# The Private and External Costs of Germany's Nuclear Phase-Out

Stephen Jarvis, Olivier Deschenes, and Akshaya Jha*

**Abstract**

Many countries have phased out nuclear power in response to concerns about nuclear waste and the risk of nuclear accidents. This paper examines the shutdown of more than half of the nuclear production capacity in Germany after the Fukushima accident in 2011. We use hourly data on power plant operations and a machine learning approach to estimate the impacts of the phase-out policy. We find that reductions in nuclear electricity production were offset primarily by increases in coal-fired production and net electricity imports. Our estimates of the social cost of the phase-out range from 3-8 billion euros per year. The majority of this cost comes from the increased mortality risk associated with exposure to the local air pollution emitted when burning fossil fuels. Policymakers would have to significantly over-estimate the risk or cost of a nuclear accident to conclude that the benefits of the phase-out exceed its social costs. We discuss the likely role of behavioural biases in this setting, and highlight the importance of ensuring that policymakers and the public are informed about the health effects of local air pollution.

**JEL Codes:** Q41, Q53, C53

**Keywords:** Nuclear Power, Air Pollution, Electricity, Health

---

*Jarvis: Department of Geography and Environment, London School of Economics, Houghton Street, WC2A 2AE London. Email: S.Jarvis@lse.ac.uk. Deschenes: Department of Economics, University of California, 2127 North Hall, Santa Barbara, CA 93106, IZA, and NBER. Email: olivier@econ.ucsb.edu. Jha: H. John Heinz III College, Carnegie Mellon University. 4800 Forbes Ave, Pittsburgh, PA 15213. Email: akshayaj@andrew.cmu.edu. The authors thank seminar participants at the University of Wyoming, the University of Texas-Austin, the University of Regensburg, the Mercator Research Institute on Global Commons and Climate Change, the Venice EAERE Summer School, The Workshop in Environmental Economics and Data Science (TWEEDS) 2020, and the 2020 ASSA Conference in San Diego. The authors also wish to acknowledge the Library at the University of California, Berkeley, and the German Research Foundation (DFG) through CRC TR 224 (Project B7) for providing support for the completion of this research.

# 1 Introduction

The Intergovernmental Panel on Climate Change (IPCC, 2018) and the International Energy Agency (IEA, 2019) both envisage nuclear power continuing to play an important role in mitigating climate change over the coming decades. This is because nuclear electricity generation produces minimal carbon emissions under normal operating conditions (Markandya and Wilkinson, 2007). In contrast, burning fossil fuels to produce electricity is known to emit both global pollutants that contribute to climate change and local air pollutants that have negative consequences on human health (NRC and NAS, 2010; Jaramillo and Muller, 2016; Deschenes, Greenstone and Shapiro, 2017; Holland et al., 2018).

Despite this, many countries have explicit policies in place to phase-out nuclear power entirely, including Germany, Belgium, Spain, Switzerland, and South Korea. These phase-out policies were implemented primarily in response to concerns about long-term solutions for storing nuclear waste and public fears of catastrophic nuclear accidents. These fears intensified considerably following the incidents at Three Mile Island in 1979, Chernobyl in 1986, and Fukushima in 2011.

The decision to phase-out nuclear production in many countries seems to suggest that the expected costs of nuclear power exceed its benefits. Yet, it has proven difficult to empirically quantify the full range of economic and environmental impacts from large-scale shifts away from nuclear power. This paper aims to fill this important gap by documenting how the phase-out of nuclear power in Germany affected market and environmental outcomes.

The context of Germany's nuclear phase-out affords us several advantages over previous research studying closures of nuclear power plants. First, and most importantly, Germany shut down over 12 GW of nuclear production capacity during this period, with two thirds of this occurring over a few months in 2011. This is far larger than the reductions in capacity studied by previous research that focused on the shutdown of a small number of nuclear plants in the United States (Davis and Hausman, 2016; Severnini, 2017; Adler, Jha and Severnini, 2020). Second, Germany plans to shut down all of its remaining nuclear reactors by 2022. Our study thus provides timely and policy-relevant information on the consequences of Germany's nuclear phase-out moving forward. Third, studying electricity markets in the European context gives us the opportunity to examine how cross-border trade was impacted by a large shock to production in one country. Finally, Germany's nuclear phase-out was the direct result of political actions taken following extensive anti-nuclear campaigning in Germany as well as a sudden increase in the perceived risk of nuclear power following the Fukushima accident in Japan (Goebel et al., 2015). Importantly, the phase-out was not caused by changes in the economic or environmental conditions pertaining to nuclear production in Germany.

This paper adds to the relatively small literature that explores the effects of the nuclear phase-out on the German electricity sector. For instance, both Traber and Kemfert (2012) and Knopf et al. (2014) used mixed economic-engineering models of the power sector to forecast changes to capacity investments, electricity prices, and carbon emissions. More recently, Grossi, Heim and Waterson (2017) uses an event study framework to econometrically estimate the impact of the initial nuclear plant closures in 2011 on electricity prices over a three year window between 2009 and 2012. Another paper by Grossi et al. (2018) focuses on the impacts of the phase-out on electricity prices in neighboring countries. The broad consensus across this small existing literature is that nuclear power was replaced primarily by fossil-fuel-fired production, resulting in higher electricity prices and more carbon emissions.

We expand on this existing literature by estimating the spatially disaggregated impacts of the phase-out on production costs, net electricity imports, and local air pollution. Estimating plant-level changes in output due to the phase-out is especially important because the majority of the social cost of this policy stems from the increases in local air pollution caused by replacing some of the lost output from nuclear sources with coal-fired electricity production. We also do not consider just the initial nuclear reactor closures in 2011. Since our analysis runs to the end of 2019, we incorporate the subsequent incremental shutdowns of nuclear power plants over this period.

To proceed, we use a machine learning framework to estimate each plant's counterfactual output over time under a "no phase-out" scenario. This counterfactual allows us to identify which power plants increased their output in response to the nuclear plant closures. Our approach draws upon a growing literature exploring ways to use machine learning methods to construct counterfactuals (Varian, 2014, 2016; Athey et al., 2017; Carvalho, Masini and Medeiros, 2018). A smaller number of papers have applied these techniques to evaluate policies in the energy sector. For instance, Burlig et al. (2020) study an energy efficiency program for schools in California using regularized regression methods to estimate counterfactual electricity consumption patterns. Both Souza (2020) and O'Neill and Weeks (2019) also examine programs aimed at encouraging energy savings, leveraging machine learning methods to construct counterfactuals that can identify important heterogeneity in treatment effects. We focus on a policy that affects electricity producers by altering the dispatch of power plants. Similar in spirit, Abrell, Kosch and Rausch (2019) uses a machine learning framework to predict the counterfactual dispatch of power plants as part of their analysis of the UK's carbon tax.

To conduct our analysis, we combine hourly data on observed power plant operations between 2010-2019 with a wide range of related information, including electricity demand, local weather conditions, wholesale electricity prices, fuel prices and various plant characteristics. We use these data to train a machine learning algorithm that can predict plant operations based on market conditions. This allows us to predict the quantity of electricity produced by each power plant in Germany in each hour-of-sample under two scenarios: one with the nuclear phase-out and one without it. We interpret the difference in market and environmental outcomes between these two scenarios as the impact of the phase-out policy. The results of this estimation procedure indicate that the lost nuclear electricity production due to the phase-out was replaced primarily by coal-fired production and net electricity imports. This phase-out induced increase in coal-fired production remains sizable and persists through 2019 across a variety of specifications, including a wide range of assumptions regarding the level of investment in renewables caused by the phase-out.

We use the predicted changes in plant-level electricity production due to the nuclear shutdowns to calculate the private and external costs of the phase-out. Our estimates of the social cost of the phase-out to German firms and residents range from 3-8 billion euros per year. The majority of this cost is due to the increased mortality risk from local air pollution exposure as a consequence of the increased use of fossil fuels for electricity production. We also find that the private costs of electricity production and wholesale electricity prices in Germany increased as a consequence of the phase-out, confirming and extending the results from the previous literature over a longer post phase-out period.[1]

[1] Neidell, Uchida and Veronesi (2019) and He and Tanaka (2019) document increases in mortality

Taken together, our analysis suggests that German citizens faced increased electricity prices and greater exposure to local air pollution as a result of the phase-out of nuclear power. The phase-out also made meeting national targets to reduce carbon emissions more challenging. Despite this, the phase-out continues to receive broad support, with more than 81% of German residents in favor of the policy in a 2015 survey (Goebel et al., 2015). Concerns about the risks of nuclear accidents and storing nuclear waste are central to anti-nuclear sentiment, and motivated the decision to phase-out nuclear power (Ethics Commission, 2011).

The risks of nuclear power are difficult to quantify. Nevertheless, even the largest estimates of the expected benefits from reducing nuclear risks are much smaller than our estimates of the social costs of the phase-out (D'haeseleer, 2013). For the expected benefits of the phase-out to be equal to the lower bound of our estimated social costs, policymakers would have to either exhibit a very high level of risk aversion or view nuclear accidents as being much more likely than the available evidence suggests (Wheatley, Sovacool and Sornette, 2017). Consistent with this, previous research has shown that people tend to greatly overestimate both the probability of a nuclear accident and the expected damages from such an event (Slovic, Fischhoff and Lichtenstein, 1979; Slovic and Weber, 2002; Slovic, 2010). The decision to phase-out nuclear power therefore suggests a preference for reducing exposure to low probability, highly uncertain, and potentially catastrophic nuclear accidents, even if this leads to relatively moderate damages from exposure to pollution emitted due to the prolonged reliance on fossil-fuel-fired electricity production. Such behavior is consistent with many components of prospect theory, in particular loss aversion and relying on "probability weighting" rather than objective probabilities (Kahneman and Tversky,

stemming from policy-induced increases in retail prices and reductions in electricity consumption following the phase-out of nuclear power in Japan.

1979; Barberis, 2013).

Nuclear accidents were particularly salient in the immediate aftermath of the Fukushima crisis (Ethics Commission, 2011; Goebel et al., 2015; Tanaka and Zabel, 2018). By comparison, an incremental increase in mortality risk due to local air pollution exposure is far less salient. Research on the role of framing effects has documented a behavioral bias towards placing increased weight on the more salient aspects of a decision (Kahneman, 2003). Consistent with this, policymakers appear to have made no mention of the impact of the phase-out on local air pollution when making their decision (BMWI, 2010; Ethics Commission, 2011). Further, subsequent studies of the impact of the phase-out have also focused exclusively on electricity prices and carbon emissions (Knopf et al., 2011; Traber and Kemfert, 2012; Knopf et al., 2014; Grossi, Heim and Waterson, 2017; Grossi et al., 2018).[2] The omission of local air pollution is especially troubling because the vast majority of the social costs of the phase-out originate from increases in air pollution due to increased output from fossil-fuel plants. This highlights the importance of ensuring that the public are informed about the health costs of local air pollution, and that policymakers incorporate these health costs into their decisions.

## 2 Background on Nuclear Power in Germany

The first nuclear power stations were built in Germany in the 1960s. Germany's nuclear production capacity expanded rapidly over the next three decades; the last nuclear reactor was commissioned in 1989. Despite no new reactors coming online

[2]We note that the absence of academic studies and discussions by policymakers of the potential impact of the phase-out on local air pollution does not by itself imply that the general public is unaware of this issue.

in the 1990s and 2000s, roughly 25% of Germany's electricity production came from nuclear plants prior to 2011.

Nuclear power has long been controversial in Germany. There were protests as far back as the 1970s at a number of sites where nuclear facilities were either proposed or under construction. However, the Chernobyl disaster in Ukraine in 1986 created a focal point in the politics of nuclear power in Germany. Specifically, radioactive fallout affected much of the country and led to growing public concern. In 1998, the Schröder government took power through a coalition between the Social Democratic Party (SPD) and the Green Party. Over the next two years, the Schröder government banned the construction of new reactors and negotiated a policy of phasing out nuclear power completely. This plan called for all nuclear reactors to be shut down by 2022.

The center-right Merkel government came to power in 2009. This government renegotiated the original phase-out policy by committing to extend the lifetimes of the existing reactors. This revised policy pushed back the shutdown of the last nuclear reactor into the 2030s. The extensions would be 8 years for the eight older reactors built up to and including 1980, and 14 years for the nine newer reactors built after 1980. However, the specter of nuclear disaster rose again due to the Fukushima incident on March 11, 2011. In response, public opposition to nuclear power intensified again, with an estimated 250,000 people taking to the streets nationwide to protest in the days and weeks following March 11, 2011. The resulting political pressure forced the Merkel government to declare a moratorium on planned extensions at existing nuclear power plants almost immediately after the Fukushima incident. In addition, eight older reactors were taken offline for testing.

By May of 2011, German policymakers decided to return to a version of the original plan: phase out all nuclear power by 2022. Specifically, of the seventeen reactors operating in 2011, the eight reactors already temporarily offline were closed immediately (8.4 GW of capacity), a ninth reactor was closed in 2015 (1.3 GW), a tenth in 2017 (1.3 GW), an eleventh in 2019 (1.4 GW), and the final six reactors (8.1 GW) will close in 2022. Our sample period ends in 2019. Consequently, our empirical analysis focuses on the closure of the nuclear reactors in 2011, 2015, 2017 and 2019, but does not forecast the impact of the planned closures in 2022.

The phase-out of nuclear power is part of a wide-ranging transformation of Germany's energy sector known as the *Energiewende*. The primary goal of this policy is to reduce Germany's carbon emissions by at least 80% by 2050 relative to 1990 levels (BMWi, 2018). To achieve this, Germany has undertaken major investments in renewable electricity production, transmission grid infrastructure, and energy efficiency measures. The sweeping scope of the *Energiewende* policy highlights the importance of accounting for a host of potential time-varying confounders when assessing the impact of the nuclear phase-out. This motivates the development of our machine learning approach.

# 3 Data Description and Summary Statistics

This paper brings together a wide range of publicly available data on the German power sector. First, we obtain data on the hourly operation of the electricity grid in Germany from the European Network of Transmission System Operators for Electricity (ENTSOE). This includes hourly data on total electricity demand, aggregate electricity production by source type, and imports and exports in and out of Germany at border points. ENTSOE also provide data on unit-level electricity production for all power plants with production capacity greater than 100MW.

However, the ENTSOE data are only available from 2015-2019. In order to study the nuclear phase-out beginning in March 2011, we supplement the ENTSOE data with data on hourly total production by source (e.g., nuclear, coal, natural gas, oil, etc.) from the European Energy Exchange (EEX) for 2010-2019. A key advantage of our machine learning approach is that it allows us to combine hourly source-level data on electricity production from 2010-2019 with data on hourly plant-level output from 2015-2019 in order to construct a prediction of each plant's output in each hour-of-sample for the entire 2010-2019 sample period.

We also integrate data from Germany's four different transmission system operators (TSOs) that are each responsible for a different geographical area on the German grid: Amprion, TenneT, TransnetBW and 50Hertz. Each TSO reports hourly production from wind and solar sources for the period 2010-2019. The TSOs also provide data at the hourly level on electricity imports and exports in and out of Germany at border points, as well as the hourly total quantity of electricity demanded for their portion of the grid.

We utilize data on hourly wholesale electricity prices for both Germany and neighboring countries. These data are collected by ENTSOE and are accessible through Thomson Datastream.

We construct each fossil-fuel-fired plant's marginal cost in each day using data on input fuel prices and carbon emission prices gathered from the following two sources. First, Thomson Datastream provides data on daily natural gas prices in Germany. The Intercontinental Exchange (ICE) lists monthly coal and oil prices as well as the monthly permit prices for carbon dioxide emissions set by the European

Union Emissions Trading System (EU-ETS).[3] Assumptions on other components of variable costs and fixed costs for all sources (including nuclear, wind and solar) are taken from a range of industry reports and are discussed in Section 6.2.

Our analysis of the environmental impacts of the nuclear phase-out also combines data from multiple sources. The European Environment Agency (EEA) reports annual carbon dioxide emissions for each power plant that participates in the EU-ETS. The EEA also reports annual plant-level data on fuel inputs and local air pollution emissions.[4] Daily station-level weather data comes from Germany's national meteorological service (DWD) and daily ambient air pollution monitor data are from the German Environment Agency (UBA). Finally, we compile other electricity sector data and power plant level characteristics from a variety of different sources (Open Power System Data, 2018; BNetzA, 2020; Egerer, 2016*a*).

Table 1 provides summary statistics for the aggregate electricity sector and by type of plant in 2010 (the first year in our sample) and 2019 (the last year in our sample). The top panel documents the drastic shifts in the electricity production mix over this time period. Nuclear production almost halved while production from renewable resources more than doubled. Production from coal plants declined while production from natural gas plants grew, due primarily to relative changes in coal prices versus natural gas prices as well as increases in carbon permit prices. Since total electricity demand has remained largely flat over the 2010s, the additional production from renewables has also served to expand Germany's position as a net exporter of power to neighboring countries. Online Appendix Figure A.1 presents a more detailed breakdown of the quantity of electricity produced by different types

[3]Specifically, we use the futures price corresponding to the nearest term monthly contract (i.e., the front month futures contract).

[4]These data are collected as part of monitoring for the EU Large Combustion Plant Directive.

Table 1: Summary Statistics

| | **2010** | **2019** |
|---|---|---|
| **Production and Net Imports Annual Totals (TWh)** | | |
| Nuclear | 134.7 | 69.5 |
| Hard Coal | 93.9 | 53.4 |
| Lignite | 130.9 | 104.2 |
| Gas | 53.6 | 75.5 |
| Oil | 1.9 | 3.1 |
| Hydro, Solar and Wind | 76.4 | 193.5 |
| Net Electricity Imports | -2.4 | -33.8 |
| **Total Number of Plants** | | |
| Nuclear | 15 | 7 |
| Hard Coal | 67 | 56 |
| Lignite | 31 | 34 |
| Gas | 176 | 209 |
| Oil | 33 | 30 |
| **Total Capacity (GW)** | | |
| Nuclear | 19.2 | 9.5 |
| Hard Coal | 25.8 | 23.7 |
| Lignite | 20.2 | 21.0 |
| Gas | 22.8 | 26.2 |
| Oil | 3.8 | 3.7 |
| **Marginal Cost (2019 € per MWh)** | | |
| Nuclear | 11.2 | 10.0 |
| Hard Coal | 46.2 | 43.6 |
| Lignite | 38.8 | 44.2 |
| Gas | 57.2 | 46.9 |
| Oil | 145.8 | 141.5 |
| **Wholesale Electricity Price** (2019 €/MWh) | 53.21 | 37.21 |

**Notes:** This table reports summary statistics for Germany's electricity generation sector in 2010 and 2019. Annual electricity production data is taken from BNetzA (2020). If data for 2019 are not yet available, we extrapolate values based on more up-to-date data from Fraunhofer ISE. The power plants in our sample are those listed in BNetzA's Power Plant List, compiled by Open Power System Data (2018). Electricity prices are from ENTSOE and Thomson Reuters. Marginal costs are calculated by the authors based on various sources, full details of which can be found in Section 6.2. All monetary values are in constant 2019 €.

of sources in Germany over 2010-2019.

Table 1 also documents that wholesale electricity prices fell by 30% from 2010 to 2019. The main driver of this decline in wholesale prices was the large increase in zero marginal cost production from wind and solar resources. A secondary driver was the fall in the price of natural gas. Natural gas plants are often the marginal source of generation in Germany, and the average marginal cost of gas-fired production fell by 17% between 2010 and 2019.

While wholesale prices decreased from 2010 to 2019, the revenues earned by plants outside of the wholesale market have increased significantly. Renewables in particular receive guaranteed payments well in excess of the wholesale price. For example, wholesale prices were roughly €37 per MWh in 2019 yet renewable resources received additional subsidy payments averaging €130/MWh in this year.[5] Fossil plants also earn sizeable non-market revenues, in large part due to the subsidies paid to combined heat and power plants that produce both electricity and useful heat. Due to this dramatic increase in payments made outside of the wholesale market, the prices faced by end-use consumers of electricity in Germany have actually risen from 2010 to 2019 despite the reduction in wholesale electricity prices (BNetzA, 2020).

The second and third panels of Table 1 report the total number of plants and the total amount of production capacity at the beginning of 2010 versus the beginning of 2019 for each of the major conventional sources of power in Germany: nuclear, hard coal, lignite, natural gas, and oil. The shutdown of nuclear plants is evident from the decline in both the number of plants and the amount of nuclear capacity.[6] The

[5]These subsidies range from €60/MWh for onshore wind to as much as €270/MWh for solar panels (BNetzA, 2020).

[6]One nuclear reactor closed at the end of 2019. We also exclude the two reactors at the Kruemmel power plant from our analysis because this plant was already in long-term shutdown since 2009.

capacity of fossil fuel plants has remained largely flat. As documented in the fourth panel of Table 1, outside of renewables, nuclear plants have the lowest marginal costs, followed by lignite, hard coal, natural gas and lastly the small number of peaking oil plants.

