Discussion Paper Series – CRC TR 224

Discussion Paper No. 806
Project B 07

Electricity Generation, Carbon Pricing, and Air Quality: Evidence from the EU ETS

Nico Lukas Kasparetz[1]
Laure de Preux[2]
Ulrich J. Wagner[3]

November 2026

[1]Department of Economics, University of Mannheim. Email: nico.kasparetz@uni-mannheim.de
[2]Imperial College Business School, Imperial College London, and CEPR. Email: l.depreux@imperial.ac.uk
[3]Department of Economics, University of Mannheim, and CEPR.Email: ulrich.wagner@uni-mannheim.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Electricity Generation, Carbon Pricing, and Air Quality: Evidence from the EU ETS*

Nico Lukas Kasparetz† Laure de Preux‡ Ulrich J. Wagner§

November 17, 2026

Abstract

We develop an integrated framework to quantify the short-run impacts of carbon pricing on climate and air pollution externalities in the electricity sector. Combining machine learning and chemical transport modeling, we track policy impacts all the way from electricity dispatch and emissions of CO_2 and air pollutants at the generating unit to formation of secondary $PM_{2.5}$, population exposure and associated premature deaths. Using hourly data from 732 fossil generating units in 13 European countries, we estimate that carbon pricing under the European Union Emissions Trading System (EU ETS) reduced CO_2 emissions by 5.7% between 2016 and 2021. The abatement channel is coal-to-gas substitution which occurs within and between interconnected electricity markets. Carbon pricing reduces $PM_{2.5}$-related mortality due to co-emissions by 8.9%. Monetized health co-benefits exceed the incremental generation costs by more than €5 billion, implying a net welfare increase in Europe on top of global climate benefits. Counterfactual analyses reveal that net co-benefits per tonne amount to 60% of what is feasible with socially optimal pollution taxes, and that marginal co-benefits are weakly decreasing in the permit price. Overall, our results reinforce the role of cap-and-trade as an efficient policy instrument to internalize both carbon and air pollution externalities of power generation.

Keywords: climate change; air pollution; health benefits; emissions trading; electricity

JEL-codes: H23, Q41, Q53, Q54, Q58, L94

*We are grateful to Yixuan Gu and Daven Henze for graciously providing us with adjoint sensitivities based on Gu et al. (2023) and to Siyu Zhang for help with processing those. We thank Leo Launhardt, Michael Indenkämpen, Freya Rauls and Yuqi Wang for excellent research assistance. We have benefited from valuable comments and suggestions from workshop participants at ETH Zurich, ZEW Mannheim, Tsinghua University, and University of Mannheim, in particular from our discussant Marten Ovaere. Support by the German Research Foundation (DFG) through CRC TR 224 (Project B07) is gratefully acknowledged. This project received financial support from the European Research Council under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 865181). Declarations of interest: none.

†Department of Economics, University of Mannheim. nico.kasparetz@uni-mannheim.de

‡Imperial College Business School, Imperial College London, and CEPR. l.depreux@imperial.ac.uk

§Department of Economics, University of Mannheim, and CEPR. ulrich.wagner@uni-mannheim.de

1 Introduction

To avoid dangerous levels of global warming beyond 1.5°C, greenhouse gas (GHG) emissions must drop to net zero within the next thirty years (IPCC, 2023). The electricity sector plays a critical role for achieving this. On one hand, electricity generation accounts for about 40% of global carbon dioxide (CO_2) emissions per year (International Energy Agency, 2026). On the other hand, emissions reductions by sectors like transportation, buildings and manufacturing increasingly rely on electrification. The combination of ambitious abatement targets, heavy reliance on fossil fuels and growing electricity demand suggests that decarbonizing the electricity sector will entail substantial costs. Since 2005, several of the world's largest electricity markets–including the European Union (EU), California, and China–have been regulated with cap-and-trade policies which encourage cost-effective CO_2 abatement. This instrument is particularly well-suited to the electricity sector because firms have the incentives and the ability to factor carbon prices into their production decisions, and because carbon leakage poses smaller challenges than for industrial emitters. However, additional efficiency implications may arise as CO_2 is jointly released with air pollutants that pose direct risk to human health.

