to observations for which the wind speed is below 5 m/s (about 8% of the data). For this subset, we do not expect a change in wind to have a significant effect on wind energy because the power curve is very flat in this range (zero first stage), and we want to reject that a change in wind speed still has a significant effect on electricity demand in the reduced-form equation. Because a direct relationship between wind speed and electricity demand can be suspected in populated areas, we use the population-weighted wind speed for these significant tests. This is the same wind speed variable we used as an instrument in Section 5.5 based on the entire dataset, yielding the same demand elasticity estimate as for wind energy as an instrument.

Test results. Compared to the full dataset, the first-stage effect of mean wind speed on electricity prices in the subsample decreases by a factor of four and becomes statistically insignificant. All other

coefficients remain qualitatively unchanged. In the reduced-form equation, the effect of the mean wind on load also becomes statistically insignificant. This indicates the exclusion restriction holds—that is, wind affects load only through prices.

5.7 Extreme events

Extreme events. Two types of situations may give some reason for concern regarding the exogeneity and exclusion restriction of our instrument: hours with negative prices and hours with extremely high wind speeds. We exclude those observations in turn.

Negative prices. We exclude negative prices because they could lead to the curtailment of wind energy, breaking the exogeneity requirement for wind energy as an instrument. As our estimate does not change substantially (Figure 14), we conclude that wind energy does fulfil the exogeneity requirement.

Storms. The exclusion of extreme wind speeds is due to the concern that storms could interrupt electricity demand from services such as railways, thereby reducing demand through another channel than electricity prices. This would violate the exclusion restriction (see Section 3.2). We exclude observations corresponding to the highest 1% wind speeds based on different aggregations of the spatially resolved wind speed time series across Germany (population-weighted mean, median, 90% quantile, maximum). Our estimates are robust to these changes in the dataset.




Figure 14: Estimates of linear (left) and log-linear (right) demand elasticity obtained using the 2SLS model excluding hours with negative prices and the 1% highest wind speeds from the dataset compared to the main estimates based on the entire dataset. The whiskers indicate the 95% confidence interval.

5.8 Control variables

Overview. This section briefly discusses the estimated linear coefficients of our control and dummy variables in the first and second stages. We also present the estimated nonparametric effects of our main control variables, solar energy and ambient temperature. Overall, the estimated coefficients and nonparametric relationships are plausible, which increases our confidence in the validity of our main results on the price elasticity of demand.

First-stage solar and temperature effects. In the first stage, the effect of solar power is very similar to that of wind energy for both linear and nonparametric specifications (compare Table 2, Figure 6, and Figure 15a). This makes sense, as wind and solar power equally constitute additional supply at low variable costs, shifting the supply curve outward and leading to a decline in prices. The linear

coefficients of the heating and cooling degrees are both positive, which is in line with our expectation that prices increase with heating and cooling due to an increase in demand. The nonparametric specification suggests a nonlinear relationship, with the effect of temperature on price being higher at temperatures above 20°C and below 0°C, which may reflect that both heat pumps and air conditioning become less efficient at more extreme temperatures (Figure 15b).