# 4 Event Study Framework and Results

In response to the Fukishima nuclear accident in Japan, the German government suddenly and unexpectedly shut down eight nuclear reactors on March $15^{th}$ 2011. We can thus analyze the impact of these closures on market outcomes using the event study framework formulated in Davis and Hausman (2016) and more recently implemented by Grossi, Heim and Waterson (2017). Specifically, we apply this event study framework to estimate how total electricity production from each fuel type $i$ in each hour-of-sample $t$ responds to changes in electricity demand before versus after March 15th, 2011.

The independent variables of interest are equally-spaced bins of net electricity demand interacted with an indicator for observations after March $15^{th}$ 2011. For the purpose of this event study, "net electricity demand" is defined as electricity demand net of production from renewable sources. We consider net demand because production from renewable sources has near-zero marginal costs and is "non-dispatchable": wind and solar resources are only able to produce when the wind is blowing or the sun is out.[7] In order to implement the event study approach, we restrict the sample to observations less than 12 months before or after March $15^{th}$ 2011 and estimate

[7]In making this assumption, we follow Davis and Hausman (2016).

the following regression:

$$G_{i,t} = \sum_b [\mathbf{1}\{L_t \in B_b\}(\alpha_{i,b} + \beta_{i,b} \cdot \mathbf{1}\{t \geq 3/15/2011\})] + \gamma_m + \epsilon_{i,t} \tag{1}$$

where $G_{i,t}$ is the total quantity of electricity produced by fuel type $i$ in hour-of-sample $t$. $L_t$ is net demand in hour $t$, and $\mathbf{1}\{L_t \in B_b\}$ is an indicator that takes on the value one if $L_t$ is in bin $B_b$ and is zero otherwise. Next, the indicator $\mathbf{1}\{t \geq 3/15/2011\}$ takes on the value one if the observation corresponds to an hour-of-sample on or after March $15^{th}$ 2011 and is zero otherwise. Finally, we include month-of-year fixed effects (i.e.: $\gamma_m$) and cluster standard errors by week-of-sample.

Figure 1 plots the coefficient estimates of interest (i.e.: $\hat{\beta}_{i,b}$) along with their 95% confidence intervals. Panel (a) of this figure shows that hourly electricity production from nuclear sources dropped by roughly 5 GWh on average across all levels of net demand. Panels (b)-(d) demonstrate that this lost nuclear production was offset in large part by increases in electricity production from fossil fuel sources. Specifically, production from lignite increased by roughly 1 GWh on average at low levels of net demand. Production from hard coal increased by 2-3 GWh on average across all levels of net demand. Finally, gas-fired electricity generation also increased by roughly 2 GWh on average, and by as much as 6 GWh for hours-of-sample with very high net demand.

There are several limitations to this event study approach. First, hourly plant-level data on electricity production are not available prior to 2015. Consequently, the event study framework cannot be used to explore heterogeneity in how different plants respond to the nuclear phase-out beginning in 2011. This heterogeneity is especially important because the amount of local air pollution emitted per MWh

Figure 1: Event Study Estimates of the Effect of the 2011 Nuclear Plant Closures on Nuclear and Fossil-Fuel Electricity Production

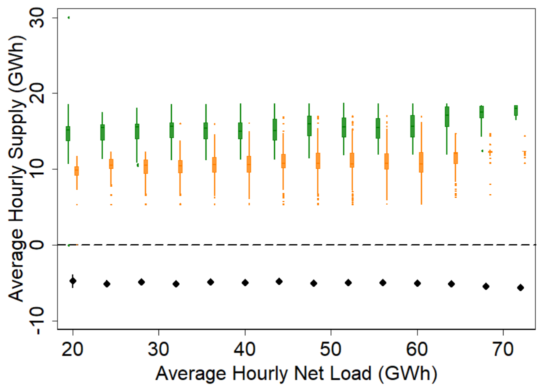


(a) Nuclear

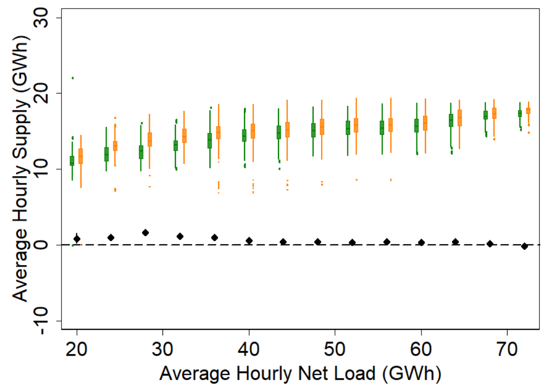


(b) Lignite

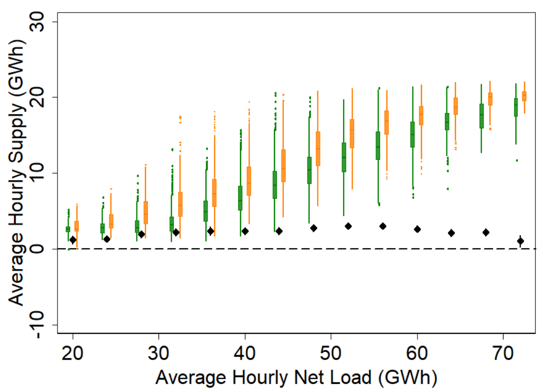


(c) Hard Coal

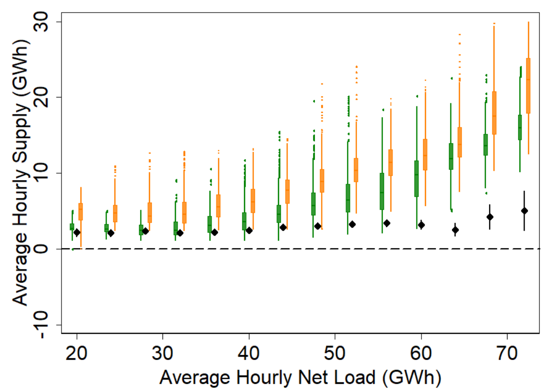


(d) Natural Gas

This figure plots the results from an event study analysis of the effects of the nuclear phase-out in Germany in 2011. Values are plotted to show how the impact varies between source types and across each bin of net demand (i.e., electricity demand minus production from renewables). The pre period spans the year prior to March 15, 2011, and the post period spans the year after. The box plots in green (shifted left) correspond to the pre period electricity production. The box plots in orange (shifted right) correspond to the post period electricity production. The points in black correspond to the estimated coefficients of the difference, with 95% confidence intervals constructed based on standard errors that were clustered by week-of-sample. Panel (a) presents the estimates for nuclear production, while Panels (b)-(d) present the corresponding estimates for production from lignite, hard coal, and natural gas, respectively.

of production can vary significantly across plants burning the same type of fuel. In addition, the monetary damage from local air pollution emissions is also tied directly to the number of people exposed to this pollution; the same level of pollution emissions from two different plants can have very different damages based on the number of people living near each of these plants.

Second, the event study framework relies on the assumption that changes in power plant operations around March 15, 2011 are caused by the nuclear reactor closures rather than changes in other factors that determine production behavior. To ensure that this assumption holds, we examine the impact of the phase-out in a fairly narrow window around the initial 2011 shutdowns. Focusing on this narrow window allows us to argue that electricity suppliers could only respond to the nuclear shutdowns in the very short-run by adjusting output. However, subsequent nuclear plant shutdowns occurred incrementally and were pre-announced. As such, firms may have been able to take actions in anticipation of these later closures.

Finally, other important economic factors changed over our 2010-2019 sample period. For example, coal and natural gas plants had similar marginal costs in 2011. However, coal prices decreased precipitously from 2011-2015 while natural gas prices increased over this period. Coal plants were thus increasingly more likely to produce in place of natural gas plants from 2011-2015 even absent any changes in nuclear power production. In addition, many older coal and gas plants were retired between 2010 and 2019, and a number of new fossil-fuel-fired plants came online during this period as well. Summarizing, it is unlikely that market outcomes before versus after March 2011 were driven solely by the phase-out, especially as we look further in time after the 2011 shutdown decision.

Combined, these concerns motivate the development of a machine learning approach that can estimate the spatially disaggregated impacts of the nuclear phase-out over a longer time horizon controlling flexibly for changes over time in other important economic factors.

# 5 Machine Learning Methodology and Validation

We develop a machine learning approach to more credibly estimate the economic and environmental impacts of the series of nuclear plant closures that occurred between 2011 and 2019. The first subsection discusses the construction of the dependent and independent variables while the second subsection provides details on the random forest algorithm used to estimate the relationship between these variables. We next describe how we use our estimates to calculate hourly output for each plant in each of two scenarios: with the nuclear phase-out versus without the nuclear phase-out. The final two subsections discuss robustness to changes in other economic factors due to the phase-out and performance of the model respectively.

## 5.1 Constructing the Training Dataset

We train our machine learning algorithm to predict power plant operations using a data set of roughly 6.5 million observations. The goal is obtain estimates of the hourly quantity of electricity produced by each "dispatchable" plant in the sample in scenarios with versus without the nuclear phase-out. "Dispatchable" plants include all fossil fuel power plants (i.e., lignite, hard coal, natural gas, oil) and each border point.

All other sources of power production (i.e., nuclear, wind, solar, hydro, biomass, waste and pumped storage) are treated as "non-dispatchable". Other than any ex-ante adjustments made to explicitly account for the impact of the phase-out on total production from nuclear and renewables sources, we assume that the levels of output from nondispatchable sources are determined exogenously and are thus equal to observed output levels for both the phase-out and no-phase-out scenarios. The net electricity demand that must be met by "dispatchable" sources is equal to total electricity demand minus the output from nondispatchable sources.

Strictly speaking, many "non-dispatchable" plants can actually vary their output in response to market conditions. However, this tends to be the exception rather than the rule. The bulk of non-dispatchable output comes from wind and solar resources. These resources have near-zero marginal costs and so will generally produce as much as possible based on how much the wind is blowing or the sun is shining. Renewable sources are also generally remunerated through fixed price feed-in-tariffs, reducing their incentive to respond flexibly to market conditions. The operations of many other types of nondispatchable sources are driven by important factors outside of the electricity market, such as seasonal rainfall for hydro and the supply characteristics of input fuels for waste. As such, it seems reasonable to take the operations of these plants as largely invariant to the changes in market conditions in our model. Further details on how we model output from non-dispatchable sources in the phase-out and no-phase-out scenarios are provided later in this section and in Online Appendix Section B.

Hourly data on plant-level electricity production are available for all EU member states since 2015 from ENTSOE.[8] We also model electricity imports and exports at

[8]The data include only plants with capacity greater than 100MW. This covers 100% of production from nuclear plants, 95% from lignite plants, 85% from hard coal plants, 50% from gas plants and 45% from oil plants. We treat the operating behavior of these plants as being representative

each border interconnection between Germany and its neighboring countries. For example, consider the hourly net electricity imports from France to Germany. If France exports 50 MWh of electricity to Germany, this border point would be modeled as "producing" 50 MWh. Conversely, if France imports 50 MWh of electricity from Germany, this border point would be modeled as "producing" -50 MWh.

Our analysis is primarily based on changes occurring within Germany's national borders. This is due both to data availability and our aim of evaluating whether a national policy decision, the nuclear phase-out, ultimately benefited stakeholders in Germany. In reality though, the German electricity grid operates in a manner that is fully integrated with neighboring Luxembourg and Austria. Rather than extending our analysis to directly model the behavior of plants located outside of Germany, we capture these interconnections through our modeling of cross-border flows at border points. For instance, our model is able to capture that net imports from Austria to Germany are sizable in the winter due to output from seasonal hydro plants across the border in the Alps. We believe our approach is sufficient to capture the core dynamics of interest. Moreover, moving to explicitly modeling plants in Luxembourg and Austria would be unlikely to significantly alter our findings in large part because border flows with these two countries are small relative to the role of in-country production in Germany.

The dependent variables considered in the machine learning approach are the production levels from each of the power plants and border points in the sample. In all cases, we normalize the relevant dependent variable by dividing output by the maximum production capacity of each power plant or the maximum transfer capacity

of the remaining plants with capacity less than 100MW, conditional on a range of other plant characteristics including technology type, combined heat and power functionality, and location.

of the border point. Our algorithm focuses on dependent variables that are bounded between 0 and 1; we rescale the flows from border points from their original scale of -1/1 to 0/1 when applying the algorithm. We refer to this rescaled output as the *operating rate*.

The independent variables include net electricity demand, local weather, each plant's marginal cost, the availability of other power plants, an indicator for whether the plant operates as a combined heat and power plant, and a wide range of power plant characteristics such as fuel type, efficiency, technology, and location. We estimate a model that predicts the operating rate for each power plant and border point in each hour using these independent variables. Importantly, we have data on all of the independent variables from 2010-2019. This allows us to predict hourly electricity production for all plants over the 2010-2019 period, despite only observing hourly plant-level production from 2015 onward.

We also build a predictive model for wholesale electricity prices. However, there is no cross-sectional variation in these prices; the hourly wholesale electricity price is the same throughout Germany. In this case, the independent variables for the time-series model of electricity prices include electricity demand, national average weather, and the marginal cost of the plant with the largest marginal cost among plants operating in the hour-of-sample.[9]

[9] In order to assuage concerns that the marginal cost of the marginal unit is a function of the phase-out policy, we calculate this magnitude separately for the phase-out and no-phase-out scenarios based on the intersection between the supply curve and the net demand implied by the scenario considered. See Online Appendix Section B.2.2 for further discussion.

## 5.2 Random Forest Algorithm

We predict outcomes using a Random Forest regression algorithm (Breiman, 2001). In particular, we use the Quantile Regression Forest algorithm (Meinshausen, 2006). Random forests are especially well-suited for our empirical context for several reasons.

First, each plant's production is based on a potentially complex combination of factors such as the marginal costs and availability of other plants, electricity demand at different locations, and transmission constraints. Consequently, the relationship between plant-level production and the independent variables listed above is likely to be highly non-linear and include multiple interactions. Random forest methods are well-suited to use variation in the data in order to find these interactions rather than pre-specifying how independent and dependent variables relate using polynomials or splines as in a more standard regression framework.[10]

Second, the Random Forest algorithm ensures that the support of possible predictions of an outcome is bounded by the support of the observed values of this outcome variable in the training data set. This prevents nonsensical predictions such as plants producing negative amounts of electricity or producing greater than their capacity (e.g., operating rates above 100% or below 0%).[11]

Third, the Quantile Regression Forests algorithm produces predictions for the full conditional distribution of the outcomes rather than just their expected value. This property allows us to make corrections to ensure that the total electricity supply implied by our predictions equals electricity demand. It is typically impossible to impose the constraint in this kind of empirical approach that electricity supply matches

[10]Capturing these kinds of non-linear effects is not a capability unique to random forests. Other machine learning algorithms (e.g. kernel-based methods) can also capture non-linearities.

[11]Random forests are not alone in being able to bound variables in this way. Logistic regression methods, for example, also have this desirable property through their use of a sigmoid function.

demand. However, the quantile random forest algorithm allows us to find, for each month-of-sample and each scenario considered, the percentile of the distribution of predicted plant-level output that ensures that total electricity supply equals demand.

Online Appendix Section B.4 provides more details on the determinants of the percentile chosen (e.g., net demand) as well as sensitivity analyses that show our main results are not driven by differences in the percentiles chosen across the factual phase-out versus counterfactual no-phase-out scenarios.

## 5.3 Implementation for our Application

We use the trained machine learning model to construct two data series. First, we predict hourly plant-level electricity production at each fossil fuel plant and border point using the observed values of the independent variables over 2010-2019. This provides us with electricity production levels at each plant and border point in the "factual" scenario with the nuclear phase-out. The machine learning model is necessary for estimating plant-level production even in the factual scenario because there is no hourly plant-level production data prior to 2015.

Second, we use the model to estimate hourly production for the same set of plants and border points in the counterfactual scenario where there was no nuclear phase-out. Put another way, we predict plant-level production and point-level flows assuming that the nuclear reactors that were shut down in 2011, 2015, 2017 and 2019 would have remained operational until the end of 2019. To do this, we first calculate the amount of electricity these nuclear plants would have produced in each hour-of-sample if they had remained online. We assume that the nuclear plants that were shut down would have operated at 80% of their capacity on average over the period to the end of 2019.[12] We also assume that all of the reactors that were shut down between 2011 and 2019 would have continued to operate until at least the end of 2019 in the absence of the phase-out.[13,14] The resulting estimate of counterfactual nuclear output in the absence of the phase-out is also adjusted to account for seasonal fluctuations in output that primarily reflect the timing of annual maintenance in the summer months.[15] We subtract our estimated counterfactual nuclear output from net electricity demand, thus reducing the production needed from the remaining plants and border points.

Our exposition has thus far focused on hourly plant-level production and point-level net flows. However, we utilize a similar approach to assess the impact of the phase-out on wholesale electricity prices. Further details on this, and the implementation of our machine learning algorithm in general, can be found in Online Appendix Section B.

[12]Many nuclear plants in Germany consistently achieve operating rates of around 90%. However, we opt for a lower value to capture the fact that the plants that were first to shut down tended to be older. In addition, even absent the phase-out, these plants would have needed to be taken offline for a period of time for reactor extension upgrades.

[13]The original reactor extension policy envisaged increasing the operating lifetimes of Germany's nuclear reactors by between 8-14 years. In addition, the market conditions for nuclear plants have not been as challenging in Europe as in the United States, largely because natural gas prices in Europe have been systematically higher than those in the United States until recently. Consequently, even the oldest plants that were shut down in 2011 would have been likely to remain in operation until the end of 2019 absent the phase-out.

[14]The Kruemmel nuclear power plant was placed in long-term shutdown in 2009; this plant is not assumed to be turned back on in the "no-phase-out" scenario.

[15]We make this adjustment based on observed fluctuations in monthly total nuclear production from 2012 to 2014 because there were no nuclear shutdowns during this period.

## 5.4 Accounting for Other Impacts of the Phase-Out

Our approach adjusts the level of net demand faced by fossil fuel plants and border points to reflect the nuclear production lost due to the phase-out. In doing so, we hold independent variables other than net demand fixed at their historically observed values for both the phase-out and no-phase-out scenarios. This assumption seems reasonable for independent variables such as those based on plant characteristics, temperature, and the seasonality of demand. Our approach also follows previous literature in assuming that fuel prices are unaffected by the phase-out (e.g., Grossi, Heim and Waterson (2017)).

However, the phase-out potentially impacted the level of investment in wind and solar resources. Namely, the incentives to invest in renewables might not have been as strong in the absence of the phase-out. To account for this, we assume that annual renewable production in the no-phase-out scenario would have been 30 TWh lower by 2020. We chose 30 TWh based on changes made to Germany's renewable energy targets in response to the phase-out decision.[16] A 30 TWh per year reduction amounts to a 15% decrease in renewable production by the end of our study period.

We decrease the aggregate level of renewable production in the counterfactual no-phase-out scenario by proportionally adjusting observed total hourly output from wind and solar resources.[17] To demonstrate the sensitivity of our findings to different

[16]Specifically, in 2010, Germany planned on producing at least 30% of its electricity from renewables by 2020. However, this target was increased to 35% following the 2011 phase-out decision (Jacobs, 2012). The difference between these two targets requires a change in annual renewable production of roughly 30 TWh per year by 2020.

[17]We assume that the decrease in renewable production in the absence of the phase-out grows linearly from 0 TWh per year in 2010 to the assumed amount of 30 TWh per year in 2020. To take 2015 as an example, in our baseline scenario, we assume that 15 TWh of the renewable power produced in that year was due to the nuclear phase-out. Total wind and solar production in that year was 114 TWh; therefore, in the counterfactual without the nuclear phase-out, the new

assumptions on how much of the investment in renewables is due to the phase-out, we also explore a low-case scenario where there is no response from renewable investment and a high-case scenario where the response from renewables is twice as large.[18]

It is also plausible that the phase-out altered the incentives to invest in fossil fuel power plants. Prior studies have demonstrated that, if the phase-out had not occurred, the amount of fossil-fuel generating capacity necessary to ensure that demand is met during peak hours in Germany would have been 4 GW lower by 2020 (Traber and Kemfert, 2012) and 8 GW lower by 2030 (Knopf et al., 2014). We account for this phase-out-induced increase in investment in fossil generating capacity.[19] However, the resulting increases in fossil investment costs are quite small, and do not affect the conclusions drawn from the analysis.

Finally, phase-out induced increases in wholesale electricity prices may have decreased total consumer demand for electricity. However, changes in wholesale prices are likely to have only a muted impact on customer demand because the commercial and residential customers that make up around half of Germany's total demand are highly price-inelastic.[20] Phase-out-induced increases in electricity prices would lead

total renewable production would be 114-15=99 TWh. This new total is 84% of the original; we multiply the observed values for renewable production in each hour in 2015 by 0.84 to calculate the counterfactual no-phase-out hourly levels of renewable production.

[18]Previous research on the phase-out typically assumes that investment in renewables did not accelerate due to the nuclear plant closures (Traber and Kemfert, 2012; Knopf et al., 2014).