These co-emissions create scope for ancillary health benefits, especially when a policy mandates CO_2 emissions reductions across all polluters which simultaneously reduce co-emitted air pollutants. However, with market-based climate policies this cannot be taken for granted. In the electricity sector, carbon pricing reduces CO_2 emissions by shifting generation to less carbon-intensive power plants—both within and between interconnected electricity markets. Whether such reallocation increases or decreases total population exposure to air pollution depends not only on plants' pollution intensities and proximity to population centers but also on weather and atmospheric chemistry, all of which matter for pollution transport and the formation of secondary air pollutants like fine dust ($PM_{2.5}$). Unless air pollution externalities are fully internalized by existing regulations, this creates regulatory trade-offs or complementarities which ought to be considered when designing mutually consistent climate and air pollution control policies for the power sector.

This paper develops a new framework for measuring such trade-offs in the context of electricity firms regulated under the EU Emissions Trading System (ETS), a cap-and-trade system for CO_2 and other GHG emissions. To overcome the challenges presented by this endeavor, we draw on elements from program evaluation, integrated assessment and atmospheric chemistry. We start by estimating the impact of carbon pricing on short-run production decisions by fitting a machine-learning model to hourly production data from 732 fossil-fuel fired generating units located in 13 countries. The model accounts for economic determinants of dispatch (the merit order), grid congestion, and cross-border electricity trade. For each generating unit, we map electricity generation predicted by the model to emissions of CO_2 and $PM_{2.5}$ precursor pollutants, using emissions data collected under the auspices of various EU air quality directives. To track pollution from these sources to the place of impact, we use a custom-built chemical transport model capable of mapping changes in generation into changes in ambient concentrations of $PM_{2.5}$ on a grid covering all of Europe. Our final health outcome, total premature deaths due to $PM_{2.5}$, is estimated by combining data on population and ambient pollution concentrations in each grid cell with dose-response functions from the literature.

We estimate the causal impact of the ETS on mortality and other outcomes by comparing electricity market outcomes predicted with observed carbon prices to predictions with counterfactual carbon prices. Predictions are based on a quantile-regression random forest model trained on more than 28 million observations, which allows carbon prices to flexibly interact with unit, firm and market-level determinants of dispatch. We examine short-run impacts of carbon pricing through counterfactual scenarios that hold electricity demand and the power plant fleet fixed. The key mechanism is policy-induced fuel switching from carbon-intensive hard coal and lignite towards relatively low-carbon natural gas.

Our results show that the EU ETS induced significant fuel switching both within and between electricity markets, effectively reallocating generation across Europe's fossil power plant fleet in major ways. These shifts between fuels and across space had economically significant effects on two environmental outcomes of key policy interest. First, we attribute a 5.7% reduction in CO_2 emissions to the introduction of carbon prices. Second, we credit the associated reduction in co-pollution emissions with a 8.9% reduction in $PM_{2.5}$-related premature deaths, suggesting that policies targeting air pollution and climate externalities of electricity generation are complements. This is not self-evident because mortality changes depend on complex interactions of the market-based policy instrument with abatement costs, co-pollution intensities, topography and weather (which determine dispersion of those co-pollutants), and the population density exposed in places downwind from power plants.

We investigate this further by decomposing the mortality change into changes in scale, fuel mix, and sorting. We find that scale and fuel mix effects both reduce mortality. The former reduces co-emissions by lowering CO_2 emissions per kilo-watt hour, the latter by increasing the share of relatively clean natural gas in the fuel mix, which lowers co-pollution emitted per tonne of CO_2. Sorting turns out to have ambiguous effects, however. On the one hand, carbon pricing favors generation at more efficient units that emit fewer air pollutants per unit of output (pollution intensity sorting). On the other hand, output is shifted to units in locations where co-pollution causes higher damages (damage sorting), and disproportionately away from units where pollution-intensity and damages are negatively correlated (joint sorting). The countervailing effect of damage and joint sorting negates one third of the maximum mortality reductions that would be achievable if scale, fuel mix and sorting effects were aligned.