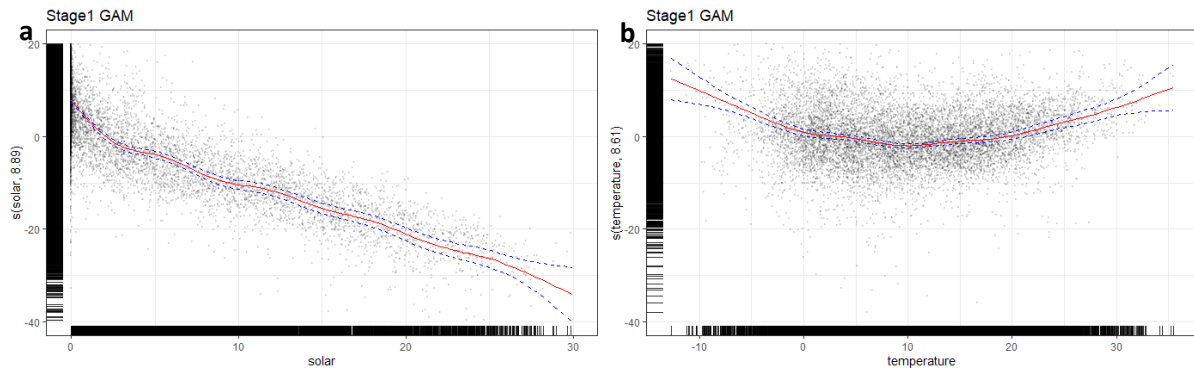


Figure 15: Nonparametric covariates in the first stage: solar (left) and temperature (right)

First-stage fuel and carbon price effects. The coefficients of the EUA in both linear and nonparametric specifications indicate that electricity prices increase by about 1 €/MWh per 1 €/t increase in carbon prices, which seems plausible given that the specific emissions of European power plants are in the range of 0.4–1.2 t/MWh (Ruhnau et al., 2022). The relationship between fuel and power prices depends on the efficiency of different power plants and the share of hours when those plants are marginal, price-setting plants. The efficiency of coal-fired power plants is in the range of 38–46% (Hirth et al., 2021), such that we would expect that a 1 €/MWh increase in coal prices would cause a 2.4 €/MWh increase in electricity prices during those hours when coal-fired power plants are setting the price. The fact that our estimate is lower (about 1.7) seems reasonable because coal-fired power plants were setting the prices only during some hours of our observed period. Our estimated coefficient for the gas price is even smaller (0.35), which may reflect the fact that modern gas-fired power plants are more efficient (up to 60%) and that gas-fired power plants were setting electricity prices during fewer hours of our observed period. The nonparametric specification for coal and gas prices also shows a plausible positive relationship with electricity prices.

Time dummies and trends in the first stage. The estimated dummies for hour, week, month, and year in the first-stage regression are displayed in the Appendix (Figure A1). Both the 2SLS model and the GAM match well with the temporal price patterns described in Section 2. The nonparametric time trend in the GAM is also significant, capturing exceptional price movements, such as the spike at the beginning of 2017 (Figure 3).

Second-stage solar and temperature effects. In the second stage of the 2SLS model, solar has a negative effect on demand. This supports the hypothesis that our demand data may not fully include demand that is served by behind-the-meter solar; therefore, an increase in behind-the-meter production may appear in our demand data as a decrease in demand. The GAM indicates that this effect is most pronounced at low solar levels (Figure 16a). The coefficients of heating and cooling degrees are both positive, as expected, with the coefficient of heating being larger than that of cooling. This matches well with the U-shape of the estimated nonparametric relationship between temperature and demand. The vertex of the curve coincides with our assumed heating and cooling threshold of 15°C (Figure 16b). In the range of 0–10°C, the spline is almost linear, with a slope that is

similar to the 2SLS estimate (310 MW/°C). Below 0°C, the slope increases (but fewer observations are available).

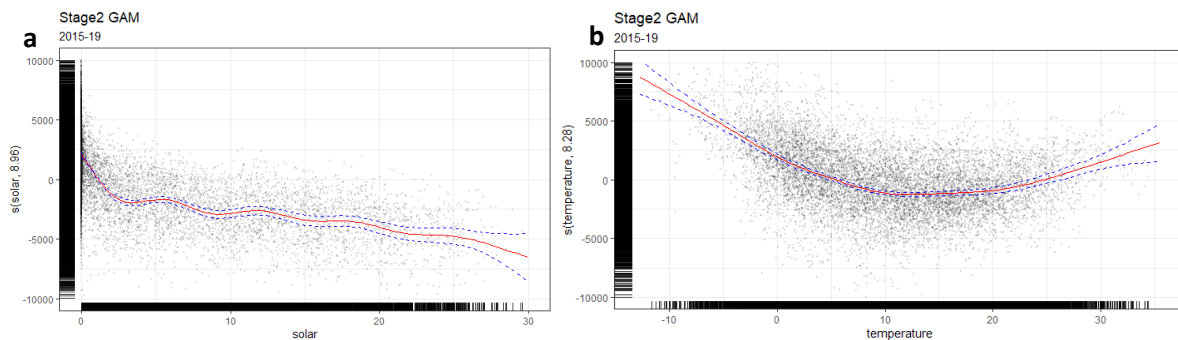


Figure 16: Nonparametric covariates in the second stage: solar (left) and temperature (right)