[19]We assume that 4 GW less fossil generating capacity would be needed in the no phase-out scenario by the end of 2019. The scale of the proportional adjustment in observed fossil capacity in each year is based on assuming the decrease in fossil capacity grows linearly from 0 GW in 2010 to the assumed amount of 4 GW in 2020. For example, we assume that 2 GW of the fossil capacity online in 2015 was due to the nuclear phase-out. Total fossil capacity in 2015 was 80 GW; in the counterfactual without the phase-out, the new fossil capacity is 80-2=78 GW. This new total is 97% of the original, and we multiply the observed fossil capacity in 2015 by 0.97.

[20]Larger industrial customers are more price-elastic, and the rates they pay for electricity are also more responsive to changes in wholesale electricity prices. This is in part because the prices paid by larger industrial customers do not incorporate any portion of the renewable subsidies that smaller customers fund through their bills (BNetzA, 2020). Still, even a conservative assumption

to decreases in net electricity demand, and thus would have the same directional impact as the various renewable investment scenarios we explore. Consequently, the renewables sensitivity analyses should help capture the range of possible phase-out-induced changes in consumer electricity demand.

## 5.5 Model Performance

Figure 2(a) compares observed hourly plant-level operating rates (i.e., percentage of capacity utilized) with the predictions from the machine learning model. Specifically the predicted electricity production (scaled on the y-axis) is plotted against the observed production (x-axis) so that observations on the 45 degree line indicate perfect prediction accuracy. Each pixel in the figure represents the predicted vs. observed operating rate in increments of 2% and darker areas correspond to a higher number of plant-hour (or plant-year) observations.

We check the out-of-sample cross-validated performance to avoid overfitting and to give a fair assessment of how the model may perform when used to make predictions for the counterfactual no-phase-out scenario. For the hourly data in panel (a), the cross-validated out-of-sample $R^2$ is 0.41.

However, even this small level of prediction error understates the relevant prediction accuracy of the machine learning model. Specifically, we will primarily use the predictions from our model to compare outcomes under the phase-out and no-phase-out scenarios at the plant-year level. We therefore also evaluate the predictive performance of the model at this level of aggregation. Specifically, Figure 2(b) plots predicted versus observed annual average operating rates. As the figure shows, the

regarding the price-elasticity of these consumers is unlikely to result in a large change in demand given that we estimate a relatively small increase in wholesale prices due to the phase-out.

performance is substantially improved, with most of the observations clustered close to the 45 degree line, and the areas of systematic error largely disappear. The cross-validated out-of-sample $R^2$ rises to 0.84.

Lastly, we also use the machine learning model to predict counterfactual hourly wholesale prices. This model performs well, with a cross-validated out-of-sample $R^2$ of 0.88. By far the most important predictor in this model is the marginal cost of the marginal plant. This is consistent with how prices are determined in the wholesale market.

# 6 Results on the Impact of the Phase-Out

This section presents the main results on the full range of impacts of the nuclear phase-out over the 2010-2019 analysis period. Specifically, we compare the market and environmental outcomes with versus without the nuclear phase-out using the predictions from our machine learning model.

## 6.1 Generation, Net Imports, and Prices

Figure 3(a) reports the monthly average difference in predicted production and net imports (in TWh) with the phase-out minus without the phase-out policy. We report monthly average differences in fossil-fuel-fired electricity production (grey diamonds), net imports (red circles), renewable electricity production (green triangles) and nuclear electricity production (purple squares). The start of the nuclear phase-out in March 2011 is marked by the vertical black dashed line.

Total nuclear production decreases by roughly 4 TWh per month immediately

Figure 2: Machine Learning Model Performance: Plant-Level Electricity Production

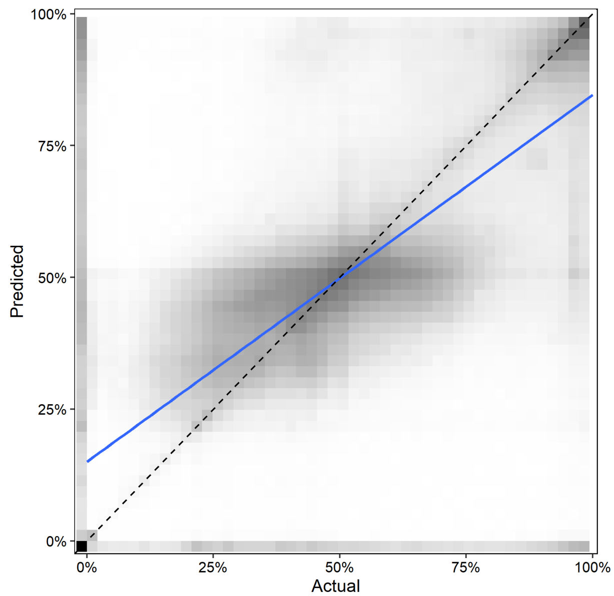


(a) Hourly Data

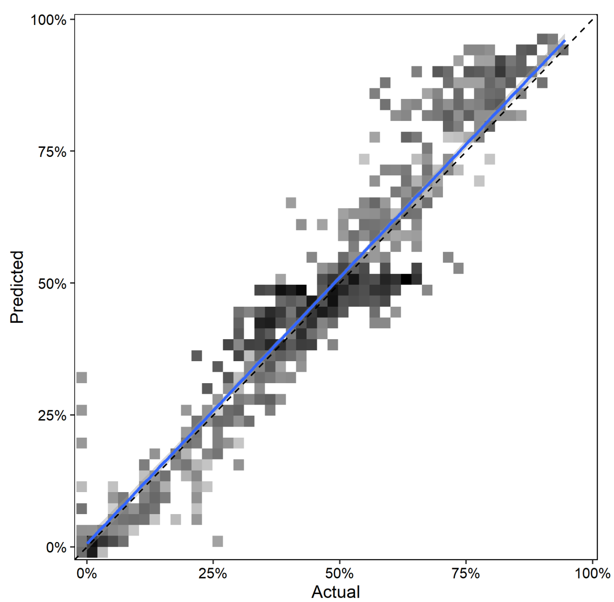


(b) Annual Data

This figure illustrates the accuracy of the plant-level predictions from the machine learning model presented in Section 5. The model predicts the operating rate of each power plant in each hour, where a value of 0% means that a plant is offline and a value of 100% means that the plant is running at maximum capacity. Values on the 45 degree line indicate perfect accuracy, and we summarize this both visually with the blue linear fitted line and by computing measures of R-Squared. We compute these metrics using out-of-sample five-fold cross-validation. Darker areas indicate higher numbers of plant-hour (or plant-year) observations. Each pixel represents the predicted vs. actual operating rate in increments of 2%. Panel (a) shows prediction accuracy at an hourly timescale; the R-squared is 0.41. Panel (b) shows prediction accuracy after taking annual averages of our hourly predictions; the R-squared is 0.84.

after the announcement of the policy, with an additional reduction of 2 TWh per month later in the sample period as more nuclear plants are taken offline. The observed seasonal fluctuations of this impact are due to the fact that nuclear reactors typically schedule their maintenance and refuelling outages in the summer months.

In our baseline scenario, we assume that some of this lost nuclear production was replaced by accelerated investment in renewable sources due to the phase-out policy. Consistent with this assumption, Figure 3 documents a steady rise in phase-out-induced increases in renewable electricity production, with an additional 2.5 TWh per month by 2020.

We use our machine learning model to estimate the remaining contribution of various dispatchable sources in replacing the lost nuclear production. We find that the phase-out caused a large increase in fossil-fuel-fired electricity production of 2-3 TWh per month, which persists over our entire sample period. The phase-out also caused a smaller increase in net imports of electricity of around 1 TWh per month. The largest contributors to net imports were the Czech Republic and France, which is consistent with the higher levels of market integration for these countries highlighted by Grossi et al. (2018).

Figure 3(b) reports the estimated impact of the nuclear phase-out on wholesale electricity prices in Euros per MWh. The estimates show that the phase-out resulted in an increase in wholesale prices of around 1 euro per MWh, although this difference diminished in recent years, and was in fact negative for part of 2018. The estimates also suggest that the phase-out may have exacerbated episodic increases in prices.

Table 2 presents the annual averages of electricity production and prices over the 2012-2019 sample period with versus without the phase-out. The estimates reveal

Figure 3: Estimated Impact of the Nuclear Phase-Out on Electricity Production and Prices

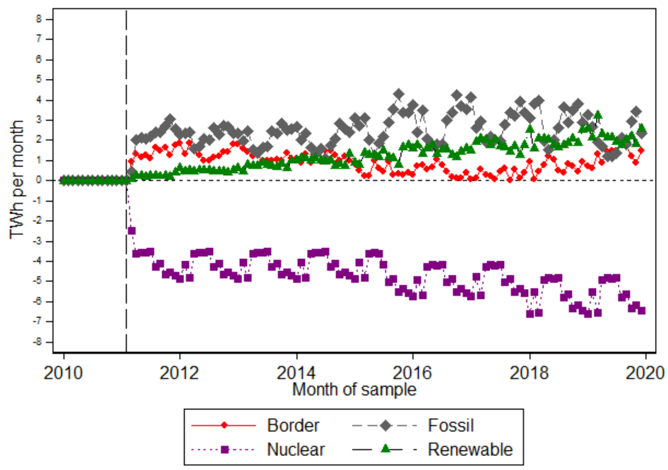


(a) Electricity Production

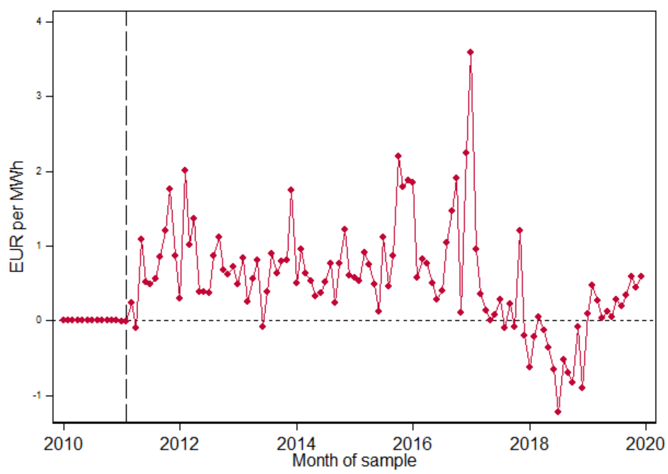


(b) Wholesale Electricity Prices

This figure plots the monthly difference between the predictions from our machine learning model with the phase-out minus without the phase-out. The start of the phase-out in March 2011 is marked by the vertical black dashed line. Panel (a) reports the estimates for all fossil-fuel-fired electricity production (grey diamonds), net imports (red circles), renewable electricity production (green triangles), and nuclear production (purple squares). Panel (b) presents the change in wholesale electricity prices.

Table 2: Estimated Impact of the Nuclear Phase-Out on Wholesale Electricity Prices, Electricity Production by Source, and Net Imports

| | **Average with Phase-Out (1)** | **Average without Phase-Out (2)** | **Change (3)** | **Change (%) (4)** |
|---|---|---|---|---|
| **Production (TWh/Year)** | 433.1 | 433.4 | -0.3 | -0.1 |
| Nuclear | 81.5 | 139.1 | -57.6 | -41.4 |
| Lignite | 147.2 | 139.6 | 7.6 | 5.4 |
| Hard Coal | 86.1 | 69.4 | 16.7 | 24.1 |
| Gas | 31.6 | 26.7 | 4.9 | 18.4 |
| Oil | 7.7 | 6.5 | 1.2 | 18.5 |
| Net Electricity Imports | -36.8 | -47.4 | 10.6 | 22.4 |
| Wind + Solar | 116.0 | 99.5 | 16.5 | 16.6 |
| **Wholesale Prices (Euros/MWh)** | 37.9 | 37.3 | 0.6 | 1.6 |

**Notes:** This table reports annual average electricity production by source and wholesale electricity prices with versus without the nuclear phase-out, as estimated using our machine learning approach. These annual averages are calculated using data spanning from 2012 to 2019.

that the phase-out caused wholesale electricity prices to increase by €0.6 per MWh on average, a 1.6% increase relative to the prices that would have prevailed on average if the phase-out had not occurred. Nuclear production fell by an average of 57.6 TWh per year during the phase-out period, corresponding to a 41.4% decline. Average annual generation from fossil-fuel plants increased by 16.7 TWh for hard coal, 7.6 TWh for lignite, and 4.9 TWh for gas. The phase-out also caused net imports to increase by 10.6 TWh per year on average. Under our base-case assumption for renewables investments, we see an increase in average annual renewable production of 16.5 TWh due to the phase-out.

A caveat with our approach is that, following prior research that has modeled the German electricity sector (Egerer, 2016*b*), we do not adjust wholesale prices when calculating the predictions of net imports with versus without the nuclear phase-out.

However, Grossi et al. (2018) demonstrates that prices in neighboring countries rose due to the phase-out in the short-run. Consequently, since we do not adjust wholesale prices in neighboring countries when calculating predictions, the phase-out-induced increase in the quantity of net imports reported in Table 2 is likely an upper bound. Our estimate of the overall social cost of the phase-out to *German producers and residents* is thus likely to be conservative because the bulk of our estimated cost comes from increases in local air pollution due to increases in coal-fired production in Germany rather than increases in net imports.

## 6.2 Private Costs and Benefits

Table 2 showed that the phase-out caused fossil-fuel and renewable generation to increase, and nuclear production to decrease. Table 3 reports the changes in average annual total variable costs, fixed costs, and revenues implied by this change in generation mix. The central component of variable costs is calculated by multiplying each plant's hourly production with our estimate of its marginal cost in the hour. For fossil fuel plants, marginal costs are calculated as the sum of the plant's fuel cost per MWh and an assumed amount of variable operating and maintenance costs that differs by fuel type. Fuel costs are calculated by converting the price of the relevant input fuel to euros per MWh based on the calorific content of the fuel and the plant's thermal efficiency (i.e., how well the plant converts units of input heat to units of electricity output).

Other variable operating and maintenance costs, as well as any fixed investment and maintenance costs, vary by source type. These magnitudes are taken from a recent study on generation costs (EIA, 2019). The key benefit of this study is that it draws on a wide range of reported cost data from actual plants. It is thus better suited for our retrospective analysis of the phase-out policy.[21]

We assume that nuclear plants have a marginal cost of approximately €10/MWh based on prior research on Germany's power sector (Egerer, 2016*a*). This is confirmed by company reports from two European nuclear plant operators, EON and EDF, which also have marginal fuel costs of approximately €10/MWh over our analysis period. The same industry reports indicate that the other fixed operating, maintenance and investment costs at these plants likely amount to a further €20/MWh, resulting in overall costs for the continued operation of existing nuclear plants of roughly €30/MWh. Beyond this, we also account for the costs of the nuclear power plant lifetime extensions that were explicitly affected by the phase-out policy. Keppler (2012) argues that extending the lifetime of the nuclear reactors in Germany would have required investments of roughly €500 million per reactor, or €8.5 billion in total. These investment costs were avoided due to the nuclear phase-out.

Wind and solar power are assumed to have zero marginal operating costs. To account for the fixed operating and capital costs of renewables, we rely on levelized cost values for wind and solar plants from the International Renewable Energy Agency (IRENA, 2020). These values are specific to Germany and capture annual average costs for plants built in each year. Lastly, for net imports, we quantify the costs as being the net quantity of imported electricity multiplied by the wholesale price in the relevant neighboring country.

Revenues are calculated as the product of plant-level production and wholesale electricity prices. Due to a lack of data, we do not include any additional revenues

[21]The EIA values are similar to corresponding estimates from other studies which are commonly used in the literature (e.g., Lazard (2019)).

plants may receive outside of the wholesale market, such as capacity payments, ancillary services payments, and subsidies. As such, our measure of revenues understates the actual revenues accrued by power producers. This is particularly relevant for renewable sources which derive a large portion of their revenues from subsidies.

Table 3 reports estimates of the impact of the nuclear phase-out on revenues and private costs for power plants. The table is structured like Table 2. The first result, not surprisingly, is that the nuclear phase-out had a large effect on the revenues of the nuclear plants that were shut down. Specifically, total wholesale revenues across all nuclear plants declined by €2.05 billion per year on average as a result of the phase-out. Despite this, wholesale revenues are still larger than nuclear plants' variable and annualized fixed costs in both scenarios, indicating the continued profitability of nuclear power over our analysis period.

The revenues previously earned by the shut-down nuclear plants were primarily redistributed to both fossil plants (hard coal, lignite, and natural gas plants) and to renewables. The increased use of these other sources led to a net increase in variable costs. This was largely driven by fossil fuel plants, with €0.20, €0.51 and €0.20 billion per year in additional variable costs for lignite, hard coal and natural gas power plants. There was also an increase in fixed costs. This was largely driven by new renewable plants, with €1.35 billion per year in additional fixed costs outweighing the savings from cancelling the nuclear reactor lifetime extensions.

The redistribution of profits among electricity producers has interesting implications for the political economy surrounding the nuclear phase-out. In particular, the four large firms that owned nuclear plants in Germany - E.ON, RWE, EnBW and Vattenfall - all publicly opposed the policy. These firms have since been awarded €2.4 billion in compensation from the German government to cover losses they in-

Table 3: Estimated Impact of the Nuclear Phase-Out on Wholesale Revenues and Private Costs

| | **Average with Phase-Out (1)** | **Average without Phase-Out (2)** | **Change (3)** | **Change (%) (4)** |
|---|---|---|---|---|
| **Wholesale Revenues** ($\frac{\textbf{Bn. Euros}}{\textbf{Year}}$) | 17.30 | 17.43 | -0.13 | -0.7 |
| Nuclear | 2.99 | 5.04 | -2.05 | -40.6 |
| Lignite | 5.43 | 5.06 | 0.37 | 7.3 |
| Hard Coal | 3.19 | 2.54 | 0.65 | 25.6 |
| Gas | 1.18 | 0.98 | 0.20 | 20.4 |
| Oil | 0.28 | 0.23 | 0.05 | 21.5 |
| Wind + Solar | 4.22 | 3.57 | 0.65 | 18.2 |
| **Variable Costs** ($\frac{\textbf{Bn. Euros}}{\textbf{Year}}$) | 7.82 | 6.92 | 0.90 | 13.0 |
| Nuclear | 0.78 | 1.34 | -0.56 | -41.8 |
| Lignite | 3.70 | 3.50 | 0.20 | 5.7 |
| Hard Coal | 2.63 | 2.12 | 0.51 | 24.1 |
| Gas | 1.38 | 1.18 | 0.20 | 17.0 |
| Oil | 0.94 | 0.77 | 0.16 | 20.6 |
| Net Electricity Imports | -1.61 | -2.00 | 0.39 | 19.5 |
| Wind + Solar | 0.00 | 0.00 | 0.00 | . |
| **Fixed Costs** ($\frac{\textbf{Bn. Euros}}{\textbf{Year}}$) | 21.98 | 21.71 | 0.28 | 1.3 |
| Nuclear | 1.57 | 2.73 | -1.16 | -42.5 |
| Lignite | 1.03 | 1.00 | 0.03 | 3.0 |
| Hard Coal | 1.29 | 1.26 | 0.04 | 3.2 |
| Gas | 0.63 | 0.62 | 0.02 | 3.2 |
| Oil | 0.10 | 0.10 | 0.00 | 0.0 |
| Net Electricity Imports | 0.00 | 0.00 | 0.00 | . |
| Wind + Solar | 17.35 | 16.00 | 1.35 | 8.4 |

**Notes:** This table reports average annual total wholesale revenues and private costs under the phase-out and no phase-out scenarios. All entries are annual averages over 2012-2019 based on predictions from our machine learning model. Wholesale revenues are defined as the product of each plant's hourly production with the hourly wholesale electricity price. Variable costs are the product of each plant's hourly production with its marginal cost in the hour. Marginal costs include fuel costs per MWh as well as variable operating and maintenance costs (in €/MWh). Fixed costs are the product of each plant's capacity with its annual fixed capital expenditure as well as fixed operating and maintenance costs (in €/MW/yr).

curred as a result of the phase-out. However, it is possible that their opposition was tempered somewhat by the fact that, in addition to their nuclear plants, all four companies had large fossil plant portfolios both in Germany and across Europe. Our finding that fossil plants played a large role in replacing the lost nuclear production indicates that any reduction in their profits due to the nuclear closures may have been cushioned by the increased profitability of their fossil plants. Moreover, if these firms' nuclear plants had remained open, they would have been subject to a new nuclear fuel tax. This tax would have greatly reduced the inframarginal rents earned by nuclear plants. So, while public concern about nuclear risks was likely the key determinant of the phase-out decision in 2011, the policy might also have been more aligned with the interests of firms than it first seemed.

Lastly, Germany remains a net exporter throughout our sample period, both with and without the phase-out. The phase-out did reduce net exports, and valuing the change at the electricity prices of the relevant neighboring countries resulted in a net increase in variable costs (i.e., lost export revenues) of €0.39 billion per year. This calculation holds fixed the prices of neighboring countries at observed levels across the phase-out and no-phase-out scenarios. However, even if we assume that the increases in prices in neighboring countries due to the phase-out estimated in Grossi et al. (2018) persisted to the end of 2019, this would only increase the total private costs to Germany by about 3.3%.[22] Even this 3.3% increase in private costs is likely to be an upper bound because: (1) as mentioned above, the phase-out-induced increase in the quantity of net imports we estimate is likely to be an upper bound

[22]Grossi et al. (2018) uses an approach similar to the event study analysis in Section 4, finding that the phase-out led to increases in wholesale prices in neighboring countries in 2011 and 2012. This would increase the price of additional imports to Germany by just under 10%. If we assume those price increases persist to the end of 2019, and apply them to the 10.6 TWh per year of additional net imports in Table 2, this would still only increase total private costs by 3.3%.

and (2) the increases in prices found by Grossi et al. (2018) for 2011 and 2012 may have dissipated from 2013-2019.