In spite of such adverse sorting effects, the health co-benefits we attribute to the EU ETS more than justify the associated abatement costs—irrespective of any global benefits due to social costs of carbon avoided. At €35.9 per tonne of CO_2 avoided, the estimated average health co-benefit exceeds estimated average abatement cost of €6.5 by a comfortable margin. The sign and magnitude of the net benefit strongly suggests that pre-existing regulation for air pollution emissions, which we cannot perfectly observe, had not been set at socially optimal levels, and allows the conclusion that carbon trading among European power plants did indeed enhance efficiency by internalizing air pollution externalities.

Given the considerable heterogeneity in pollution and damage intensities across units in our sample, carbon prices must be considered a rather blunt instrument to address air pollution externalities. Notwithstanding this, additional counterfactual analysis shows that carbon prices under the EU ETS internalized almost 60% of the mortality reduction that could be achieved

by Pigouvian air pollution taxes differentiated by pollutant, location and calendar month of emissions. The reason is that scale, fuel mix, and intensity sorting effects of carbon pricing all increase health co-benefits. By additionally targeting particularly damaging generators, the Pigouvian tax minimizes adverse mortality effects of damage and joint sorting. While the former effects (scale, fuel mix, and pollution intensity sorting) are likely to hold in settings other than the EU ETS, the sign and magnitude of the latter effects (damage and joint sorting) fundamentally depend on the spatial distribution of the power plant fleet.

A forward-looking counterfactual considers the effect of higher-than-actual carbon prices. We simulate outcomes for the observed time series of EU ETS permit prices, which varied between €4 and €88 during the sample period, and impose a permit price floor of €50. We find significantly larger reductions in CO_2 emissions (-13.9%) and $PM_{2.5}$-related mortality (-20.4%) in this scenario. Relative to the scenario with actual carbon prices, the incremental abatement costs 80% more and yields a 13% lower co-benefit per tCO_2. While net co-benefits remain positive in this scenario, the marginal net co-benefits of fuel switching diminish as the carbon price rises. At very high carbon prices, the realization of additional health co-benefits will then depend on changes to elements of the energy system which are largely held fixed in this study, such as the supply of electricity from renewables, nuclear, and hydropower, or the large-scale adoption of storage or carbon-capture-and-storage technologies.

Related Literature Our study provides the first micro-founded evaluation of the ancillary health benefits of climate policy, focusing on power plants operating in the world's first and most stringent carbon market. In so doing, our paper builds on prior research on the EU ETS, on the external costs of electricity generation, and on integrated assessment.

An early literature in the tradition of integrated assessment models (IAM) has sought to quantify the ancillary benefits of climate policy in terms of avoided pollution by simulating emissions paths with and without the policy (Aunan et al., 2007; Burtraw et al., 2003; Ekins, 1996; Ekin, 1996; Pittel and Rübbelke, 2008; Rypdal et al., 2007; van Vuuren et al., 2006; Driscoll et al., 2015). Given the uncertainty about future regulations, technology, economic development, demography, and natural activities (Davis et al., 2000), it is very difficult to ascertain the path of emissions *ex ante* with confidence. Our study relaxes these assumptions by letting ex-post data inform an empirical model of how hourly generation at hundreds of generating units varies with the carbon price. The granularity of our approach is particularly well-suited to our question because it allows us to model decentralized pollution releases that are characteristic for compliance with market-based regulatory instruments.

A related strand of literature also focuses on ex-post evaluation of emissions trading schemes (Carlson et al., 2000; Fowlie et al., 2012), and in particular, the EU ETS. Econometric analyses based on country-level emissions data credit the policy with significant emissions abatement of not only CO_2 (Bayer and Aklin, 2020) but also air pollutants with potential health co-benefits potentially in the hundreds of billions of Euros (Basaglia et al., 2024). This motivates further micro-level research around the health co-benefits of carbon abatement in the EU ETS. Existing micro-econometric studies of the EU ETS have analyzed CO_2 abatement in manufacturing (Martin et al., 2016; Colmer et al., 2024) and in the power sector (Abrell et al., 2022; Zaklan,

2023), but not health co-benefits, which is the subject of this paper.