Time dummies and nonparametric time trends in the second stage. As in the first stage, the second-stage time dummies are displayed in the Appendix (Figure A2) and reflect the time temporal patterns presented in Section 2. The nonparametric time variable features very large confidence intervals, indicating weak time trends on top of the yearly, monthly, daily, and hourly variations.

6 Discussion and conclusion

Summary. We estimate that a 1 €/MWh increase in German wholesale prices triggers a reduction in aggregate demand of 67–80 MW (linear estimates using the 2SLS/2SGAM) or 0.12–0.14% (log-linear estimates). In the following, we put these numbers into perspective, compare them with estimates in the literature, discuss their relevance, and draw some conclusions.

Absolute demand response. While our estimated demand elasticity in terms of €/MWh might seem small at first glance, it should be noted that wholesale electricity prices fluctuate widely over short time scales. To put this into perspective, we derive the absolute price response to the expected range of wind-induced fluctuations in wholesale prices. We focus on wind-induced variation to reflect that our estimation technique only captures the response to price changes that can be explained by wind energy. As an example, we consider the 5–95th percentile of wind energy generation, which corresponds to a range of 27 GW. Using our (first-stage) estimate of a 0.94 €/MWh decrease in wholesale electricity prices per 1 GW increase in wind energy generation, 27 GW of wind variation corresponds to a 26 €/MWh variation in wholesale prices. Using the demand response estimates from our linear model specification, this means that the absolute demand response to wind-induced changes in wholesale prices is about 2 GW. This is 2.6% of peak demand and is quite high, given the widespread assumption that electricity demand is inelastic in the very short term.

Regulatory framework. When interpreting our estimates, two more points are worth keeping in mind. First, only a fraction of customers in Germany are currently exposed to wholesale prices. The percentage change in demand per €/MWh of the exposed consumers must be higher than our estimates, given that the responsiveness of those with flat retail tariffs is zero. Second, our estimates are conditional on the current regulatory situation in Germany. Industrial customers receive a discount on grid fees if they consume a constant amount of electricity from the grid, which likely dampens their response to wholesale prices (Richstein and Hosseinioun, 2020). If more customers were exposed to wholesale prices and regulatory barriers to flexibility were reduced, the price responsiveness of the aggregate demand would likely increase.

Dimensionless elasticity parameter. To make our results comparable to previous literature, we derive a dimensionless elasticity parameter. Recall that directly estimating this parameter from a log-log specification of the demand–price relationship is impossible due to negative prices. At a specific point on the linear demand curve, the dimensionless elasticity can be derived as follows:

$$\varepsilon = \beta_1 \cdot \frac{Price}{Demand} \quad (5)$$

Hence, while the dimensionless elasticity is constant in a log-log specification, it depends on the price and demand levels in the linear model. For average price and demand, the derived elasticities for the 2SLS model and the GAM are -0.051 and -0.045 , respectively.

Literature comparison. This is the parameter most often reported in previous publications. Our estimates are consistent with those of Knaut and Paulus (2016) of -0.02 to -0.13 , who used a similar method. Bönnte et al. (2015) found an elasticity that is an order of magnitude larger (-0.43), which is likely explained by the fact that they looked at ask bids at the power exchange, not aggregate demand. The estimate from Lijesen (2007) is an order of magnitude smaller (-0.0014), which might be the result of a bias introduced by a violation of the exclusion restriction of the instrument used. Kulakov and Ziel (2019) and Damien et al. (2019) also found smaller or even positive estimates, possibly because the simultaneity problem of price and demand was not properly addressed.