## 6.3 External Costs

This subsection presents an analysis of the environmental costs associated with the increase in fossil-fuel-fired production caused by the nuclear phase-out. Specifically, burning fossil fuels emits both global pollutants such as carbon dioxide that contribute to climate change and local air pollutants that adversely impact the health of exposed populations.

We begin by describing how we estimate the change in carbon emissions due to the phase-out. We first calculate the change in the amount of fuel burned by each fossil fuel plant due to the phase-out using our predictions of each plant's hourly production and thermal efficiency. We then use the carbon intensity of different fuels documented in industry reports to convert changes in fuel burned to changes in plant-level $CO_2$ emissions.[23]

We also estimate the change in emissions of air pollutants due to changes in each plant's production levels caused by the phase-out. Similar to the approach for $CO_2$ emissions, we translate changes in fuel use into changes in emissions using plant-level emissions rates for each local air pollutant from the EU Large Combustion Plant Directive (LCPD). The LCPD database provides annual plant-level data on fuel inputs and emissions of sulfur dioxide ($SO_2$), nitrogen oxides ($NO_x$) and particulate matter (PM). The LCPD data covers the vast majority of large fossil plants in

[23]The carbon intensities we use are 93.6 tCO2/TJ for hard coal, 55.9 tCO2/TJ for gas and 74.0 tCO2/TJ for oil. We consider three different intensities for lignite depending on the mining region that the plant sources its coal from. These are 113.3 tCO2/TJ (Rhineland), 111.2 tCO2/TJ (Lusatian) and 102.8 tCO2/TJ (Central).

Germany.[24] We assign the small number of plants not in the LCPD database an emissions factor based on the average emissions factor of plants with the same fuel type.

We next monetize the damages caused by local air pollution emissions. For this, we rely on two studies that estimate the health impacts of ambient air pollution in Europe (EEA, 2014; Jones et al., 2018). In particular, Jones et al. (2018) provide estimates of the annual health damages from the ambient air pollution emitted by roughly four hundred of the largest coal-fired power plants in Europe. We use these data to convert our predicted increases in plant-level kilotons of $SO_2$, $NO_x$ and PM emissions into monetized health damages.[25]

Table 4 presents the results of this analysis. Specifically, this table reports the fuel-specific average annual total emissions of $CO_2$ (in Megatonnes, Mt) and three local air pollutants: $SO_2$, $NO_x$, and PM (in Kilotonnes, Kt). Lignite and hard coal plants are by far the two largest polluters, contributing more than 90% of the emissions of each of the pollutants listed. Emissions from lignite and hard coal plants also contribute the most in terms of mortality and monetized pollution damages.

In aggregate, the phase-out led to an increase in $CO_2$ emissions of 26.2 Mt per year among lignite, hard coal, gas, and oil plants. This corresponds to a 11% increase relative to the scenario without the nuclear phase-out. This increase in $CO_2$ emissions

[24]The data covers 99% of lignite capacity, 98% of hard coal capacity, 90% of gas capacity and 91% of oil capacity.

[25]We assume that increased emissions at a given fossil fuel plant in Germany would have the same health damages per unit of emissions as if they were emitted at the nearest location for which we have health damages estimates. The capacity weighted average distance between each of the power plants in our data set and the closest of the 400 locations with damage estimates from Jones et al. (2018) is: 0.9km for lignite plants, 0.2km for hard coal plants, 12km for gas plants, and 20km for oil plants. Consistent with Jones et al. (2018) and many other studies, our approach assumes that the marginal damage per unit of pollutant emissions is linear and additive.

Table 4: Estimated Impact of the Nuclear Phase-Out on Annual $CO_2$ Emissions and Local Air Pollution Mortality Damages

| | Average with Phase-Out (1) | Average without Phase-Out (2) | Change (3) | Change (%) (4) |
|---|---|---|---|---|
| **$CO_2$ Emissions ($\frac{Mt}{Year}$)** | 264.5 | 238.3 | 26.2 | 10.99 |
| Lignite | 167.5 | 159.1 | 8.5 | 5.34 |
| Hard Coal | 77.6 | 62.8 | 14.8 | 23.56 |
| Gas | 13.2 | 11.2 | 2.0 | 17.79 |
| Oil | 6.1 | 5.2 | 1.0 | 19.40 |
| **$SO_2$ Emissions ($\frac{Kt}{Year}$)** | 127.2 | 115.7 | 11.4 | 9.85 |
| Lignite | 86.2 | 82.3 | 3.9 | 4.74 |
| Hard Coal | 35.0 | 28.2 | 6.8 | 24.15 |
| Gas | 0.9 | 0.8 | 0.1 | 12.35 |
| Oil | 5.0 | 4.4 | 0.6 | 13.51 |
| **$NO_x$ Emissions ($\frac{Kt}{Year}$)** | 175.8 | 158.8 | 17.1 | 10.77 |
| Lignite | 111.5 | 105.8 | 5.6 | 5.29 |
| Hard Coal | 48.7 | 39.5 | 9.2 | 23.30 |
| Gas | 9.0 | 7.8 | 1.3 | 16.71 |
| Oil | 6.7 | 5.6 | 1.0 | 17.73 |
| **PM Emissions ($\frac{Kt}{Year}$)** | 4.61 | 4.18 | 0.42 | 10.04 |
| Lignite | 3.02 | 2.88 | 0.14 | 4.86 |
| Hard Coal | 1.41 | 1.15 | 0.26 | 22.60 |
| Gas | 0.06 | 0.05 | 0.01 | 21.37 |
| Oil | 0.12 | 0.11 | 0.02 | 18.70 |
| **Mortality ($\frac{Excess\ Deaths}{Year}$)** | 6,852.8 | 6,053.0 | 799.8 | 13.21 |
| Lignite | 3,779.1 | 3,590.6 | 188.5 | 5.25 |
| Hard Coal | 2,616.7 | 2,074.0 | 542.7 | 26.17 |
| Gas | 267.2 | 228.6 | 38.6 | 16.88 |
| Oil | 189.8 | 159.7 | 30.1 | 18.85 |
| **Monetized Excess Mortality ($\frac{Bn.\ Euros}{Year}$)** | 17.24 | 15.22 | 2.01 | 13.20 |
| Lignite | 9.51 | 9.03 | 0.47 | 5.20 |
| Hard Coal | 6.58 | 5.22 | 1.36 | 26.07 |
| Gas | 0.67 | 0.58 | 0.10 | 17.39 |
| Oil | 0.48 | 0.40 | 0.08 | 19.91 |

**Notes:** This table reports estimates for emissions of $CO_2$ as well as three local air pollutants: $SO_2$, $NO_x$, and PM, and estimates of the mortality damages from all three of these local air pollutants. All values are average annual totals based on predicted plant-level generation from 2012 to 2019 for fossil-fuel-fired plants in Germany. Emissions are the product of each plant's hourly generation with our estimate of their emissions rate. We ignore other small sources of emissions from biomass, landfill gas or waste plants. The estimates in the table are emissions and damages in Germany and do not consider changes in emissions in neighboring countries due to changes in net imports. The air pollution damages reported in the last row of the table only include the monetized costs associated with premature mortality due to air pollution exposure.

was primarily attributable to an increase in emissions from hard coal plants, with lignite, gas, and oil making up the remainder. Valuing these carbon emissions at a social cost of carbon of $125/tCO2 would imply that the phase-out resulted in climate damages of €3.0 billion per year (Carleton and Greenstone, 2021). However, increases in carbon emissions in Germany might be offset by decreases in carbon emissions elsewhere since German power plants are covered by the EUETS (a carbon cap-and-trade market spanning the European Union). For this reason, we exclude climate damages from our preferred estimates of the social cost of the phase-out.

However, previous work suggests that political considerations play a role in the total number of credits allocated as part of the EUETS (Koch et al., 2016). The phase-out may have led to more credits being put into circulation through this political channel. Consequently, our core estimate of the social cost of the phase-out, which does not include climate damages, represents a lower bound; the social cost may be larger than our estimate if regulators increased the EUETS cap due to political pressure from German policymakers following the phase-out.

The phase-out also led to a roughly 10-11% increase in the emissions of each of the three ambient air pollutants considered. Again, this increase is due primarily to increased emissions from hard coal plants. However, the phase-out also led to large proportional increases in emissions from gas and oil plants, as shown in column (4).

The bottom two rows of Table 4 report the estimated impacts of the phase-out on annual mortality and on annual monetized mortality damages associated with excess emissions of $SO_2$, $NO_x$, and PM. From 2012 to 2019, emissions from fossil plants in the phase-out scenario account for almost 6,900 excess deaths per year, corresponding to a monetized damage of €17.2 billion in mortality costs each year. By comparing these estimates to the estimates from the no-phase-out scenario, we can attribute 800 excess deaths per year on average to the phase-out. This corresponds to €2.0 billion per year in monetized damages, and represents a 13% increase in damages relative to the scenario without the nuclear phase-out.[26] 91% of the monetized mortality damages from the phase-out are driven by emissions from lignite (23% of damages) and hard coal plants (68% of damages).

Our focus thus far has been on external costs incurred as a consequence of phase-out-induced changes in output from plants in Germany.[27] However, the phase-out also resulted in an increase in net imports. This will have changed external costs in neighboring countries. To get a sense of the potential external costs associated with the increase in net imports, we calculate the emissions intensity of imports for each neighboring country and then follow the same valuation approach set out above. We find that increased net imports resulted in additional monetized damages of €0.25 billion per year. The vast majority of this is due to increased emissions in the Czech Republic, and to a lesser extent Denmark and the Netherlands. Despite being the largest source of additional imports, France has a relatively clean electricity supply mix and so does not see a significant increase in external costs. Overall, the external costs of imports are small relative to the impacts from changes in production in Germany.

[26] We use an estimate of the value of statistical life (VSL) of €2.56 millions for Germany which follows from Jones et al. (2018) and convert to 2019 values by adjusting for inflation. In the robustness analysis below, we also consider a larger VSL of €6.7 millions taken from Viscusi and Masterman (2017) which is more in line with valuations undertaken by the US EPA.

[27] The pollution transport model used by Jones et al. (2018) captures the fact that emissions at power plants in Germany can cause health impacts in neighboring countries. We are unable to fully disentangle these cross-border dependencies with the available information, so a small portion of damages from emissions from German plants is incurred by people in neighboring countries.

## 6.4 Alternative Estimates of Local Air Pollution Costs

The health damages in Table 4 are calculated by linking data on plant-level emissions rates to pollution exposure using atmospheric chemistry modeling. As a complement to these estimates, Online Appendix Section C presents a secondary approach to estimate the external costs imposed by increased ambient air pollution caused by the phase-out. In this approach, we use granular air pollution monitor data to examine how changes in output from hard coal and lignite plants affect ambient $PM_{2.5}$ concentration levels. The estimates from this approach indicate that the phase-out resulted in 334 additional excess deaths per year.

Taken together, the results in Table 4 and Online Appendix Table C.2 paint a consistent picture of the impact of the nuclear phase-out on pollution-caused excess death. Our estimates attribute 330 to 800 excess deaths per year to the increase in air pollution caused by the nuclear phase-out. This amounts to monetized health damages of €0.9 to €2.0 billion per year during the 2012-2019 period. Our preferred estimate is the €2.0 billion per year in damages calculated based on reported emissions (Table 4). This is because the analysis using reported emissions considers a more complete set of pollutants and draws on a more sophisticated analysis of pollution transport and exposure. The results presented in Online Appendix Table C.2 based on our estimated relationship between $PM_{2.5}$ levels and electricity production serve as a valuable complementary validation exercise based on an entirely distinct approach.

## 6.5 Social Costs

This subsection brings the analysis together by summarizing the benefits and the full range of private and external costs of the nuclear phase-out. The private costs of the phase-out consist of the aggregate changes in the variable and fixed costs incurred by power plants in Germany as well as any costs from changes to net imports. The external costs of the phase-out are measured by the damages from mortality and morbidity caused by the additional air pollution attributable to phase-out-induced changes in the electricity production mix.[28] The sum of measured private and external costs is our estimate of the social cost of the nuclear phase-out.

Table 5 reports the estimates of the aggregate annual private and external costs of the phase-out. The top panel reports social costs as well as each private and external cost component for the base-case assumptions on the growth of renewables driven by the phase-out and the value of statistical life (VSL) used to monetize pollution-caused premature mortality. The phase-out led to a replacement of low cost nuclear production with higher cost sources such as fossil fuels and net imports. This resulted in private costs increasing by €1.2 billion (or 4.1%) on average per year in Germany, mostly due to an increase in variable costs. This increase in private costs, however, is smaller than the €2.1 billion annual increase in external costs associated with the phase-out. Specifically, burning fossil fuels to produce electricity rather than using nuclear plants led to an increase in local air pollution emissions, which in turn led to increases in premature mortality and adverse morbidity events. Overall, we estimate that the annual social cost of the nuclear phase-out is €3.3 billion per year.

[28]Morbidity costs per MWh of output are reported in Jones et al. (2018); we apply the same methodology as for mortality in Table 4 to calculate aggregate morbidity costs.

Table 5: Impact of the Phase-Out on Annual Average Private and External Costs

| | **Average with Phase-Out (1)** | **Average without Phase-Out (2)** | **Change (3)** | **Change (%) (4)** |
|---|---|---|---|---|
| **Social Costs** ($\frac{\textbf{Bn. Euros}}{\textbf{Year}}$) | 48.27 | 44.95 | 3.32 | 7.4 |
| | | **Private** | | |
| Variable | 7.82 | 6.92 | 0.90 | 13.01 |
| Fixed | 21.98 | 21.71 | 0.28 | 1.29 |
| | | **External** | | |
| Mortality | 17.24 | 15.22 | 2.01 | 13.20 |
| Morbidity | 1.23 | 1.10 | 0.13 | 11.82 |

| | **Average with Phase-Out (1)** | **Average without Phase-Out (2)** | **Change (3)** | **Change (%) (4)** |
|---|---|---|---|---|
| Base VSL, Low Renewables | 48.27 | 44.82 | 3.45 | 7.7 |
| **Base VSL, Base Renewables** | **48.27** | **44.95** | **3.32** | **7.4** |
| Base VSL, High Renewables | 48.27 | 45.05 | 3.21 | 7.1 |
| High VSL, Low Renewables | 76.14 | 68.04 | 8.10 | 11.9 |
| High VSL, Base Renewables | 76.14 | 69.57 | 6.57 | 9.4 |
| High VSL, High Renewables | 76.14 | 71.01 | 5.13 | 7.2 |

**Notes:** This table reports annual averages over 2012-2019 of the private and external costs incurred with versus without the phase-out. Private costs are the variable and fixed costs associated with power plants in Germany plus any changes in net imports (valued at the wholesale electricity price). External costs consist of excess mortality and morbidity costs from air pollution emissions. The "Social Costs" row reports the sum of private and external costs.

The bottom panel reports the estimates of annual social costs with and without the phase-out for six scenarios based on assumptions on the growth of renewables induced by the phase-out (the low-case, base-case, and high-case scenarios described in Section 5.4) and the Value of Statistical Life (VSL) used to monetize the excess deaths (base and high). Following Jones et al. (2018), the base-case VSL we use is €2.56 million. We also consider an alternative VSL of €6.7 million, taken from Viscusi and Masterman (2017), which is in line with the approach taken by the United States Environmental Protection Agency.[29]

We also examine the implications of 3 scenarios for renewables investment in response to the phase-out. Our base-case scenario assumes that the phase-out induced investments in renewables to produce an additional 30 TWh per year by 2020. In the low-case renewables scenario we assume no additional renewables investment due to the phase-out, and in the high-case scenario, we assume that the phase-out caused investments in renewables sufficient to increase production by an additional 60 TWh per year by 2020.[30]

Column (3) of the bottom panel of Table 5 reports the estimated change in annualized social costs (sum of private and external costs) attributable to phase-out. Across the six scenarios, the social costs range from €3.21 to €8.10 billion per year. The €3.32 billion estimate reported in the top panel of the table can be seen under the "Base VSL, Base Renewables" row. A higher assumed VSL necessarily leads to higher external costs and hence total costs. This is evident by comparing the rows

[29]The lower VSL we consider is derived from studies based on stated preference methods and may thus suffer from hypothetical bias. In contrast, the high-case estimate of the VSL is derived from studies using revealed preference methods (Viscusi and Masterman, 2017).

[30]It is possible that changing the pace and scale of investment in renewable sources could have additional dynamic impacts on costs, such as promoting learning-by-doing or spurring new grid infrastructure investments. These dynamic impacts are not incorporated into our analysis, and are likely second order relative to the capital and operating costs associated with renewable investments.

with “base VSL” versus “high VSL” in the bottom panel.

Looking across the “base VSL” rows, we observe that social costs fall as the assumed growth of renewables attributable to the phase-out increases. Higher renewable growth leads to increased private costs. This is because investment in renewables is expensive relative to increased production from fossil fuel plants. However, higher renewable growth also decreases external costs. This is because the additional renewables displace fossil fuel production, decreasing the monetized health impacts from local air pollution. We find that the increase in private costs across the “low”, “base”, and “high” renewables scenarios is smaller than than the reduction in monetized health impacts as renewables grow from the low-case to the high-case. This effect is even more pronounced in the “high VSL” rows where we assume that the premature mortality due to air pollution exposure is more costly.

There are some noteworthy limitations to the analysis in Table 5. While nuclear power plants emit virtually no global or local air pollution, nuclear energy does come with catastrophic accident risk and requires storing the waste that results from nuclear production. Estimates from the literature suggest that the external costs of nuclear power due to waste storage and accident risk fall between €1-4 per MWh (D’haeseleer, 2013). This wide range is due to differing estimates of accident probabilities and severity, as well as varying assumptions on discount rates. If we value the external costs of nuclear power at €4 per MWh, the expected benefits from the nuclear phase-out are very small at just €0.2 billion per year. This is clearly far smaller than the €3-8 billion per year in annual social costs estimated above. Even if we assume that the external costs of nuclear power are an order of magnitude larger than those from existing literature, the expected benefits of the nuclear phase-out are still smaller than our central estimate of the social cost of the phase-out policy.

# 7 Conclusions and Policy Discussion

Following the Fukushima disaster in 2011, German authorities made the unprecedented decision to immediately shut down almost half of the country's nuclear power plants, and to shut down all of the remaining nuclear power plants by 2022. We quantify the economic and environmental costs of this decision. Our analysis indicates that the phase-out of nuclear power has come with an annual social cost to Germany of roughly €3-8 billion per year. The majority of this cost is due to the 800 excess deaths per year resulting from the local air pollution emitted by coal-fired power plants operating in place of the shutdown nuclear plants. The scale of the social costs from the nuclear phase-out exceeds even the largest estimates of the expected benefits from the phase-out.

In light of the scale of the costs we identify, it is striking that the nuclear phase-out continues to receive widespread support, with more than 81% in favor in a 2015 survey (Goebel et al., 2015). Concerns about the risks of nuclear accidents and storing nuclear waste are at the core of the anti-nuclear sentiment, and motivated the decision to phase-out nuclear power (Ethics Commission, 2011). Thus far we have compared the social costs of the phase-out against its *expected* benefits. However, nuclear accident risks are highly uncertain and the costs associated with nuclear waste disposal are also arguably uncertain. It is therefore possible that a sufficiently risk-averse policymaker could phase-out nuclear to avoid the tail risks associated with nuclear accidents and waste disposal, even though the costs of the phase-out are higher than its benefits in expectation.

To get a sense of the level of risk aversion implied by the decision to proceed with the nuclear phase-out, we calculate the probability of a major nuclear accident that would result in the expected benefits from the phase-out being equal to its costs. For this back-of-the-envelope calculation, assume that, absent the phase-out, nuclear plants would have been shut down in the same order but approximately ten years later, with the last plants closing in 2032 instead of 2022. With the phase-out starting part way through 2011, this gives roughly two decades over which the policy would reduce nuclear production. Our estimate for the cost of the phase-out is €3-8 billion per year; this implies a cumulative cost of the phase-out of around €60-160 billion over this twenty year period. The estimated cost of the Fukushima accident is 35-80 trillion yen, or €280-640 billion (JECR, 2019). Assume for simplicity that there can either be no accidents or there can be one Fukushima magnitude accident during this twenty year period. The probability of this Fukushima-scale accident occurring would have to be anywhere from around 1 in 10 to as high as 1 in 2 for the expected benefits of the phase-out to be equal to the expected costs.[31] A 1 in 10 chance, let alone a 1 in 2 chance, is far greater than even the most conservative estimates of the probability of an accident of this magnitude occurring in Germany. For instance, Wheatley, Sovacool and Sornette (2017) estimates that there is a 50% chance that a Fukushima event (or larger) occurs every 60-150 years across the entire *global* fleet of nuclear reactors. Germany had less than 4% of the world's nuclear reactors in 2011.