While the efficiency benefits of cap-and-trade increase with abatement cost heterogeneity, empirical evidence suggests that heterogeneity in environmental damages can work in the opposite direction unless appropriate exchange rates are applied for permit trades between air sheds (Mauzerall et al., 2005; Muller and Mendelsohn, 2009). Moreover, cap-and-trade programs have come under scrutiny on suspicion of shifting pollution emissions—and the environmental damages associated with them–disproportionately towards minorities and low-income households (Fowlie et al., 2012; Grainger and Ruangmas, 2018; Deschenes and Weber, 2025). Empirical research on the air pollution effects of carbon trading has so far focused on carbon markets in California (Cushing et al., 2018; Hernandez-Cortes and Meng, 2023) and China (Zhang and Almond, 2026). We contribute to this literature by measuring and comparing CO_2 abatement costs and monetized health benefits in the context of the world's first and largest carbon market.

Finally, our paper builds upon and extends the empirical research on the social cost of electricity generation. While a series of papers has estimated the external costs of electricity generation in the U.S. (e.g., Holland and Mansur, 2008; Holland et al., 2016; Borenstein and Bushnell, 2022), we focus on the EU where a lack of a spatially differentiated, marginal damage estimates has hitherto prevented a similar endeavour. To estimate causal impacts of policy on electricity generation, we adopt a machine learning approach following recent research on the UK carbon price floor (Abrell et al., 2022) and Germany's nuclear exit (Jarvis et al., 2022). In broadening the scope of this literature from national policies to an EU wide carbon pricing policy, our paper makes two critical methodological innovations. First, we adapt the model framework developed by Jarvis et al. (2022) for a single market to accommodate settings with up to 21 interconnected electricity markets that trade electricity with each other. We develop a novel algorithm for counterfactual analysis whereby predictions of electricity generation and cross-border electricity flows across European electricity markets are such that markets clear and a new, counterfactual equilibrium obtains. Second, we use a full-complexity chemical transport model to compute $PM_{2.5}$-related mortality impacts that vary by pollutant, source, and calendar month. In so doing, our study demonstrates how recent advances in atmospheric chemistry can be used to increase the rigor of counterfactual air quality scenarios in economics, which has thus far heavily relied on reduced-complexity models such as mass transport or source-receptor matrices.

2 Background

2.1 The EU Emissions Trading Scheme

The EU ETS is a cap-and-trade system for CO_2 emissions, initially focused on power generation and energy-intensive manufacturing industries. Launched in 2005 by the Emissions Trading Directive[1], the policy has been covering roughly 12,000 installations in 31 countries.

[1]Directive 2009/29/EC of the European Parliament and of the Council amending Directive 2003/87/EC so as to improve and extend the greenhouse gas emission allowance trading scheme of the Community (2009) OJ L 140, 5.6.2009, p. 63-87 (Emissions Trading Directive).

Participation in the EU ETS is mandatory for all combustion installations with a rated thermal input exceeding 20 megawatts (MW), which includes conventional power plants. In addition, the scheme covers large installations engaging in activities that are intensive in CO_2 emissions such as mineral oil refineries, coke ovens, iron and steel, and factories producing cement, glass, lime, bricks, ceramics, and pulp and paper. Since 2012, the system also includes airlines and aluminium manufacturing installations, as well as several greenhouse gases other than CO_2.

Participating installations are required to surrender a pollution permit, known as an EU allowance (EUA), for each metric ton of CO_2 emitted per year. These allowances are distributed to firms in part for free or through an auctioning system. The annual number of permits issued—the 'cap'—is chosen so as to limit the total amount of pollution below the levels that would arise under business-as-usual conditions. Initially exceeding 2 billion tons of CO_2 per year, the cap declined at an annual rate of 1.74% until 2020 and has since been decreasing by 2.2% per year, so as to ensure the EU interim target for 2030 of a CO_2 emissions reduction by 43% from 2005 levels.

The mechanical decline in the cap induces scarcity of pollution permits and puts upward pressure on the price for emission permits as firms trade on the carbon market. In this way, the EU ETS has established a uniform carbon price that signals the opportunity cost of emitting CO_2 to all market participants. The CO_2 price provides each of them with an incentive to reduce emissions up until the point at which they are indifferent between buying one permit at the market price and paying the cost of reducing its emissions by one additional ton of carbon.