Further research. We focused on the demand response to hourly fluctuations in wholesale prices. Further research may expand on this by analyzing the demand response at different time scales. For example, the recent increase in fuel prices may be used as an instrument to analyze price elasticity at a monthly resolution. In addition, future could look at estimating intertemporal price elasticity of demand. We take the first steps towards that using our first differenced approach but a more thorough treatment is beyond the scope of this paper and should be explored in future research.

Appendix A

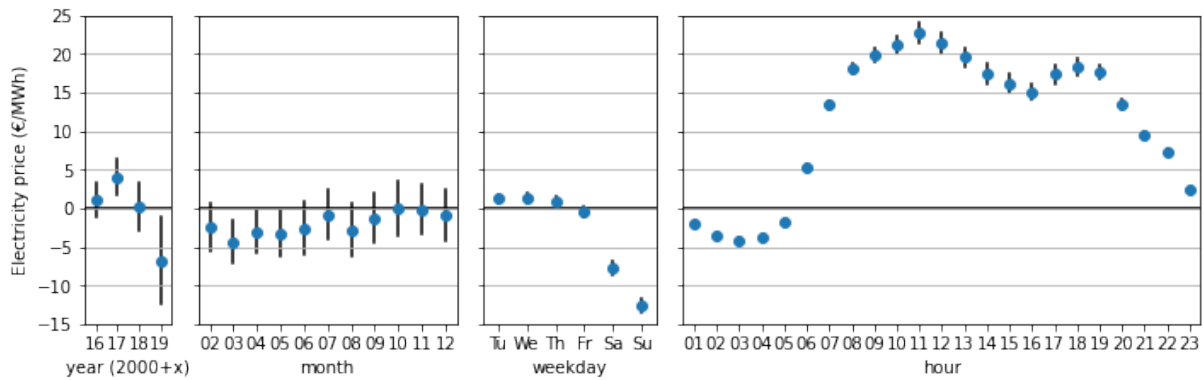


Figure A1: 2SLS time dummies in the first stage (year, month, weekday, and hour)




Figure A2: 2SLS time dummies in the second stage (year, month, weekday, and hour)




Figure A3: Estimates of log-linear demand elasticity based on the dataset being split into years, seasons, weekdays, weekends, and daytime and nighttime hours compared to the estimate based on the entire dataset (“all”). The whiskers indicate the 95% confidence interval.

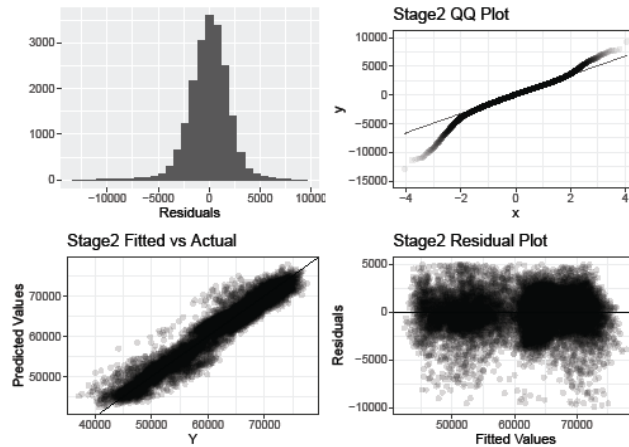


Figure A4: Model diagnostics for the linear model specification in 2SLS for **daytime** hours. The model overall shows a good fit: the residuals exhibit normal distribution, the residuals vs. fitted values plots appear randomly scattered and the predicted vs. actual plots are as expected.