We can also examine how costly a single nuclear accident would have to be in order to rationalize the phase-out decision based on expected costs and benefits. Using the estimates from Wheatley, Sovacool and Sornette (2017), a Fukushima event (or larger) might plausibly occur in Germany over a twenty year period with a probability between 1 in 150 and 1 in 375.[32] This nuclear accident would have

[31] $\frac{\text{€160 billion}}{\text{€280 billion}} = 0.57$, which we approximate as 1 in 2, and $\frac{\text{€60 billion}}{\text{€640 billion}} = 0.094$, which we approximate as 1 in 10

[32] $(50\% \times 4\% \times \frac{20 \text{ years}}{60 \text{ years}}) = 0.0067 = 1$ in 150, and $(50\% \times 4\% \times \frac{20 \text{ years}}{150 \text{ years}}) = 0.0027 = 1$ in 375. This assumes reactors in Germany are just as prone to accidents as the average reactor in the world.

impose social costs between €9 trillion and €60 trillion for the expected benefits of the phase-out to be equal to its costs.[33] This is orders of magnitude greater than current estimates of the cost of the Fukushima accident. Both of these calculations, while simple in their approach, indicate that policymakers would have to exhibit an extremely high level of risk aversion in order the rationalize the phase-out decision.

However, key behavioural factors may affect decision-making in this setting. Prospect theory predicts that people will over-weight the risk associated with low-probability events, and tend to be risk averse when they face even a small chance of incurring a large loss (Kahneman and Tversky, 1979; Barberis, 2013). In our case, the choice is between a high probability of incurring a moderate loss due to increased electricity prices and worse air pollution versus a low probability of incurring a very large loss due to a nuclear accident. Here, loss aversion rationalizes a preference for avoiding the "low probability, high damage" option. Indeed, existing evidence suggests that people tend to greatly overestimate both the probability of a nuclear accident and the expected damages from such an event, including a number of studies focused on Germany (Slovic, Fischhoff and Lichtenstein, 1979; Slovic and Weber, 2002; Slovic, 2010).

The extreme level of risk aversion required to justify the phase-out decision also points to another behavioral explanation: salience. While the scientific literature on the harmful effects of air pollution is now definitive, there is still relatively limited public understanding of the scale of the adverse health consequences of local air pollution exposure. This might be due to the difficulty in attributing any single death entirely to air pollution exposure from a given power plant. Instead, air pollution concentration levels are the result of a wide range of different emitters and air

[33] $\frac{\text{€60 billion}}{0.0067}$ = €9 trillion, and $\frac{\text{€160 billion}}{0.0027}$ = €60 trillion.

pollution has a small but persistent effect on mortality risk. Similarly, the costs of climate change will primarily be borne by future generations, and linking a future climate event to the carbon emissions from a power plant smokestack today is even less straightforward. In contrast, a nuclear accident is a highly visible event that can be clearly linked back to a nuclear reactor.

Local air pollution emissions are almost certainly less salient than carbon emissions, particularly at the time when the phase-out decision was made. Germany has long had an ambitious set of policies to reduce carbon emissions as part of its *Energiewende* program. There is widespread support for these policies to tackle climate change, even when the upfront costs have been substantial. For example, renewable subsidies in Germany exceeded €25 billion per year in 2017, and the charges to fund these payments now make up about a quarter of the electricity price paid by residential households (BNetzA, 2020). The issue of local air pollution, on the other hand, has not received the same level of national attention. In fact, policymakers appear to have made no mention of air pollution at the time of the decision (BMWI, 2010; Ethics Commission, 2011). Subsequent studies of the impact of the phase-out have also focused exclusively on electricity prices and carbon emissions (Knopf et al., 2011; Traber and Kemfert, 2012; Knopf et al., 2014; Grossi, Heim and Waterson, 2017; Grossi et al., 2018).

The relative salience of the different impacts of the nuclear phase-out was further exacerbated in our setting by the Fukushima crisis in 2011. The German government clearly acknowledged this when making their decision, stating that "the risks of nuclear energy have not changed since Fukushima, but the perception of the risks has" (Ethics Commission, 2011). For example, Tanaka and Zabel (2018) demonstrated how the increased salience of nuclear accident risks due to the Fukushima incident affected house prices in the U.S. Interestingly, the study finds that the effects largely dissipated by two years after the Fukushima accident. This highlights an important challenge for policymakers - highly salient events can galvanize the public into supporting large-scale policy changes, but they can also make it harder to weigh the merits of those policy changes in the moment.

Regardless of the underlying causes, it is clear that Germans care deeply about climate change, yet they are distinctly anti-nuclear. Policymakers in Germany and around the world thus face a difficult tradeoff. On the one hand, there is a strong case to be made that nuclear power still has an important role to play in the shift away from carbon-intensive fossil fuels (IPCC, 2018; IEA, 2019). Many citizens are also willing to incur substantial costs upfront to reduce the risk of climate change. However, many of those same citizens have historically been more concerned about the risks of nuclear energy than about the health impacts of local air pollution from using fossil fuels. Clearly, it is vital that economists are able to credibly estimate and convey the costs and benefits of large policy changes. Beyond that, our findings underscore the importance of understanding the behavioral factors that can influence the way those costs and benefits are evaluated by policymakers and the public.

# References

**Abrell, Jan, Mirjam Kosch, and Sebastian Rausch.** 2019. "How Effective Was the UK Carbon Tax?-A Machine Learning Approach to Policy Evaluation." *A Machine Learning Approach to Policy Evaluation (April 15, 2019). CER-ETH–Center of Economic Research at ETH Zurich Working Paper*, 19: 317.

**Adler, David B, Akshaya Jha, and Edson Severnini.** 2020. "Considering the nuclear option: Hidden benefits and social costs of nuclear power in the US since 1970." *Resource and Energy Economics*, 59: 101127.

**Athey, Susan, Mohsen Bayati, Nikolay Doudchenko, Guido Imbens, and Khashayar Khosravi.** 2017. "Matrix Completion Methods for Causal Panel Data Models." *arXiv*, arXiv–1710.

**Barberis, Nicholas C.** 2013. "Thirty Years of Prospect Theory in Economics: A Review and Assessment." *Journal of Economic Perspectives*, 27(1): 173–196.

**Beelen, Rob, Ole Raaschou-Nielsen, Massimo Stafoggia, Zorana Jovanovic Andersen, Gudrun Weinmayr, Barbara Hoffmann, Kathrin Wolf, Evangelia Samoli, Paul Fischer, Mark Nieuwenhuijsen, et al.** 2014. "Effects of long-term exposure to air pollution on natural-cause mortality: an analysis of 22 European cohorts within the multicentre ESCAPE project." *The Lancet*, 383(9919): 785–795.

**BMWI.** 2010. "Monitoring Report 2020." BMWI Report.

**BMWi.** 2018. "Sixth Energy Transition Monitoring Report: The Energy of the Future." Federal Ministry of Economic Affairs and Energy (BMWi) Report.

**BNetzA.** 2020. "Monitoring Report 2020." Bundesnetzagentur Report.

**Breiman, Leo.** 2001. "Random Forests." *Machine Learning*, 45(1): 5–32.

**Burlig, Fiona, Christopher Knittel, David Rapson, Mar Reguant, and Catherine Wolfram.** 2020. "Machine learning from schools about energy efficiency." *Journal of the Association of Environmental and Resource Economists*, 7(6): 1181–1217.

**Carleton, Tamma, and Michael Greenstone.** 2021. "Updating the United States Government's Social Cost of Carbon."

**Carvalho, Carlos, Ricardo Masini, and Marcelo C Medeiros.** 2018. "Arco: an artificial counterfactual approach for high-dimensional panel time-series data." *Journal of econometrics*, 207(2): 352–380.

**Davis, Lucas, and Catherine Hausman.** 2016. "Market Impacts of a Nuclear Power Plant Closure." *American Economic Journal: Applied Economics*, 8(2): 92–122.

**Deryugina, Tatyana, Garth Heutel, Nolan H Miller, David Molitor, and Julian Reif.** 2019. "The mortality and medical costs of air pollution: Evidence from changes in wind direction." *American Economic Review*, 109(12): 4178–4219.

**Deschenes, Olivier, Michael Greenstone, and Joseph S. Shapiro.** 2017. "Defensive Investments and the Demand for Air Quality: Evidence from the NOx Budget Program." *American Economic Review*, 107(10): 2958–89.

**D'haeseleer, William.** 2013. "Synthesis on the Economics of Nuclear Energy." DG Energy Report.

**EEA.** 2014. "Costs of air pollution from European industrial facilities 2008–2012." European Environment Agency EEA Technical Report 20/2014.

**Egerer, Jonas.** 2016*a*. "Open Source Electricity Model for Germany (ELMOD-DE)." DIW Berlin, German Institute for Economic Research Data Documentation 83.

**Egerer, Jonas.** 2016*b*. "Open Source Electricity Model for Germany (ELMOD-DE)." DIW Berlin, German Institute for Economic Research Data Documentation 83.

**EIA.** 2019. "Generating Unit Annual Capital and Life Extension Costs Analysis." Energy Information Administration Report.

**Ethics Commission.** 2011. "Germany's energy transition - A collective project for the future." Ethics Commission for a Safe Energy Supply Report.

**Goebel, Jan, Christian Krekel, Tim Tiefenbach, and Nicolas R Ziebarth.** 2015. "How natural disasters can affect environmental concerns, risk aversion, and even politics: evidence from Fukushima and three European countries." *Journal of Population Economics*, 28(4): 1137–1180.

**Grossi, Luigi, Sven Heim, and Michael Waterson.** 2017. "The impact of the German response to the Fukushima earthquake." *Energy Economics*, 66: 450 – 465.

**Grossi, Luigi, Sven Heim, Kai Hüschelrath, and Michael Waterson.** 2018. "Electricity market integration and the impact of unilateral policy reforms." *Oxford Economic Papers*, 70(3): 799–820.

**He, Guojun, and Takanao Tanaka.** 2019. "Energy Saving Can Kill: Evidence from the Fukushima Nuclear Accident."

**Holland, Stephen P, Erin T Mansur, Nicholas Muller, and Andrew J Yates.** 2018. "Decompositions and Policy Consequences of an Extraordinary Decline in Air Pollution from Electricity Generation." National Bureau of Economic Research Working Paper 25339.

**IEA.** 2019. "World Energy Outlook 2019." International Energy Agency Report.

**IPCC.** 2018. "Global Warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty." Intergovernmental Panel on Climate Change Report.

**IRENA.** 2020. "Renewable Power Generation Costs in 2020." International Renewable Energy Agency Technical Report.

**Jacobs, David.** 2012. "The German Energiewende – History, Targets, Policies and Challenges." *Renewable Energy Law and Policy Review*, 3(4): 223–233.

**Jaramillo, Paulina, and Nicholas Muller.** 2016. "Air pollution emissions and damages from energy production in the U.S.: 2002–2011." *Energy Policy*, 90(C): 202–211.

**JECR.** 2019. "Follow up Report of Public Financial Burden of the Fukushima Nuclear Accident." Japan Center for Economic Research Report.

**Jones, Dave, Charles Moore, Will Richard, Rosa Gierens, Lauri Myllvirta, Sala Primc, Greg McNevin, Kathrin Gutmann, Anton Lazarus, Christian Schaible, and Joanna Flisowka.** 2018. "Last Gasp: the coal companies making Europe sick." Europe Beyond Coal.

**Kahneman, Daniel.** 2003. "Maps of Bounded Rationality: Psychology for Behavioral Economics." *The American Economic Review*, 93(5): 1449–1475.

**Kahneman, Daniel, and Amos Tversky.** 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47(2): 263–291.

**Keppler, Jan Horst.** 2012. "The economic costs of the nuclear phase-out in Germany."

**Knopf, Brigitte, Michael Pahle, Hendrik Kondziella, Fabian Joas, Ottmar Edenhofer, and Thomas Bruckner.** 2011. "Scenarios for Phasing Out Nuclear Energy in Germany." PIK Report.

**Knopf, Brigitte, Michael Pahle, Hendrik Kondziella, Fabian Joas, Ottmar Edenhofer, and Thomas Bruckner.** 2014. "Germany's Nuclear Phase-out: Sensitivities and Impacts on Electricity Prices and CO2 Emissions." *Economics of Energy & Environmental Policy*, 0(Number 1).

**Koch, Nicolas, Godefroy Grosjean, Sabine Fuss, and Ottmar Edenhofer.** 2016. "Politics matters: Regulatory events as catalysts for price formation under cap-and-trade." *Journal of Environmental Economics and Management*, 78: 121–139.

**Lazard.** 2019. "Levelized Cost of Energy Analysis - Version 13.0." Lazard.

**Markandya, Anil, and Paul Wilkinson.** 2007. "Electricity generation and health." *The lancet*, 370(9591): 979–990.

**Meinshausen, Nicolai.** 2006. "Quantile Regression Forests." *J. Mach. Learn. Res.*, 7: 983–999.

**Neidell, Matthew J, Shinsuke Uchida, and Marcella Veronesi.** 2019. "Be Cautious with the Precautionary Principle: Evidence from Fukushima Daiichi Nuclear Accident." National Bureau of Economic Research.

**NRC and NAS.** 2010. "Hidden Costs of Energy: Unpriced Consequences of Energy Production and Use." National Research Council (US). Committee on Health, Environmental, and Other External Costs and Benefits of Energy Production and Consumption. National Academies Press.

**O'Neill, Eoghan, and Melvyn Weeks.** 2019. "Causal Tree Estimation of Heterogeneous Household Response to Time-Of-Use Electricity Pricing Schemes."

**Open Power System Data.** 2018. "Data Package Conventional power plants."

**Severnini, Edson.** 2017. "Impacts of nuclear plant shutdown on coal-fired power generation and infant health in the Tennessee Valley in the 1980s." *Nature Energy*, 2(4): 17051.

**Slovic, Paul.** 2010. "The psychology of risks." *Saúde e Sociedade*, 19(4): 737–747.

**Slovic, Paul, and Elke U Weber.** 2002. "Perception of risk posed by extreme events." *Regulation of Toxic Substances and Hazardous Waste (2nd edition)(Applegate, Gabba, Laitos, and Sachs, Editors), Foundation Press, Forthcoming.*

**Slovic, Paul, Baruch Fischhoff, and Sarah. Lichtenstein.** 1979. "Rating the risks." *Environment*, 21(3): 14–39.

**Souza, Mateus.** 2020. "Predictive Counterfactuals for Treatment Effect Heterogeneity in Event Studies with Staggered Adoption." *Working Paper*.

**Tanaka, Shinsuke, and Jeffrey Zabel.** 2018. "Valuing nuclear energy risk: Evidence from the impact of the Fukushima crisis on U.S. house prices." *Journal of Environmental Economics and Management*, 88(C): 411–426.

**Traber, Thure, and Claudia Kemfert.** 2012. "German Nuclear Phase-out Policy: Effects on European Electricity Wholesale Prices, Emission Prices, Conventional Power Plant Investments and Eletricity Trade." DIW Berlin, German Institute for Economic Research Discussion Papers of DIW Berlin 1219.

**Varian, Hal R.** 2014. "Big data: New tricks for econometrics." *Journal of Economic Perspectives*, 28(2): 3–28.

**Varian, Hal R.** 2016. "Causal inference in economics and marketing." *Proceedings of the National Academy of Sciences*, 113(27): 7310–7315.

**Viscusi, W. Kip, and Clayton J. Masterman.** 2017. "Income Elasticities and Global Values of a Statistical Life." *Journal of Benefit-Cost Analysis*, 8(2): 226–250.

**Wheatley, Spencer, Benjamin Sovacool, and Didier Sornette.** 2017. "Of Disasters and Dragon Kings: A Statistical Analysis of Nuclear Power Incidents and Accidents." *Risk Analysis*, 37(1): 99–115.

# List of Figures



# List of Tables

# Appendices

# A Appendix Tables and Figures

Appendix Figure A.1 presents annual total electricity production in Germany by source as well as total imports and exports. This figure documents the precipitous drop in nuclear production following the 2011 closure of nine reactors as well as the rapid growth in production from wind and solar resources over our 2010-2019 sample period.

Appendix Figure A.2 presents a snapshot of the electricity supply curve based on our estimates of plants' short-run marginal costs in 2011. We assume that biomass, waste, hydroelectric, wind and solar resources have zero marginal operating cost. Marginal costs for fossil fuel plants are calculated as the sum of fuel costs and an assumed amount of variable operating and maintenance costs that differs by fuel type. Lastly, we assume that nuclear plants have a marginal operating cost of approximately €10/MWh based on prior research on Germany's power sector (Egerer, 2016*a*). This is confirmed by company reports from two European nuclear plant operators, EON and EDF, which also have marginal fuel costs of approximately €10/MWh.

Figure A.1: Electricity Supply by Source: 2010-2019

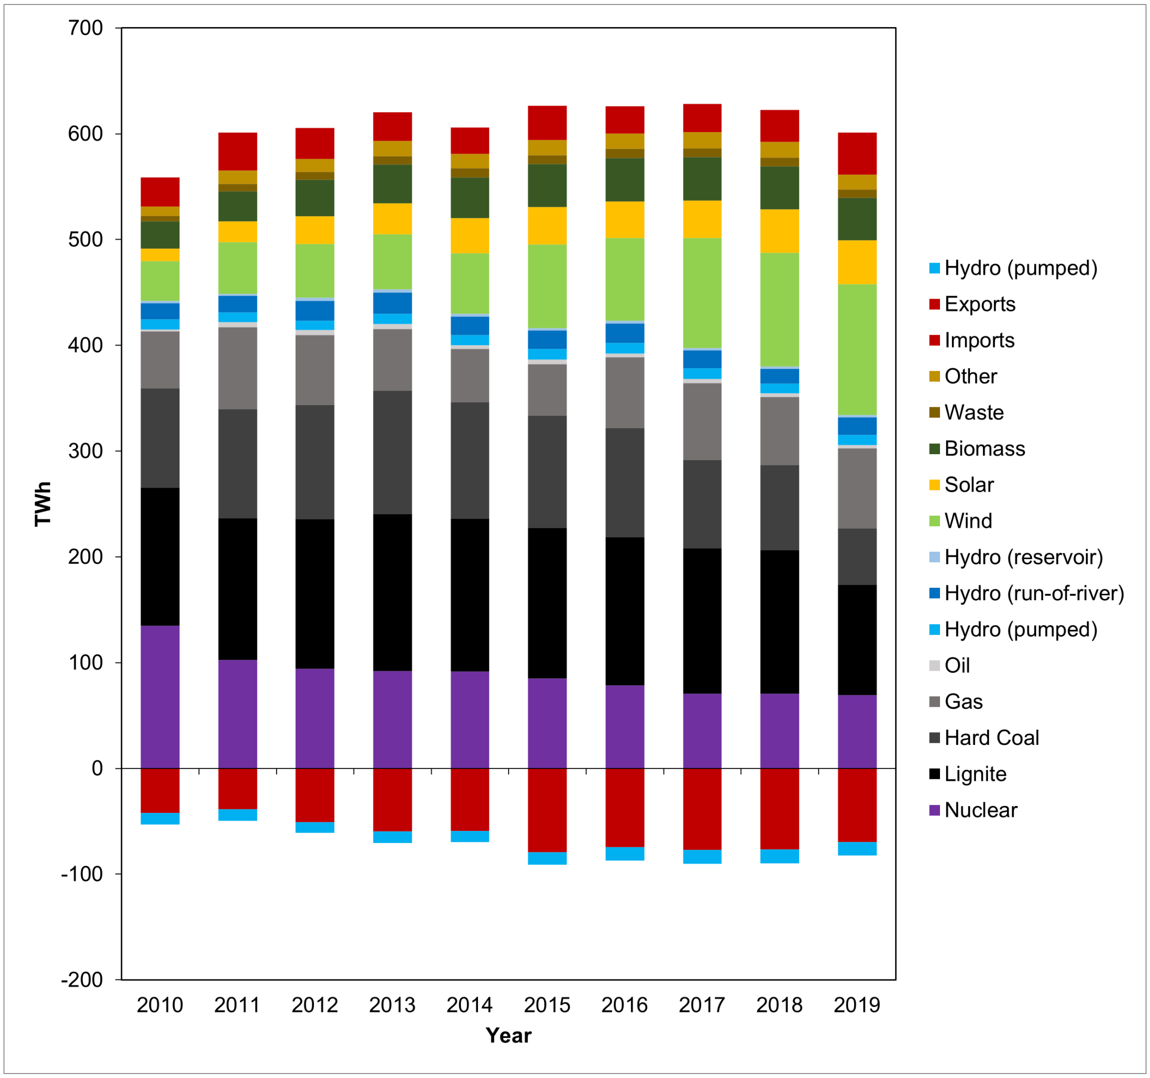


**Notes:** This figure plots the annual total quantity of electricity produced by different sources in Germany from 2010-2019. This includes the annual total quantity of electricity imports and exports for this same sample period. The data underlying this figure are from BNetzA Monitoring Reports (BNetzA, 2020).