As the world's first carbon trading policy for individual firms, the EU ETS has created a multi-billion Euro market for emission permits. With EUA prices nearing €100, it is arguably the world's most stringent carbon pricing policy. The EU ETS is the cornerstone of the EU's policy roadmap towards carbon neutrality and widely regarded as a prototype for a future global carbon market.[2]

2.2 Air pollution control policies

Pollution emissions from fossil fuel–fired power plants have long been subject to both national and European regulations. A series of increasingly stringent EU regulations has targeted large combustion plants (LCPs), i.e., industrial combustion plants with a total rated thermal input equal to or greater than 50MW. The central legal instrument during the early years of the EU ETS was the Large Combustion Plant (LCP) Directive.[3] It imposed emission limit values for key pollutants such as SO_2, NO_x, and PM. The LCP directive also required older plants to either retrofit pollution-control technologies (e.g. flue gas desulfurization) or close by 2015.

From 2016 onwards, the LCP was effectively replaced by the Industrial Emissions Directive (IED).[4] The IED required permits based on best available techniques (BAT), with stricter and more harmonized emission limit values. For power plants, this meant tighter operational conditions and a stronger link between environmental performance and permitting. From around 2015 onward, operators had to comply with increasingly strict limits derived from BAT

[2] A comprehensive review of the history and structure of the EU ETS can be found in Ellerman et al. (2016).

[3] European Union, Directive 2001/80/EC on the Limitation of Emissions of Certain Pollutants into the Air from Large Combustion Plants, OJ L309 (2001): 1–21.

[4] European Union, Directive 2010/75/EU on Industrial Emissions, OJ L334 (2010): 17–119.

reference documents, which further accelerated emission reductions and technological upgrades.

The Pollutant Release and Transfer Register (PRTR) Regulation complements the LCP and Industrial Emissions Directives by requiring operators of power plants and large industrial facilities to report annual emissions of a broad range of air pollutants (subject to reporting thresholds).[5] Unlike the two directives, which set emission limits and permitting conditions, the PRTR focuses on transparency and public access to data through an EU-wide database, henceforth referred to as the E-PRTR. The more integrated regulatory framework introduced by the IED has established standardization, monitoring, and verification requirements that have benefited data quality in the E-PRTR.

Based on these data, LCP emissions have fallen dramatically since the onset of the EU ETS. From 2004 until 2024, SO_2 and dust emissions fell by 94% and NO_x emissions by 74% (European Environment Agency, 2025). The main reasons were technological changes and a reduction in the use of coal and other fossil fuels. Against this background, it is important not to confound the air quality impacts of those dedicated policies to those of carbon pricing. Our subsequent analysis disentangles these concurrent policy impacts by focusing on the short-run effects of carbon pricing on pollution emissions which are driven by fuel switching in response to daily variation in carbon prices. We take as given any longer-term decisions on plant closures and new plant openings, which might be driven by retrofits or BAT regulations stipulated in LCP and IED. This is conservative in the sense that it ignores such extensive-margin decisions even when they were indeed driven by carbon pricing.

National policies to limit SO_2 and NO_x emissions from power plants in European countries have generally operated alongside EU directives, but some countries have gone far beyond common policies in terms of imposing stricter standards, deploying economic instruments, and reorienting structural energy policies. First, when implementing EU directives through environmental permits, national authorities sometimes impose stricter conditions or faster compliance timelines, such as the Dutch Emission Limits Decree or the Austrian Boiler Emissions Act. Second, some countries have been pricing pollution via taxes on SO_2 and NO_x emissions, most prominently Sweden since the early 1990s. Third, many countries adopted technology- and fuel-based measures to reduce emissions. These include setting emission limits that effectively require the installation of flue gas desulfurisation (FGD) for SO_2 control and selective catalytic reduction (SCR) for NO_x. Examples include Germany's Large Combustion Plant Ordinance from the 1980s, Czechia's Clean Air Act from 1991, and Spain's National Emission Reduction Plan from 2007. Fourth, since the 2010s, national strategies increasingly rely on structural and economic measures, such as coal phase-outs, emissions ceilings, and climate-energy policies. Overall, most national policies have tended to promote cleaner technologies and accelerate the shift away from high-emission fossil fuels. However, prominent examples exist of national policies working in the other direction, such as Germany's "nuclear exit" (Jarvis et al., 2022).

[5] European Parliament and Council, Regulation (EC) No 166/2006 concerning the establishment of a European Pollutant Release and Transfer Register, OJ L33 (2006): 1–17.