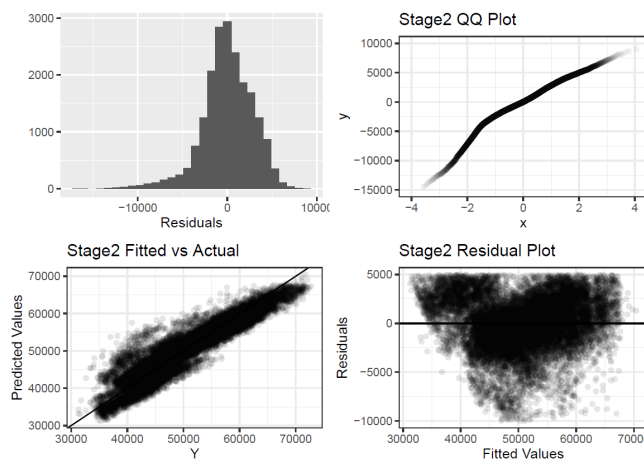


Figure A5: Model diagnostics for the linear model specification in 2SLS for **nighttime** hours. The model fit is considerably weaker compared to daytime hours: the residuals are left-skewed, the residuals vs. fitted values plots appear show distinct pattern and the QQ plots show large deviances.

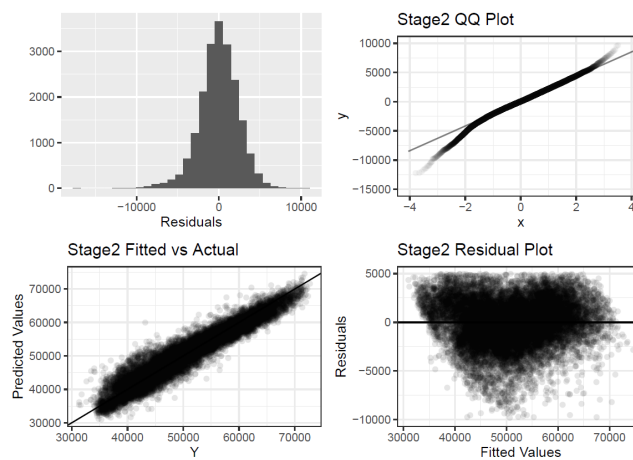


Figure A6: Model diagnostics for the nonparametric model specification in 2SGAM for **nighttime** hours. The model overall shows a better fit: the residuals exhibit normal distribution, the residuals vs. fitted values plots appear randomly scattered, the QQ plot shows lower deviance and the predicted vs. actual plots are as expected.

Appendix B

Serial correlation. The time series data for electricity price, demand, wind energy and our weather-related control variables show serial correlation. This serial correlation can arise on account of both seasonal patterns in the time series data, and correlation between hourly and daily values of these data. Although we control for seasonal dummy variables, the residuals from 2SLS models are also serially correlated. Figure A7 displays the serial correlation in residuals from the first and (linear) second stage of the 2SLS model without any correction for serial correlation. Presence of serial correlation can result in incorrect standard errors. To correct for serial correlation, all confidence intervals reported for the 2SLS estimates in Table 3 are based on HAC standard errors that are robust to serial correlation and heteroscedasticity. The HAC confidence intervals (95% CI = $[-91.3, -67.8]$) are wider than those from the 2SLS model (95% CI = $[-83.6, -75.5]$) but are statistically significant. We also calculated the confidence intervals using the Cochrane–Orcutt (CORC) estimator, which results in a confidence interval that is higher than that obtained from the HAC standard errors (95% CI = $[-125.1, -98.2]$). The estimates reported by us in Table 3 are thus conservative.




Figure A7: ACF plots for the linear first stage specification (left) and the linear second stage specification (right). As expected, serial correlation shows a strong 24-hour seasonality. The residuals of the loglinear second stage behave similarly to the linear second stage.

First differences. Another way to address persistence in the linear specification of the 2SLS models is the first differences approach, which is presented in Subsection 5.4. Figure A8 confirms that the serial correlation in the residuals of the model with first differences is lower than the model with level variables. Note that, however, the adjusted R^2 is also lower: 0.45 instead of 0.76 for the first stage and 0.69 instead of 0.89 for the linear second stage (0.73 instead of 0.90 for the loglinear second stage). The partial R^2 of wind energy in the first stage is also lower (0.03 instead of 0.46), but the corresponding F-statistic has a value of 115, which is still above the thresholds below which wind energy would be considered a weak instrument.