Figure A.2: Marginal Cost Curve in 2011

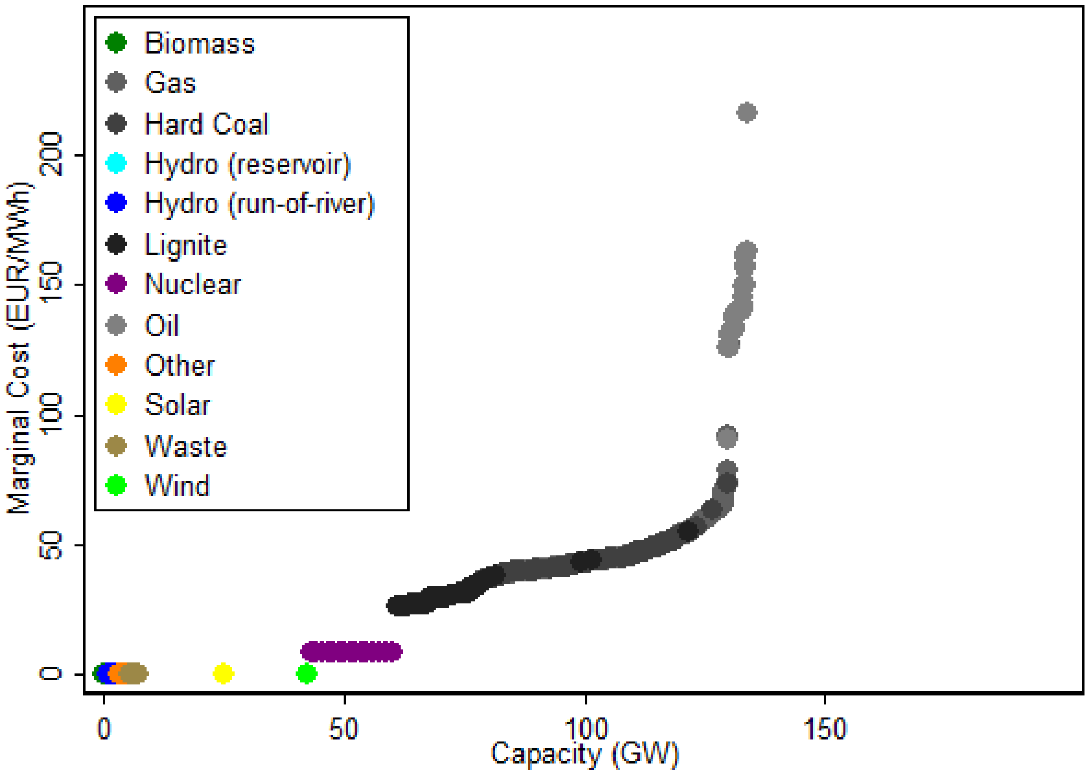


**Notes:** This figure plots the estimated marginal costs for power plants in Germany in 2011. Specifically, plants are ordered in terms of marginal cost to create an aggregate supply curve. For a given marginal cost $c$ (plotted on the y-axis), the x-axis provides the sum of the production capacity (in GW) over all plants with marginal cost less than or equal to $c$. For coal, gas and oil plants, marginal costs are calculated as the sum of fuel costs and an assumed variable operating and maintenance cost that differs by fuel type. Fuel costs are converted to euros per MWh using the plant's thermal efficiency (a measure of how well a plant converts units of input heat to units of electricity output). For this figure, we consider the fuel costs on February 1st, 2011. Nuclear plants are assigned a marginal cost of €10 per MWh. Hydro, wind and solar are assumed to have zero marginal costs. For simplicity, the small amount of remaining sources are also assigned a marginal cost of zero (i.e., biomass, waste and other). For ease of presentation, this figure does not show how electricity imports and exports factor into the aggregate supply curve; importantly, we account for imports and exports in our analysis.

# B Further Details on the Predictive Model

## B.1 Context

Studies of the electricity sector traditionally utilize an electricity dispatch model that combines engineering and economic modeling tools to simulate the operation of the power grid. These models must explicitly specify firm incentives (ex: whether/how firms exercise market power) as well as operational constraints such as transmission congestion and plants' start-up/ramping costs.

We opt to employ an empirical approach instead. Specifically, our approach seeks to recover how plants are dispatched based on a host of different variables pertaining to plant operations, demand, and electricity transmission. The primary benefit of this empirical approach is that it requires fewer assumptions regarding firm incentives or operational constraints. We allow the data to tell us how these factors impact plant operations.

That being said, this empirical approach has limitations as well. First, we can only examine scenarios that are sufficiently similar to observed outcomes. This is why other empirical models of wholesale electricity markets tend to focus either on ex-post policy assessments or identifying how marginal changes in electricity demand impact plant operations. Indeed, our paper focuses on an ex-post evaluation of the nuclear phase-out in Germany on aggregate market outcomes.

We want to highlight that empirical approaches such as ours can struggle with counterfactual scenarios far out of the span of the data used to estimate the model. As such, our empirical modeling should be seen as a complement rather than a substitute for more explicit simulation modeling of electricity markets.

## B.2 Construction of Key Variables

The most important independent variables for our analysis are net load, marginal cost and available capacity.

### B.2.1 Net Electricity Demand (or Net Load)

Net electricity demand, or net load, is defined as total electricity demand minus production from low marginal cost or non-dispatchable sources. Specifically, we subtract production from renewables (wind, solar, hydro, biomass, waste) and nuclear. This net load variable thus measures the amount of production required by "dispatchable" (typically fossil-fuel-fired) sources.[34]

Key to setting up the net demand variable in this manner is the assumption that flexible production sources will respond to marginal changes in net load in much the same way as they would to marginal changes in the component parts of net load. To provide evidence in favor of this assumption, we conduct a straightforward regression analysis based on the hourly data used to construct the net load variable. The results are shown in Appendix Table B.1.

Specifically, Appendix Table B.1 shows how "dispatchable" output responds to a 1 MWh change in net electricity demand, or its constituent parts. The dependent variable is hourly total "dispatchable" generation. This is the sum of hourly production from all fossil plants and border points. The independent variable in columns (1) and (3) is net electricity demand as defined and used in our analysis. The in-

[34]We also considered specifications that included lags and leads of net load to capture the fact that many power plants have dynamic production constraints (e.g. the speed at which they can "ramp up" their output, or the minimum amount of time they have to be offline before they can restart).

Table B.1: Response of Dispatchable Generation to Electricity Demand

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| Net Demand | 0.893*** | | 0.893*** | |
| | (0.017) | | (0.020) | |
| Demand | | 0.871*** | | 0.876*** |
| | | (0.010) | | (0.016) |
| Nuclear Output | | -0.650** | | -0.805*** |
| | | (0.217) | | (0.192) |
| Wind Output | | -0.908*** | | -0.940*** |
| | | (0.039) | | (0.023) |
| Solar Output | | -0.859*** | | -0.787*** |
| | | (0.051) | | (0.052) |
| Year FE | | | Y | Y |
| Month of Year FE | | | Y | Y |
| Hour of the Day FE | | | Y | Y |
| $R^2$ | 0.869 | 0.872 | 0.896 | 0.898 |
| Mean of Dep. Var. | 31,222.3 | 31,222.3 | 31,222.3 | 31,222.3 |
| Number of Obs. | 87,648 | 87,648 | 87,648 | 87,648 |

**Notes:** This table shows the results of regressing hourly total "dispatchable" generation on hourly net electricity demand (columns (1) and (3)), or its constituent parts (columns (2) and (4)). Columns (1) and (2) do not include any fixed effects while columns (3) and (4) include year, month of year, and hour of the day fixed effects. Standard errors are clustered by year.

dependent variables in columns (2) and (4) are the constituent parts that make up net electricity demand - namely total demand, nuclear production, wind production and solar production. Columns (1) and (2) do not include any fixed effects while columns (3) and (4) include fixed effects for year, month of year, and hour of the day. Standard errors are clustered by year for all specifications.

The coefficients of this regression tell us how “dispatchable” output responds to a 1 MWh change in net electricity demand, or to a 1 MWh change in the constituent parts of net electricity demand. The coefficients for net demand and demand are similar and are both close to one. The coefficients for nondispatchable output (nuclear, wind, and solar) are close to -1, suggesting that increases in nondispatchable output replace production from dispatchable sources roughly one for one. This demonstrates that, on average, dispatchable sources will respond to variation in net demand in much the same way as they would to variation in the constituent parts of net demand. We view this as a validation of our decision to focus on net demand rather than include each of its constituent parts in the model.

Appendix Table B.2 presents results from regressing hourly total output by source type on hourly total electricity demand, including year, month of year, and hour of the day fixed effects. The first, second, and third columns of this panel consider hourly total output from dispatchable sources (coal + oil + gas + net exports), wind sources, and solar sources respectively. Standard errors are clustered by year.

After controlling for year fixed effects, month of year fixed effects, and hour of the day fixed effects, neither the correlation between total demand and wind generation nor the correlation between total demand and solar generation are statistically significant. This suggests that generation from wind and solar resources is not responsive to the level of demand. In contrast, we cannot reject that a 1 MWh increase in total

Table B.2: Response of Dispatchable and Renewable Generation to Electricity Demand

| | (1) | (2) | (3) |
|---|---|---|---|
| Demand | 0.889*** | -0.034 | 0.000 |
| | (0.025) | (0.025) | (0.013) |
| Year FE | Y | Y | Y |
| Month of Year FE | Y | Y | Y |
| Hour of the Day FE | Y | Y | Y |
| $R^2$ | 0.606 | 0.267 | 0.647 |
| Mean of Dep. Var. | 31,222.3 | 8428.0 | 3.504.9 |
| Number of Obs. | 87,648 | 87,648 | 87,648 |

**Notes:** This table presents results from regressing hourly total output by source type on hourly total electricity demand, including year, month of year, and hour of the day fixed effects. The first, second, and third columns of this panel consider hourly total output from dispatchable sources (coal + oil + gas + net exports), wind sources, and solar sources respectively. Standard errors are clustered by year.

demand corresponds to a 1 MWh increase in generation from dispatchable sources, suggesting that the lack of correlation for wind and solar resources is not an artifact of the specification considered.

### B.2.2 Marginal Cost

A plant decides whether to produce primarily based on whether its marginal cost is less than the electricity price it will be paid for its output. In electricity markets such as Germany's, the electricity price is typically set by the highest marginal cost plant necessary to meet demand. Consequently, we first construct estimates of each plant's marginal cost over time. We then estimate the marginal cost of the clearing plant: the last fossil plant (or border point) necessary to meet net load in a given hour. Finally, we construct a "standardized" marginal cost for each plant as the plant's marginal cost minus the marginal cost of the clearing plant for that hour. Plants typically produce if this standardized marginal cost is negative, and don't produce if this standardized marginal cost is positive.

In standardizing each plant's marginal cost, one may be concerned that the marginal cost of the marginal unit is endogenous to the policy. Namely, in the absence of the phase-out, net electricity demand will shift to the left, lowering the marginal cost of the marginal unit. Thus the phase-out impacts the marginal cost of the marginal unit.

However, as illustrated in Appendix Figure B.1, it is important to emphasize that we calculate the marginal cost of the marginal unit in each hour-of-sample separately for the scenarios with versus without the phase-out. Namely, we decrease the net demand to be served by dispatchable units in the scenario without the phase-out, stack units from lowest marginal cost to highest marginal cost, and then find the marginal unit. In doing so, our measure of the marginal cost of the marginal unit for both the scenarios with versus without the phase-out are based on exogenous factors such as total demand, production from nondispatchable units such as renewables, etc.

Standardizing each plant's marginal cost by the marginal cost of the marginal unit, adjusting this standardization depending on whether we're considering the scenario with versus without the phase-out, has the added benefit of preventing the predictions from being adversely affected by secular trends in fuel prices from before the 2015-2019 sample period used to train the model.

### B.2.3 Available Capacity

Where the "marginal cost" variable captures the position of a plant in the supply curve in terms of price, the "available capacity" variable captures the position of a plant in the supply curve in terms of quantity. For each plant, we calculate the total amount of capacity from other fossil plants (or border points) with a lower marginal cost. Our "available capacity" variable is then calculated as the total amount of capacity with a lower marginal cost than the plant minus net load for that hour. Once again, plants with negative available capacity are likely to produce, while plants with positive available capacity are unlikely to produce.

### B.2.4 Role in Simulating the Phase-Out Policy

The machine learning application we use is designed to predict how dispatchable flexible sources such as fossil-fuel plants and border flows increase or decrease their output in order to meet the residual demand left after accounting for output from renewables and nuclear sources. Net load, the relative marginal cost of each plant, and the amount of available capacity are key predictors in the analysis not only because they play a significant role in explaining plant operating decisions, but also because they are the variables we modify in order to construct the counterfactual scenario.

For the scenario with the phase-out, the net load variable is the observed net load given the phase-out decision as shown in Appendix Figure B.1a. For the counterfactual scenario without the phase-out, nuclear production would have been higher and so net load would have been lower, as shown in Appendix Figure B.1b. This reduction in net load also changes the marginal cost and available capacity variables.

Specifically, if net load is lower, the marginal cost of the clearing plant would also be lower. Moreover, the amount of capacity below net load is also lower for lower values of net load. This is illustrated in Appendix Figures B.1c and B.1d.

### B.2.5 Variable Scaling and Importance

Appendix Figure B.2a illustrates the relative importance of each of our covariates. As expected, net demand, marginal cost and available capacity are all particularly important covariates. However, it is noteworthy that the two most important covariates are the type of source (i.e., lignite, hard coal, gas, oil or border point) and whether the fossil-fuel-fired plant is combined-heat-and-power. This reflects the fact that different types of electricity generators face different operational constraints. For example, many natural gas plants in Germany are combined-heat-and-power. As such, whilst they may have higher marginal costs than coal plants, they receive revenues both for their electricity output and from providing heating services. Consequently, combined-heat-and-power plants operate more frequently than would be suggested by simply comparing their marginal cost to electricity prices.

When making out-of-sample predictions using a predictive model such as this, it is important to ensure that the training data-set provides sufficient support across the predictor variables. This is because our algorithm is ill-suited to extrapolate beyond the economic conditions seen in the training data. Appendix Figure B.2 documents that there is substantial overlap in support of the key variables across the training, counterfactual, and post-2015 "missing" data. This provides evidence that assessing the impacts of the nuclear phase-out is, for most hours-of-sample, an interpolation exercise rather than an extrapolation exercise.

Figure B.1: Net Demand and Scenario Implementation

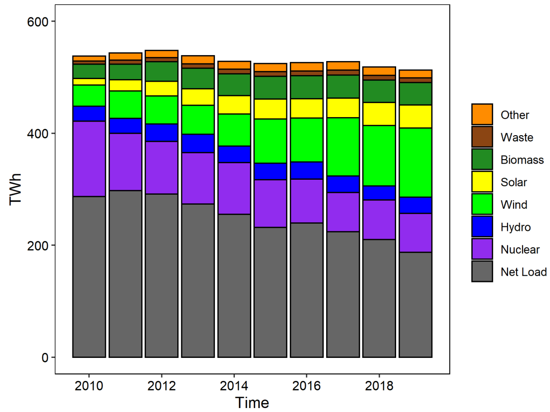


(a) Net Demand (With Phase-Out)

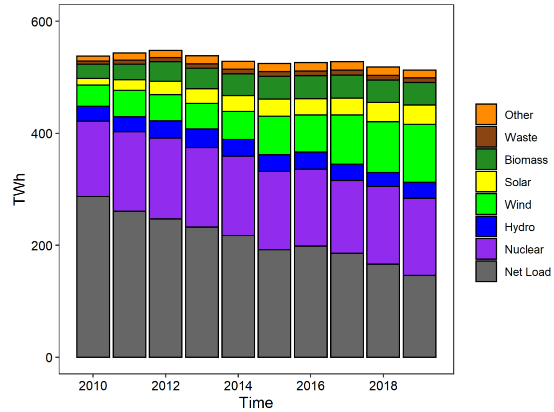


(b) Net Demand (Without Phase-Out)

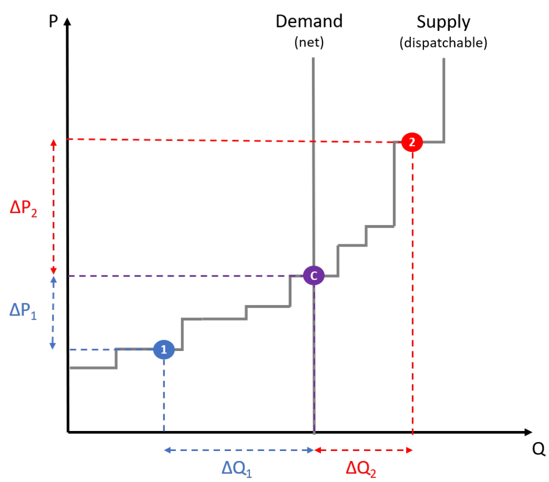


(c) Clearing Example (With Phase-Out)

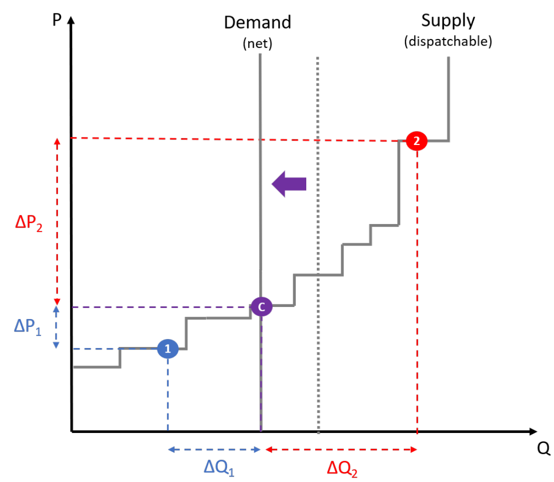


(d) Clearing Example (Without Phase-Out)

**Notes:** This figure illustrates the role of the net electricity demand variable in the analysis. Net demand is defined as total electricity demand minus production from low marginal cost or non-dispatchable sources. Specifically, we subtract production from renewables (wind, solar, hydro etc.) and nuclear. Panels (a) and (b) show the level of net demand both with and without the phase-out respectively. Comparing panel (a) to panel (b) shows that more nuclear production leads to less net demand to be satisfied in the scenario without the nuclear phase-out. Panels (c) and (d) provide an illustration of how changing net demand impacts the estimation process. This happens because altering net demand alters the position where net demand intersects with the supply curve of dispatchable capacity. The point where net demand intersects with the supply curve of dispatchable capacity indicates the plant that is "on-the-margin" (purple). Altering the clearing plant affects the relative marginal cost ($\Delta P$) and available capacity ($\Delta Q$) values for all dispatchable supply. These two variables are illustrated for a high marginal cost plant (red) and a low marginal cost plant (blue).

Figure B.2: Machine Learning Model Diagnostics

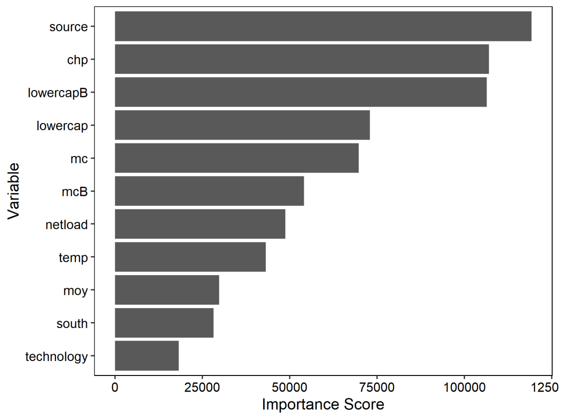


(a) Variable importance scores

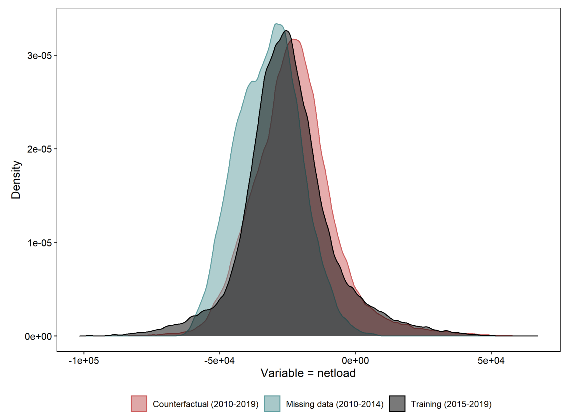


(b) Support of Net Demand

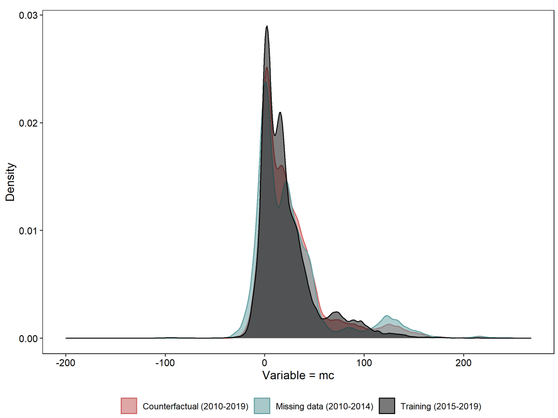


(c) Support of Marginal Cost

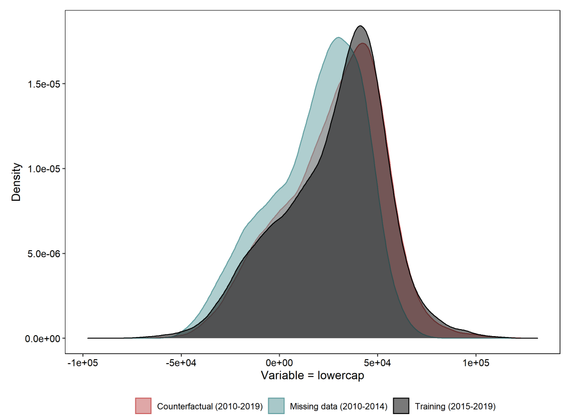


(d) Support of Available Capacity

**Notes:** This figure illustrates a range of key model diagnostics. Panel (a) shows the importance scores for each of the variables included in the estimation. Importance scores indicate the relative importance of each variable in predicting the outcome of interest. The abbreviated names in the figure are as follows: source = source type (e.g., lignite, hard coal, gas, oil or border); mc = marginal cost; mcB = marginal cost including border capacity; lowercap = available capacity; lowercapB = available capacity including border capacity; chp = combined-heat-and-power capability; technology = technology type (e.g., steam turbine, combined cycle); temp = local temperature; south = indicator for located in the south; moy = month-of-year; dow = day-of-week; hod = hour-of-day; netload = electricity load minus production from wind, solar, hydro and nuclear sources; netloadX = lag of net demand by X; netload_X = lead of net demand by X hours. Panels (b-d) show the support of three key variables: net demand, marginal cost and available capacity. The grey area shows the distribution of observations in the 2015-2019 training data-set (i.e., where we have hourly, plant-level production data). The blue area is for the missing 2010-2014 data (i.e.: where we only have hourly data on production by fuel type). The red area is for the counterfactual scenario (i.e., without the nuclear phase-out) across the full 2010-2019 analysis period.