Figure A8: Autocorrelation of the 2SLS residuals with the differenced first stage specification (left) and the differenced linear second stage specification (right). The residuals of the differenced loglinear second stage behave similarly to the differenced linear second stage.

Seasonally differenced estimates. In addition to the first differences approach explained above, we differenced each by different lags: 24th lag (corresponding to the number of observations in a day), 168th lag (corresponding to the number of observations in a week) and 8736th lag (corresponding to the number of observations in 52 weeks of a year). We take the differences of all the time series variables with the 24th lag, leaving the time-related dummy variables as level variables. Differencing with the 24th lag, the 168th lag, and the 8736th lag results in estimates of -76 , -71 , and -65 respectively, which are close to the estimate from the main model specification. The residuals from the 24th lag difference estimate show high serial correlation up to the fourth lag and the overall pattern is similar to that shown for the main specification in Figure A7 except that the cyclical peaks in correlation occur around the 20th lag and not the 24th lag. The residuals from the 168th and the 8736th lag difference estimate show a very high degree of correlation even up to the 50th lag.

These results indicate that, in addition to the seasonal patterns captured by the time dummy variables in the main model specification, there is no persistence in the series arising from seasonality (correlation in time series data across 24 hours, over the week or the year) that significantly affects our estimates. However, the serial correlation arising from persistence between consecutive hours is mitigated by taking the first differences. However, since the estimates from the first difference model are actually higher than those from our main model specification, our main results continue to remain conservative.

References

- Angrist, J.D., Krueger, A.B., 2001. Instrumental Variables and the Search for Identification: From Supply and Demand to Natural Experiments. *Journal of Economic Perspectives* 15, 69–85. <https://doi.org/10.1257/jep.15.4.69>
- Bönte, W., Nielen, S., Valitov, N., Engelmeyer, T., 2015. Price elasticity of demand in the EPEX spot market for electricity—New empirical evidence. *Economics Letters* 135, 5–8. <https://doi.org/10.1016/j.econlet.2015.07.007>
- Choi, W.H., Sen, A., White, A., 2011. Response of industrial customers to hourly pricing in Ontario’s deregulated electricity market. *J Regul Econ* 40, 303–323. <https://doi.org/10.1007/s11149-011-9169-9>
- Cialani, C., Mortazavi, R., 2018. Household and industrial electricity demand in Europe. *Energy Policy* 122, 592–600. <https://doi.org/10.1016/j.enpol.2018.07.060>
- Cramton, P., Ockenfels, A., Stoft, S., 2013. Capacity Market Fundamentals. *EEEEP* 2. <https://doi.org/10.5547/2160-5890.2.2.2>
- Csereklyei, Z., 2020. Price and income elasticities of residential and industrial electricity demand in the European Union. *Energy Policy* 137, 111079. <https://doi.org/10.1016/j.enpol.2019.111079>
- Damien, P., Fuentes-García, R., Mena, R.H., Zarnikau, J., 2019. Impacts of day-ahead versus real-time market prices on wholesale electricity demand in Texas. *Energy Economics* 81, 259–272. <https://doi.org/10.1016/j.eneco.2019.04.008>
- Davis, L., 2021. What Matters for Electrification? Evidence from 70 Years of U.S. Home Heating Choices (No. w28324). National Bureau of Economic Research, Cambridge, MA. <https://doi.org/10.3386/w28324>
- Faruqui, A., Sergici, S., 2013. Arcturus: International Evidence on Dynamic Pricing. *The Electricity Journal* 26, 55–65. <https://doi.org/10.1016/j.tej.2013.07.007>
- Gils, H.C., 2014. Assessment of the theoretical demand response potential in Europe. *Energy* 67, 1–18. <https://doi.org/10.1016/j.energy.2014.02.019>
- Graddy, K., 2006. Markets: The Fulton Fish Market. *Journal of Economic Perspectives* 20, 207–220. <https://doi.org/10.1257/jep.20.2.207>
- Gu, C., 2013. Smoothing spline ANOVA models, 2nd ed. ed, Springer series in statistics. Springer, New York, NY.
- Hirth, L., Mühlenpfordt, J., Bulkeley, M., 2018. The ENTSO-E Transparency Platform – A review of Europe’s most ambitious electricity data platform. *Applied Energy* 225, 1054–1067. <https://doi.org/10.1016/j.apenergy.2018.04.048>
- Hirth, L., Ruhnau, O., Sgarlato, R., 2021. The European Electricity Market Model EMMA. Model description. <https://emma-model.com/>.
- Hirth, L., Ueckerdt, F., Edenhofer, O., 2016. Why Wind Is Not Coal: On the Economics of Electricity Generation. *EJ* 37. <https://doi.org/10.5547/01956574.37.3.lhir>
- Hopper, N., Goldman, C., Bharvirkar, R., Neenan, B., 2006. Customer response to day-ahead market hourly pricing: Choices and performance. *Utilities Policy* 14, 126–134. <https://doi.org/10.1016/j.jup.2005.10.001>
- Huang, Y., Akbari, H., Taha, H., 1990. The Wind-Shielding and Shading Effects of Trees on Residential Heating and Cooling Requirements. *ASHRAE proceedings* 96.
- Huisman, R., Huurman, C., Mahieu, R., 2007. Hourly electricity prices in day-ahead markets. *Energy Economics* 29, 240–248. <https://doi.org/10.1016/j.eneco.2006.08.005>
- Knaut, A., Paulus, S., 2016. When are consumers responding to electricity prices? An hourly pattern of demand elasticity. *EWI Working Paper* 07.
- Knittel, C.R., Roberts, M.R., 2005. An empirical examination of restructured electricity prices. *Energy Economics* 27, 791–817. <https://doi.org/10.1016/j.eneco.2004.11.005>
- Kochems, J., 2020. Lastflexibilisierungspotenziale in Deutschland - Bestandsaufnahme und Entwicklungsprojektionen. 16. Symposium Energieinnovation.
- Kosater, P., Mosler, K., 2006. Can Markov regime-switching models improve power-price forecasts? Evidence from German daily power prices. *Applied Energy* 83, 943–958. <https://doi.org/10.1016/j.apenergy.2005.10.007>
- Kulakov, S., Ziel, F., 2019. Determining Fundamental Supply and Demand Curves in a Wholesale Electricity Market. *arXiv:1903.11383 [econ, q-fin]*.
- Labandeira, X., Labeaga, J.M., López-Otero, X., 2017. A meta-analysis on the price elasticity of energy demand. *Energy Policy* 102, 549–568. <https://doi.org/10.1016/j.enpol.2017.01.002>