Rescaling certain variables can also help to ensure that our out-of-sample predictions are not extrapolating too far outside the support of the training data.[35] The three main variables we use to approximate the interaction between supply and demand are net load, marginal costs, and the amount of available capacity. The counterfactual no-phase-out scenario contains some periods where these variables fall outside the range in the training dataset. This is potentially a concern because random forests are susceptible to bias at the edges of the feature space. For example, we might be concerned that, during periods of peak net demand, our model would artificially understate the predicted output from fossil plants. A similar problem applies during periods of extremely low net demand.

Fortunately, this problem largely affects only the subset of plants that are close to the margin during periods of extremely high or low net demand. Inframarginal plants will already be operating at their maximum and so will not be subject to this censoring effect. More importantly, these kinds of periods where we move outside of the support of the feature space in our training set are very limited. There is such substantial variation in electricity demand, production from renewables and marginal costs that there is considerable overlap in the support of these variables in our training dataset and any other periods or scenarios we are applying our model to. This includes the counterfactual no-phase-out scenario over the full 2010-2019 period, as well as the earlier 2010-2014 period where we do not have hourly plant-level data on output.

Appendix Figures B.2b, B.2c and B.2d provide a clear visual illustration of this

[35]For example, we rescale the marginal cost of each plant by the marginal cost of the marginal plant needed to clear the market. Even if fuel costs doubled from 2010-2019, for example, the rescaling would ensure that the rescaled marginal costs fed into our algorithm stay within a reasonable range over our sample period.

Table B.3: Common Support of Key Variables

| Variable | Scenario | Training Range | Training 95% | Training 80% | Overlap Coef |
|---|---|---|---|---|---|
| Net Demand | Missing data (2010-2014) | 1.00 | 0.99 | 0.78 | 0.79 |
| Net Demand | Counterfactual (2010-2019) | 1.00 | 0.97 | 0.81 | 0.90 |
| Marginal Cost | Missing data (2010-2014) | 1.00 | 0.87 | 0.72 | 0.79 |
| Marginal Cost | Counterfactual (2010-2019) | 1.00 | 0.92 | 0.78 | 0.90 |
| Available Capacity | Missing data (2010-2014) | 1.00 | 0.96 | 0.84 | 0.82 |
| Available Capacity | Counterfactual (2010-2019) | 1.00 | 0.96 | 0.81 | 0.97 |

**Notes:** This table shows a range of summary statistics to illustrate the shared support between the periods we make predictions for and the underlying training dataset we use to estimate our model. We apply our model to predict the missing 2010-2014 data (i.e., where we don't have hourly plant-level data) and the full 2010-2019 counterfactual scenario (i.e., without the nuclear phase-out). Column 3 shows the extent to which the range of the feature of interest is contained by the range of the same feature in the training set. Columns 4 and 5 are the same but use the central 95% and 80% of the training set. Lastly column 6 provides the overlap coefficient between the distribution of the feature of interest and the same feature in the training set.

overlap. Appendix Table B.3 provides summary statistics that make clear the extent to which the periods and scenarios we make predictions for lie within the feature space in our training set.

## B.3 Additional Model Performance Details

This subsection presents a number of further details on the model performance. We estimate a quantile random forest model. We tuned the main hyperparameters to ascertain the sensitivity of our results to these assumptions. The range of hyperparameters examined were combined into a grid and then fitted using five-fold cross-validation. For the number of trees, we examined 25, 50, 100, and 500, with 100 being our final choice. For the number of variables available for splitting at each tree node, we found the default value was sufficient (i.e., the (rounded down) square root of the number of variables). Using parameters at half this and twice this level had minimal impacts on accuracy. Lastly, for the minimum node size, again we found the default value of 5 was sufficient, although we did explore larger values (i.e., 10 and 50).

The results shown earlier in Figure 2 provided some insight into the performance of the random forest model. This was calculated using five-fold cross-validation to get reliable measures of out-of-sample performance. The folds were created randomly. A more stringent test of model performance would be to stratify the data along important dimensions and then construct the folds from these blocks. This better captures the challenge the model must overcome in predicting out-of-sample.

To illustrate this, we conduct an additional piece of analysis to measure model performance, where the folds are constructed based on distinct blocks of plant source types and years. The training dataset is divided up into blocks based on the five years (2015-2019) and the five source types (lignite, hard coal, natural gas, oil and border) in the training data-set. The resulting 25 blocks of the training dataset are randomly split across the five folds. The same analysis as before is then conducted whereby the model is fit on four of the folds and used to predict the fifth, before repeating the process until all folds have each been predicted for out-of-sample.

The results of this can be seen in Appendix Figure B.3. Using this alternative approach for out-of-sample cross-validation, the predictive performance of the model is largely unchanged from the values seen earlier in Figure 2.

## B.4 Matching Supply and Demand

A key reason for our use of quantile random forests is the fact that an empirical approach such as the one used in this paper does not incorporate the constraint

Figure B.3: Machine Learning Model Performance: Plant-Level Electricity Production

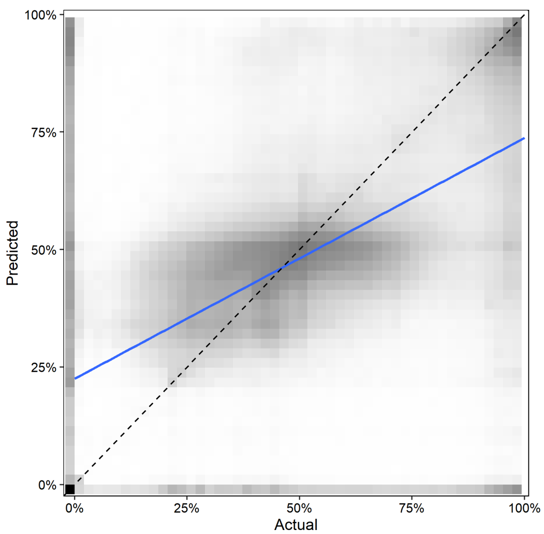


(a) Hourly Data

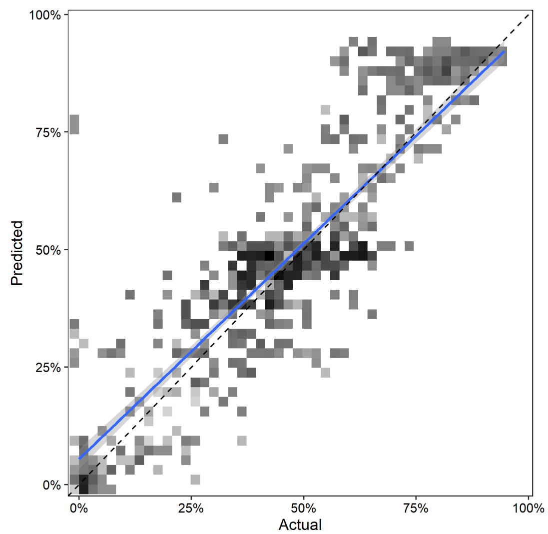


(b) Annual Data

**Notes:** This figure illustrates the accuracy of the plant-level predictions from the machine learning model presented in Section 5. The model predicts the operating rate of each power plant in each hour, where a value of 0% means that the plant is offline and a value of 100% means that the plant is running at maximum capacity. Values on the 45 degree line indicate perfect accuracy, and we summarize this both visually and by computing measures of Mean Squared Error and $R^2$. We compute these metrics using out-of-sample five-fold cross-validation. Darker areas indicate higher numbers of plant-hour (or plant-year) observations. Each pixel represents the predicted versus actual operating rate in increments of 2%. Panel (a) shows prediction accuracy at an hourly timescale. The MSE and $R^2$ are 0.10 and 0.48, respectively. Panel (b) shows prediction accuracy after taking annual averages of our hourly predictions. The number of observations is 846, and the MSE and $R^2$ are 0.02 and 0.85, respectively.

that total supply of electricity should be equal to the total demand for electricity. Ideally, we would expect the median or mean predictions from our machine learning approach to predict aggregate generation that is fairly close to the total demand that must be served.

In our case, the median predictions across all plants and hours in our analysis predicts total generation that is 9% higher than the total net electricity demand that needs to be met. In the no-phase-out scenario, our predictions are 27% above the expected level of total net electricity demand that needs to be met. To remedy this, we utilize the information our quantile regression model provides us on the full conditional distribution of potential changes to output. Our approach involves the following steps.

First, when making predictions of operating rates for each plant in each hour, we don't just make these predictions for the median. Instead, we generate predicted hourly plant operating rates for the 25th, 40th, 50th, 60th and 75th percentiles. In principle, we could generate predictions for many more percentiles. However, this becomes prohibitively time consuming, and so we focus on these five percentiles.

Second, we assign a set of weights to these five percentiles in order to recover the desired level of predicted total electricity production. To do this, we take the predictions for each of the five percentiles and sum the resulting values over all plants and hours in each month-of-sample to get monthly totals for predicted electricity production. We then compare the total monthly predicted electricity production for each percentile with the true monthly totals for net electricity demand that need to be met. Our goal is to find the percentile totals that lie on either side of the true total. So, if the 50th percentile predictions are 9% above the observed net demand that needs to be served and the 40th percentile predictions are 5% below, then our desired percentile would lie somewhere between these two.

Fourth, when we have identified the two relevant percentiles for a given month, we calculate the weight to assign to each of the two percentiles to produce a weighted average prediction that best matches supply and demand on a monthly level. In the above example, we would want to assign weights to the 40th and 50th percentiles of 0.64 and 0.36 respectively. As noted above, a more precise way to achieve the same result would be to make predictions for a larger number of percentiles such that we can simply go directly to the one that is closest to the true value. Again though, the main barrier to doing this is that making predictions for each plant-hour for each percentile is ultimately very time-consuming.

Appendix Figure B.4 shows the average weights assigned to each of the percentiles during our analysis. The panel columns capture the weights assigned in the scenario with the phase-out ("Actual") and without the phase-out ("Counterfactual"). The panel rows refer to our three renewables investment scenarios. In general, it is clear that the median predictions already perform relatively well. The tendency for our estimation to systematically overestimate the required level of supply is reflected in the largest weights being assigned to the predictions from the 40th and 50th percentiles.

A potential concern is that the percentile chosen for our analysis may depend on the scenario and time period under consideration, and that this could play an important role in driving our final results. For instance, the average percentile of the predicted output used in the no-phase-out case for the low, base, and high renewables scenarios (41.3, 42.5, and 43.7) is lower than the one used in the phase-out case (46).

To see if this presents a problem for our results, we first checked how the per-

Figure B.4: Quantile Weights for Supply/Demand Matching

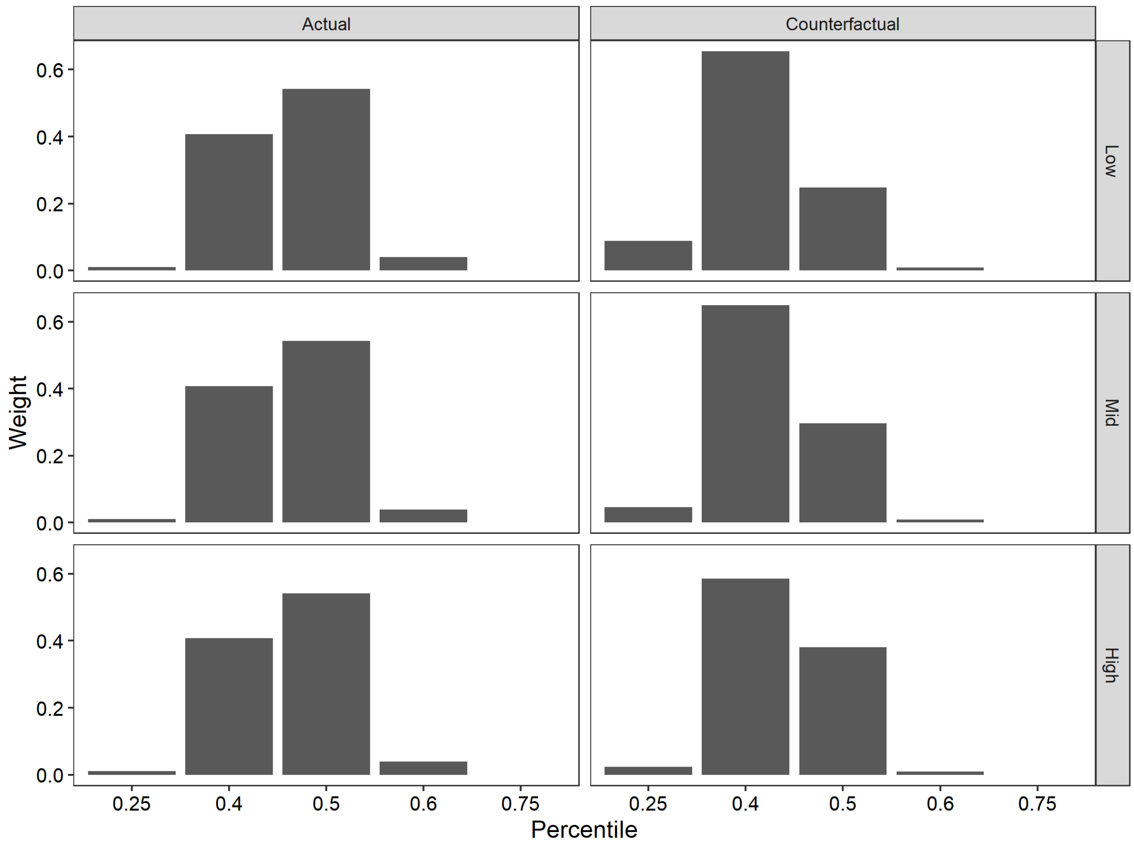


**Notes**: This figure plots the average weights applied to the five quantiles for which predictions are generated from the random forest model. The panel columns capture the weights assigned in the scenario with the phase-out ("Actual") and without the phase-out ("Counterfactual"). The panel rows refer to our three renewables investment scenarios.

centile chosen to match supply and demand varied with important covariates in the model. Here, we find that our model predictions do appear to perform worse (i.e. overprediction is largest) in the later years when net demand is very low. This is particularly the case in the "low" renewables scenario. A key driver of this is that there are not many observations in the training dataset with net demand as low as the net demand in later years in the low renewables counterfactual.

To explore the sensitivity of our results to this overprediction problem, we considered specifications in which we weight observations when estimating the model. The aim is to apply larger weights to the observations in our training dataset where our model currently struggles (e.g., where net demand is particularly low). We construct these weights focusing on the three key variables in our analysis: net demand, marginal cost, and available capacity. Our approach upweights observations in the training dataset that are more like those in the full prediction sample, and downweights observations that are less like those in the full prediction sample.[36]

The weighted estimation approach did slightly reduce the observed tendency for over-prediction, but did not radically alter it. The lack of any substantial change likely reflects the fact that, as documented in Appendix Section B.2.5, our training dataset is already fairly representative of the full range of time periods and renewable scenarios we are making predictions for. Panels A and B in Appendix Table B.4 below present estimates of the social cost of the phase-out based on the unweighted

[36]To do this, we calculate the deciles of each of the three variables using the training dataset. Interacting these three sets of deciles allows us to segment the support of the dataset into 1,000 cells. We then calculate the fraction of the overall 2010-2019 sample that falls into each of these cells as well as the corresponding fractions for the sample implied by each of our three counterfactual no-phase-out scenarios (i.e., all hours and all plants for all three renewable scenarios over 2010-2019). Then, the weight for each cell is constructed by dividing the fraction of observations in the training dataset that falls into the cell by the fraction of observations in the relevant prediction sample that falls into the cell. We also tried a weighting approach based on estimating propensity scores using the full set of covariates. This produced qualitatively similar results.

model (i.e., the baseline estimates from Table 5) and the weighted model respectively. For each model, we present estimates for the six different scenarios defined by assumptions on the value of statistical life and the percentage of phase-out-induced deployment of renewables. Comparing Panels A and B in Appendix Table B.4 for all six scenarios, the estimates of the social cost of the phase-out calculated using the predictions of output from the weighted machine learning model are very similar to those in our main specification.

As a second approach to tackling this issue, we demonstrate that our estimates of the social cost of the phase-out remain largely unchanged if we focus on the subset of months-of-sample where the percentile of predictions used in the factual scenario are very similar to the percentile of predictions used in the counterfactual no-phase-out scenario. As noted earlier, the average percentile of the predicted output used in the no-phase-out case for the low, base, and high renewables scenarios (41.3, 42.5, and 43.7) is lower than the one used in the phase-out case (46). The difference here is a gap of between 2.5 and 4.7 percentile points. If we focus on the 25% most similar months-of-sample where this difference is smallest, the average percentile of the predicted output used is close to 44 in both the phase-out and no-phase out cases across all three renewable scenarios. The difference between the percentiles used in the phase-out versus no-phase-out cases is just 0.1 to 0.4 percentile points. During these months, the level of overprediction is very similar in the phase-out and no phase-out scenarios, and so any adjustments needed to match supply and demand are also very similar.

The results in Panel C of Appendix Table B.4 demonstrate that our estimates of the social cost of the phase-out for the low, middle, and high renewables scenarios remains broadly unchanged when focusing on months where the level of adjustment

Table B.4: Estimated Annual Impact of the Nuclear Phase-Out on Private and External Costs: Sensitivity Analyses

*Panel A: Baseline Estimates*

| | **Average with Phase-Out (1)** | **Average without Phase-Out (2)** | **Change (3)** | **Change (%) (4)** |
|---|---|---|---|---|
| Base VSL, Low Renewables | 48.27 | 44.82 | 3.45 | 7.7 |
| **Base VSL, Base Renewables** | **48.27** | **44.95** | **3.32** | **7.4** |
| Base VSL, High Renewables | 48.27 | 45.05 | 3.21 | 7.1 |
| High VSL, Low Renewables | 76.14 | 68.04 | 8.10 | 11.9 |
| High VSL, Base Renewables | 76.14 | 69.57 | 6.57 | 9.4 |
| High VSL, High Renewables | 76.14 | 71.01 | 5.13 | 7.2 |

*Panel B: Estimates from Weighted Machine Learning Model*

| | **Average with Phase-Out (1)** | **Average without Phase-Out (2)** | **Change (3)** | **Change (%) (4)** |
|---|---|---|---|---|
| Base VSL, Low Renewables | 48.71 | 45.08 | 3.62 | 8.03 |
| **Base VSL, Base Renewables** | **48.65** | **45.24** | **3.40** | **7.52** |
| Base VSL, High Renewables | 48.59 | 45.40 | 3.18 | 7.0 |
| High VSL, Low Renewables | 77.44 | 68.94 | 8.49 | 12.31 |
| High VSL, Base Renewables | 77.29 | 70.46 | 6.83 | 9.69 |
| High VSL, High Renewables | 77.13 | 71.96 | 5.17 | 7.18 |

*Panel C: Estimates from "25% Most Similar Months"*

| | **Average with Phase-Out (1)** | **Average without Phase-Out (2)** | **Change (3)** | **Change (%) (4)** |
|---|---|---|---|---|
| Base VSL, Low Renewables | 46.99 | 43.89 | 3.10 | 7.1 |
| **Base VSL, Base Renewables** | **45.13** | **42.34** | **2.79** | **6.6** |
| Base VSL, High Renewables | 44.89 | 42.44 | 2.45 | 5.8 |
| High VSL, Low Renewables | 75.79 | 67.81 | 7.98 | 11.8 |
| High VSL, Base Renewables | 72.59 | 66.71 | 5.88 | 8.8 |
| High VSL, High Renewables | 68.63 | 65.00 | 3.63 | 5.6 |

**Notes:** This table reports estimates of the social cost of the nuclear phase-out for the six scenarios defined by low versus base assumptions on the value of statistical life (i.e, €2.56 million versus €6.7 million) and low, base, and high assumptions on phase-out-induced investment in renewables (see Section 5.4 for more details). In Panel A, we reproduce the baseline estimates from the bottom panel of Table 5 in the main text. In Panel B, we present results based on the weighted quantile random forest model. These weights are designed to make the training dataset more representative of the wider set of scenarios and time periods we make predictions for. In Panel C, we focus on months-of-sample where the predictions from our model require similar adjustments to match supply and demand in both the actual and counterfactual cases (the months-of-sample in bottom 25% of the absolute difference between the percentiles of predicted output used in the factual scenario with the phase-out versus the counterfactual scenario without the phase-out).

needed to match supply and demand is highly similar across both the actual and counterfactual cases (i.e., the "25% Most Similar Months" estimates). Our main estimates of the annual average increase in social costs due to the phase-out range from 3.21-8.10 billion euros, with a central estimate of 3.32 billion euros. The corresponding range of social costs when focusing on the subsample with highly comparable levels of prediction adjustment is 2.45-7.98 billion euros, with a central estimate of 2.79 billion euros. We view this evidence as indicative that the qualitative and quantitative conclusions drawn from our analysis of the nuclear phase-out are not driven by differences in the extent of overprediction for our model in the phase-out versus no-phase-out scenarios.