- Lal, A., Lockhart, M.W., Xu, Y., Zu, Z., 2021. How Much Should We Trust Instrumental Variable Estimates in Political Science? Practical Advice based on Over 60 Replicated Studies. SSRN Journal. <https://doi.org/10.2139/ssrn.3905329>
- Liebl, D., 2013. Modeling and forecasting electricity spot prices: A functional data perspective. *Ann. Appl. Stat.* 7. <https://doi.org/10.1214/13-AOAS652>
- Lijesen, M.G., 2007. The real-time price elasticity of electricity. *Energy Economics* 29, 249–258. <https://doi.org/10.1016/j.eneco.2006.08.008>
- Marra, G., Radice, R., 2011. A flexible instrumental variable approach. *Statistical Modelling* 11, 581–603. <https://doi.org/10.1177/1471082X1001100607>
- Marra, G., Wood, S.N., 2012. Coverage Properties of Confidence Intervals for Generalized Additive Model Components: Coverage properties of GAM intervals. *Scandinavian Journal of Statistics* 39, 53–74. <https://doi.org/10.1111/j.1467-9469.2011.00760.x>
- Mier, M., Weissbart, C., 2020. Power markets in transition: decarbonization, energy efficiency, and short-term demand response. *Energy Economics* 104644. <https://doi.org/10.1016/j.eneco.2019.104644>
- Olsthoorn, M., Schleich, J., Wohlfarth, K., Klobasa, M., 2019. How much load flexibility can a euro buy? Findings from a contingent valuation experiment with companies in the German commerce and services sector. *Energy Economics* 104603. <https://doi.org/10.1016/j.eneco.2019.104603>
- Patrick, R., Wolak, F., 2001. Estimating the Customer-Level Demand for Electricity Under Real-Time Market Prices (No. w8213). National Bureau of Economic Research, Cambridge, MA. <https://doi.org/10.3386/w8213>
- Richstein, J.C., Hosseinioun, S.S., 2020. Industrial demand response: How network tariffs and regulation (do not) impact flexibility provision in electricity markets and reserves. *Applied Energy* 278, 115431. <https://doi.org/10.1016/j.apenergy.2020.115431>
- Ruhnau, O., Bucksteeg, M., Ritter, D., Schmitz, R., Böttger, D., Koch, M., Pöstges, A., Wiedmann, M., Hirth, L., 2022. Why electricity market models yield different results: Carbon pricing in a model-comparison experiment. *Renewable and Sustainable Energy Reviews* 111701.
- Sauer, A., Abele, E., Buhl, H.U., Fraunhofer-Institut für Produktionstechnik und Automatisierung (Eds.), 2019. Energieflexibilität in der deutschen Industrie: Ergebnisse aus dem Kopernikus-Projekt - Synchronisierte und energieadaptive Produktionstechnik zur flexiblen Ausrichtung von Industrieprozessen auf eine fluktuierende Energieversorgung (SynErgie). Fraunhofer Verlag, Stuttgart.
- Schumacher, M., Hirth, L., 2015. How Much Electricity Do We Consume? A Guide to German and European Electricity Consumption and Generation Data. SSRN Journal. <https://doi.org/10.2139/ssrn.2715986>
- Sholahudin, S., Han, H., 2016. Simplified dynamic neural network model to predict heating load of a building using Taguchi method. *Energy* 115, 1672–1678. <https://doi.org/10.1016/j.energy.2016.03.057>
- Stock, J.H., Yogo, M., 2005. Testing for Weak Instruments in Linear IV Regression, in: Andrews, D.W.K., Stock, J.H. (Eds.), *Identification and Inference for Econometric Models*. Cambridge University Press, pp. 80–108. <https://doi.org/10.1017/CBO9780511614491.006>
- van Kippersluis, H., Rietveld, C.A., 2018. Beyond plausibly exogenous. *The Econometrics Journal* 21, 316–331. <https://doi.org/10.1111/ectj.12113>
- Wiese, F., Schlecht, I., Bunke, W.-D., Gerbaulet, C., Hirth, L., Jahn, M., Kunz, F., Lorenz, C., Mühlenpfordt, J., Reimann, J., Schill, W.-P., 2019. Open Power System Data – Frictionless data for electricity system modelling. *Applied Energy* 236, 401–409. <https://doi.org/10.1016/j.apenergy.2018.11.097>
- Wood, S.N., 2017. *Generalized Additive Models An Introduction with R*, Second Edition, 2nd ed. Chapman and Hall/CRC.
- Wood, S.N., 2006. *Generalized additive models: an introduction with R*, Texts in statistical science. Chapman & Hall/CRC, Boca Raton, FL.
- Wooldridge, J.M., 2013. *Introductory econometrics: a modern approach*, 5. ed., international ed. ed. South-Western Cengage Learning, Mason, Ohio].
- Wright, P.G., 1928. *The Tariff on Animal and Vegetable Oils*. Macmillan, New York.
- Zanin, L., Radice, R., Marra, G., 2015. Modelling the impact of women’s education on fertility in Malawi. *J Popul Econ* 28, 89–111. <https://doi.org/10.1007/s00148-013-0502-8>
- Zarnikau, J., Hallett, I., 2008. Aggregate industrial energy consumer response to wholesale prices in the restructured Texas electricity market. *Energy Economics* 30, 1798–1808. <https://doi.org/10.1016/j.eneco.2007.11.003>