## B.5 Additional Model Results

Appendix Figure B.5 shows the median model predictions for how the nuclear phase-out impacted aggregate plant-level electricity production in Germany. As expected, points on this figure tend to lie above the horizontal axis; the nuclear phase-out reduced nuclear generation, with fossil-fuel-fired production filling the gap. The largest response to the phase-out comes from the hard coal plants.

Appendix Figure B.6 illustrates which plants and border points increased production to meet the reductions in nuclear output due to the phase-out. Most of the fossil-fuel generation comes from the industrial regions in the west and south of the country. Changes to net imports come primarily at the borders between Germany and Denmark, France and the Czech Republic.

Figure B.5: Plant-level Changes in Production due to the Phase-Out

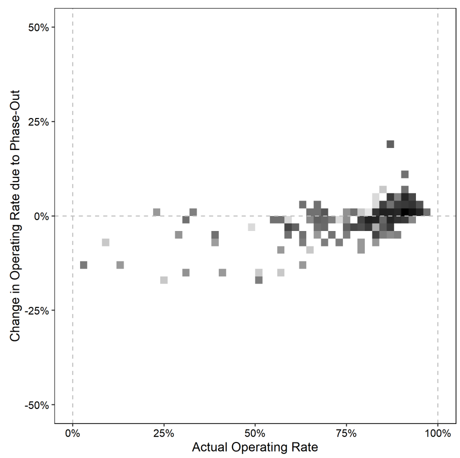


(a) Lignite

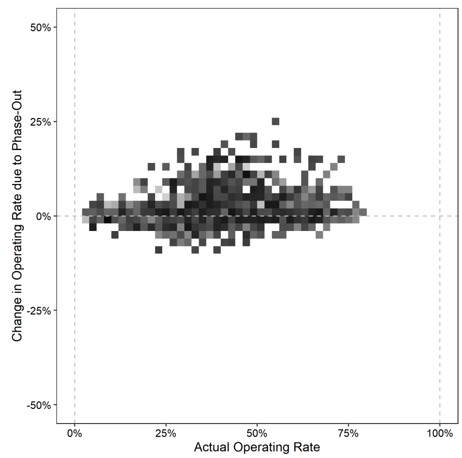


(b) Hard Coal

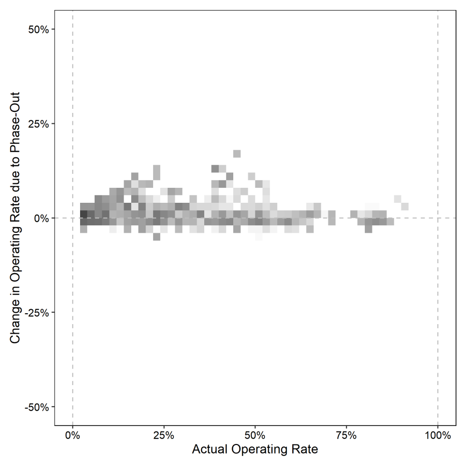


(c) Natural Gas

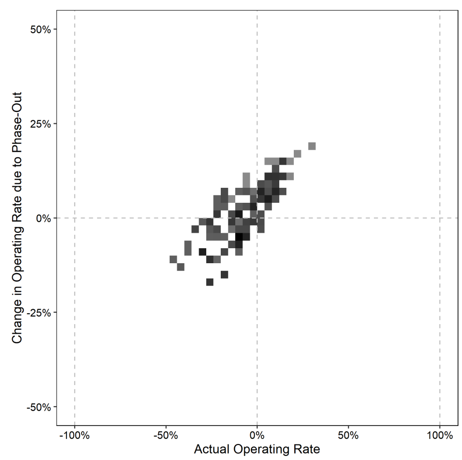


(d) Border

**Notes**: This figure illustrates the plant-level disaggregation of the machine learning prediction model results. The model predicts the operating rate of each power plant in each hour, where a value of 0% means the plant is offline and a value of 100% means it is running at maximum capacity. These figures plot plant-level annual average operating rates. The x-axis corresponds to each plant's operating rate in the baseline scenario with the phase-out. The y-axis corresponds to the impact of the phase-out on plant-level operations. This is determined by the difference between the predictions in the scenario with the phase-out versus the scenario without the phase-out. Darker areas indicate higher numbers of plant-year observations. Each panel refers to a different type of dispatchable electricity source. Panel (a) covers lignite plants, panel (b) covers hard coal plants, panel (c) covers gas plants and panel (d) covers border points. Oil plants are not shown because they are a very small portion of total capacity and are largely unaffected by the phase-out.

Figure B.6: Map of Plant-Level Changes in Production due to the Phase-Out

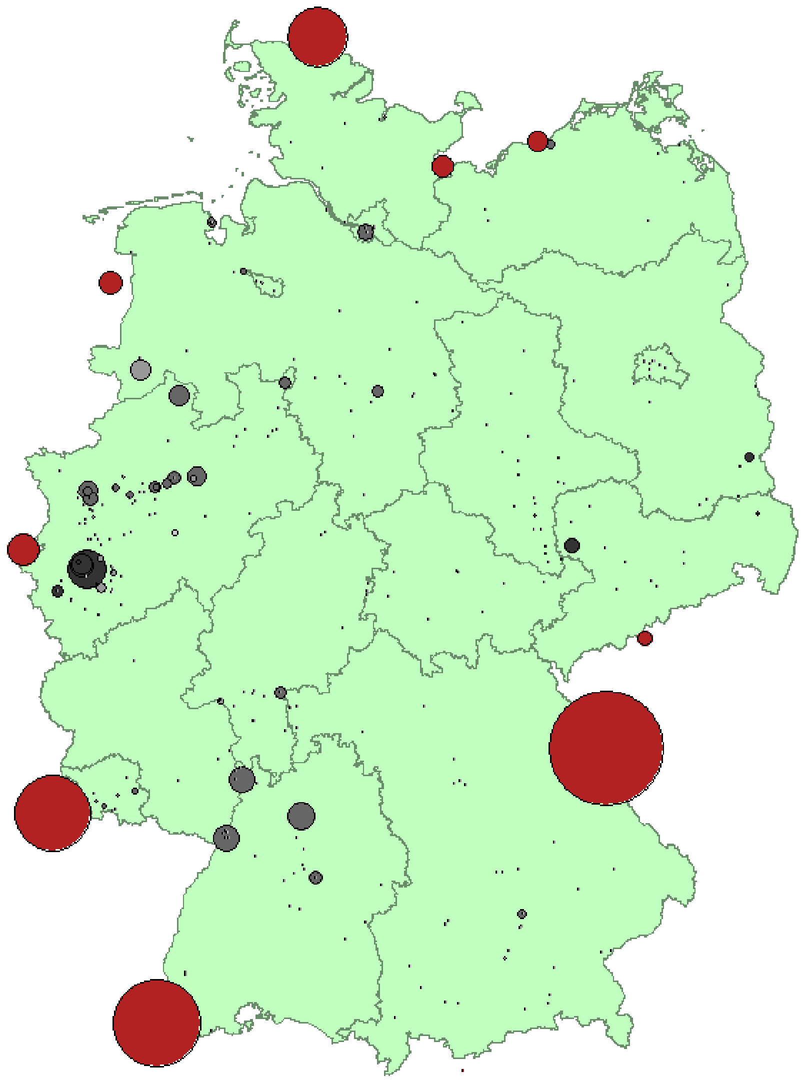

**Notes**: This map illustrates the location of the fossil-fuel-fired plants or border points that increased their electricity production as a result of the nuclear phase-out policy. The size of the circle reflects the amount of additional production provided by the fossil-fuel plant or border point. Points in red are border points and points in grey are fossil-fuel plants. Lignite plants are depicted in the darkest grey, followed by hard coal, then natural gas, and finally oil plants are depicted in the lightest grey.

# C Alternative Estimates of Local Air Pollution Costs

The health damages in Table 4 are calculated by linking data on plant-level emissions rates to pollution exposure using atmospheric chemistry modeling. As a complement to these estimates, we present a secondary approach to estimate the external costs imposed by increased ambient air pollution caused by the phase-out. To this end, we use granular air pollution monitor data to estimate how changes in output from hard coal and lignite plants affect ambient $PM_{2.5}$ concentrations. We then calculate the resulting excess mortality associated with the phase-out using similar methods to the ones underlying Table 4.

To proceed, we consider all air pollution monitors with daily $PM_{2.5}$ records in our sample period. We focus on the impacts on pollution of changes in output from hard coal and lignite plants because the evidence in Table 4 demonstrates that those 2 sources were by far the largest emitters. We construct a panel with daily ambient $PM_{2.5}$, generation, and wind direction for all generating unit/monitoring station pairs within a distance bandwidth of 250km and estimate the following regression:[37]

$$Y_{i,y,m,d} = \alpha + \beta G_{p,y,m,d} + \gamma_{i,p} + \theta_m + \psi_y + u_{i,p,y,m,d} \tag{C.1}$$

where $Y_{i,y,m,d}$ represents ambient concentrations of $PM_{2.5}$ at monitor $i$ on day $d$ of month $m$ of year $y$. $G_{p,y,m,d}$ represents daily generation at plant $p$ linked to monitor $i$. The parameter of interest $\beta$ captures the impact of an additional GWh of generation on ambient concentrations of $PM_{2.5}$ in micrograms per cubic meter. The

[37]There are 50 hard coal and lignite plants in our sample and 212 air pollution monitors with daily $PM_{2.5}$ measurements.

model includes fixed effects for each unit/monitor pair (e.g., to control for attributes specific to each pair such as generating unit efficiency, distance between unit and monitor, and the presence of other stationary sources of air pollution), as well as month-of-year and year-of-sample fixed effects (which control for seasonality and trends in air pollution and generation constant across all unit/monitor pairs).

In some specifications, we augment the model in Appendix Equation (C.1) by controlling for generation from all other units in the distance bandwidth. We also consider specifications interacting generation with indicators for the unit being upwind, downwind, or neither upwind or downwind from a given air quality monitor on a given day.[38] Based on previous literature (e.g., Deryugina et al. (2019)), we expect the impact of generation from units upwind from a pollution monitor to have a larger impact in driving ambient $PM_{2.5}$ concentrations than units downwind from the monitor.

Appendix Table C.1 presents the results from the estimation of Appendix Equation (C.1). Column (1) shows that across all unit/monitor pairs, an increase in daily generation of 1 GWh leads to a 0.178 $\mu g/m^3$ increase in daily average $PM_{2.5}$ concentrations. Evaluated at the sample means, this corresponds to an elasticity of about 0.1. Air pollution recorded at one given monitoring station is likely to reflect emissions from a host of mobile and stationary sources, including other power plants. To this end, in column (2) we add a control for generation from all other units within the 250km distance bandwidth. Thus, the regression isolates the link between ambi-

[38]Specifically, we merge daily station-level data on wind direction to pollution monitors. Then, we calculate the bearing of each unit relative to each monitor for unit/monitor pairs within 250km of each other. Finally, a unit is classified as being upwind from a monitor if the absolute difference between the wind direction from the monitor and the relative bearing between monitor and unit is less than 45 degrees. Similarly, a unit is “downwind” from a monitor (i.e., the wind blows from monitor to unit) if the absolute difference between wind direction and relative bearing is between 135 degrees and 225 degrees.

ent $PM_{2.5}$ and generation from a given unit, holding constant generation at all other units. The estimated coefficient drops to 0.057 (so about one-third in comparison to column (1)), but remains statistically significant.

The model in column (3) estimates a separate coefficient on electricity production interacted with the upwind, downwind, or neither upwind or downwind indicators described above. The evidence indicates, as expected, that generation at units upwind from a monitor is the primary driver of the observed relationship between $PM_{2.5}$ and generation shown in column (2). Column (4) further shows that increasing the maximum distance bandwidth between generating units and air pollution monitors to 300 km does not alter this conclusion.

Using the estimates from Appendix Table C.1, we can derive estimates of the increase in premature mortality due to changes in ambient $PM_{2.5}$ concentrations induced by the phase-out. To do so, we utilize the dose-response function from the ESCAPE project (Beelen et al., 2014).[39] Based on the hazard ratio estimated as part of this project, we can calculate the increase in mortality caused by the additional air pollution due to the phase-out using the following formula:

$$VSL \times \sum_{a=25}^{85} (POP_{a,r,y} \times MR_{a,y} \times \left(1 - \frac{1}{\exp(\rho \Delta PM2.5_{r,s,y})}\right)) \tag{C.2}$$

where $POP_{a,r,y}$ is the population of people aged $a$ in state $r$ in year $y$ and $MR_{a,y}$ is the mortality rate for age $a$ in year $y$; note that ages 85+ are binned into one

[39]The European Study of Cohorts for Air Pollution Effects (ESCAPE) is one of the few large-scale studies on the health impact of air pollution exposure in Europe. It is based on 22 European cohort studies with a total study population of more than 350,000 participants. Specifically, the ESCAPE project reports that the mortality rate when $PM_{2.5}$ exposure is X + 5 micrograms per cubic meter divided by the mortality rate when $PM_{2.5}$ exposure is X micrograms per cubic meter is 1.07.

Table C.1: Estimated Effect of Coal-Fired Electricity Generation on Ambient PM2.5 Concentration Levels

| Dep. Var.: $PM_{2.5}$ Concentration Levels (micrograms per cubic meter) | | | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Electricity Production (GWh) | 0.178*** | 0.057*** | | |
| | (0.026) | (0.010) | | |
| Upwind × Production (GWh) | | | 0.115*** | 0.123*** |
| | | | (0.016) | (0.016) |
| Downwind × Production (GWh) | | | 0.013 | -0.010 |
| | | | (0.011) | (0.011) |
| Neither DW nor UW × GWh | | | 0.051*** | 0.046*** |
| | | | (0.010) | (0.009) |
| GWh from Other Plants | | 0.033*** | 0.030*** | 0.026*** |
| | | (0.004) | (0.003) | (0.003) |
| GWh from Other Upwind Plants | | | 0.009*** | 0.009*** |
| | | | (0.002) | (0.001) |
| GWh from Other Downwind Plants | | | -0.003* | -0.002** |
| | | | (0.001) | (0.001) |
| Distance Bandwidth | 250 | 250 | 250 | 300 |
| Unit/Monitor Pair FE | Yes | Yes | Yes | Yes |
| Month-of-Year FE and Year FE | Yes | Yes | Yes | Yes |
| Number of Obs. | 2,341,993 | 2,341,993 | 2,341,993 | 3,108,471 |

**Notes:** This table presents the results from regressions of daily average $PM_{2.5}$ concentration levels on daily unit-level electricity production. We focus on non-CHP lignite and hard coal fired electricity generating units in Germany over the sample period 1/1/2015-12/31/2019. The unit of observation is a fossil fuel unit matched to an air quality monitor in a day-of-sample. Standard errors, reported in parentheses, are two-way clustered by unit/monitor pair and month-of-sample. For columns (1) to (3), the estimation sample includes all units and air quality monitors within 250km of each other; column (4) includes all pairs within 300km of each other. All specifications include fixed effects for each unit/monitor pair as well as month-of-year and year-of-sample fixed effects. In column (1), we regress daily average $PM_{2.5}$ on electricity generation (in GWh) for each unit/monitor pair with no additional controls. Column (2) adds the total production from all units within the relevant distance bandwidth other than the one in the unit/monitor pair. The specifications in columns (3) and (4) break down the impact of generation on $PM_{2.5}$ by whether the unit is upwind, downwind, or neither upwind nor downwind from the monitor.

category. The parameter $\rho$ corresponds to the hazard ratio estimated in Beelen et al. (2014) and $\Delta PM_{2.5}$ is the change in ambient $PM_{2.5}$ concentration levels attributed to the phase-out.

The change in ambient $PM_{2.5}$ attributed to the phase-out is calculated in three steps. First, we calculate $\Delta Gen_{i,t}$, the change in electricity generated from unit $i$ in hour $t$ due to the phase-out. We multiply this difference by 0.057, the estimated impact of fossil-fuel-fired generation on fine particulates from Column (2) in Appendix Table C.1. Finally, we take the average of the resulting changes in $PM_{2.5}$ over units of each source (either hard coal or lignite) in each state in each year of sample from 2012 to 2019.

The resulting estimates of monetized mortality damages are presented in Appendix Table C.2. Specifically, we report the average daily concentrations of $PM_{2.5}$ under the phase-out (column 1) and no phase-out (column 2) scenarios, as well as the difference between the two (column 3). The entries are annual averages over 2012-2019. The middle and bottom rows report estimates of premature mortality driven by exposure to $PM_{2.5}$ and the monetized damages from this premature mortality using the formula discussed above. The key result is that the phase-out resulted in an increase in ambient $PM_{2.5}$ concentration levels of 1.4% (0.19 $\mu g/m^3$), which translates into 334 additional excess deaths per year. This amounts to €0.9 billion in annual damages.

Taken together, the results in Table 4 and Appendix Table C.2 paint a consistent picture of the impact of the nuclear phase-out on pollution-caused excess death. Our estimates attribute 330 to 800 excess deaths per year to the increase in air pollution caused by the nuclear phase-out. This amounts to monetized health damages of €0.9 to €2.0 billion per year during the 2012-2019 period. Our preferred esti-

Table C.2: Estimated Impact of the Phase-Out on Air Pollution Mortality Damages

| | **Average with Phase-Out (1)** | **Average without Phase-Out (2)** | **Change (3)** | **Change (%) (4)** |
|---|---|---|---|---|
| **$PM_{2.5}$ Concentrations ($\frac{\text{micrograms}}{\text{cubic meter}}$)** | 14.01 | 13.83 | 0.19 | 1.37 |
| Lignite | 15.38 | 15.11 | 0.26 | 1.72 |
| Hard Coal | 12.65 | 12.54 | 0.11 | 0.88 |
| **Mortality ($\frac{\text{Excess Deaths}}{\text{Year}}$)** | | | 333.8 | |
| Lignite | | | 78.0 | |
| Hard Coal | | | 255.7 | |
| **Pollution Damages ($\frac{\text{Bn. Euros}}{\text{Year}}$)** | | | 0.90 | |
| Lignite | | | 0.21 | |
| Hard Coal | | | 0.69 | |

**Notes:** This table reports alternatives estimates of the monetary damages associated with the premature mortality resulting from the additional exposure to ambient $PM_{2.5}$ attributable to the nuclear phase-out. We use the estimated relationship between daily ambient $PM_{2.5}$ and daily generation from lignite and hard coal fired power plants reported in column (2) of Appendix Table C.1. We multiply this estimated relationship by the predicted changes in production from each plant due to the phase-out. The resulting changes in ambient $PM_{2.5}$ concentration levels are converted to a change in premature mortality using dose-response estimates from the ESCAPE project (Beelen et al., 2014). We monetize this additional premature mortality using the base-case VSL of €2.56 million for Germany. We do not report the absolute levels of mortality or damages because the baseline levels of pollution recorded at monitors are not attributable entirely to power plant activity; for example, industrial facilities and mobile sources also emit these pollutants.

mate is the €2.0 billion per year in damages calculated based on reported emissions (Table 4). This is because the analysis using reported emissions considers a more complete set of pollutants and draws on a more sophisticated analysis of pollution transport and exposure. The results presented in Appendix Table C.2 based on our estimated relationship between $PM_{2.5}$ levels and electricity production serve as a valuable complementary validation exercise based on an entirely distinct approach.