2.3 Electricity markets in Europe

The European market for electricity is subdivided into several bidding zones, i.e., geographic areas in which electricity is traded at a uniform wholesale price, irrespective of where it is generated or consumed. Most bidding zones coincide with national borders, but some span several countries, such as the joint German-Austrian-Luxembourgian zone until October 2018. Other countries are split into multiple bidding zones on account of persistent transmission bottlenecks within them. Examples include Sweden where the north is rich in hydropower while demand is concentrated in the south, two Danish power systems that belong to different synchronous grids, as well as Italy with multiple bidding zones on the peninsula and the islands.

Within a zone, the grid is treated as if it had unlimited transmission capacity (akin to a copperplate). Congestion inside a zone is resolved by the transmission system operator (TSO) rather than through prices. Neighboring bidding zones are connected by interconnectors, so that electricity can flow from zones with low prices to zones with high prices until either prices equalize or the available cross-border transmission capacity is exhausted. Since 2014, most European zones have been integrated through Single Day-Ahead Coupling, which allocates cross-border capacity implicitly as part of the joint market-clearing process.

The principal market for wholesale electricity is the day-ahead market. Each day before noon, generators submit bids stating how much electricity they are willing to supply with each of their units at which price, separately for every hour of the following day. At the same time, retailers and large consumers submit corresponding demand bids. The power exchanges aggregate these bids into supply and demand curves, and a common algorithm (EUPHEMIA) clears all coupled zones simultaneously, taking into account the available transmission capacity between them. The result is a market-clearing price for every zone and hour which is paid to all accepted bids, together with a dispatch schedule stipulating which units will be active in a given hour and their level of output. Subsequent intraday trading, re-dispatch and balancing markets allows for last-minute adjustments to this schedule, for instance in response to forecast errors or unplanned outages.

The power sector is central to European climate policy. According to official data from the European Network of Transmission System Operators for Electricity (2024), between 2016 and 2021, total electricity generation in the EU-27 plus Great Britain averaged at 3,238 TWh per year, 38.2% of which came from fossil fuels. Of all fossil generation, hard coal and lignite jointly accounted for 52.8%, natural gas for 43.8%, and the remaining 3.4% came from oil, peat and coal gas. Concurrent with the expansion of renewable energy sources and the replacement of coal-fired generation by natural gas, the fossil share of electricity generation declined from about 40% in 2016 to 35% in 2021, and total CO_2 emissions from fossil generators fell from approximately 1,050 Mt to 780 Mt. The annual average emissions of 930Mt are equivalent to 59% of all verified emissions covered by the EU ETS and 51% of the annual cap on stationary installations.

3 Data

The objective of this study is to quantify local air quality externalities from fossil-fuel fired power generation and assess how those change under different carbon pricing regimes. To illustrate our approach and the associated data requirements, we write the aggregate mortality impact of such externalities as

$$D = \sum_{i=1}^{N} e_i(y_i) \cdot \mathbf{z}_i' \boldsymbol{\lambda}_i. \tag{1}$$

where e_i denotes power unit i's CO_2 emissions as a function of its electricity output y_i, $\mathbf{z_i}$ is a $K \times 1$ vector containing co-pollution intensities at plant i for each co-pollutant k, expressed in units of co-pollutant k emitted per unit of e, and $\boldsymbol{\lambda_i}$ contains damage sensitivities for each co-pollutant, i.e. coefficients that convert pollution emissions at unit i's location into premature deaths due to $PM_{2.5}$ in Europe.

We estimate the components of this function in several steps. For given electricity output y_i, we compute fuel input requirements based on the unit's fuel type and net electrical efficiency. These are used to compute CO_2 emissions e_i using standard conversion factors. We collect data on co-emission intensities $\mathbf{z_i}$ for $K = 3$ important $PM_{2.5}$ precursors. The public health impacts of those co-emissions depend in complex ways on atmospheric dispersion, the chemistry of secondary air pollutant formation, and population exposure. We use a custom-built chemical transport model (CTM) that maps changes in source-specific co-emissions into changes in $PM_{2.5}$ exposure on a grid for Europe. Combined with age- and country specific dose response functions, the CTM provides pollutant and location specific, so-called adjoint sensitivities $\boldsymbol{\lambda}_i$ that allow us to compute marginal mortality impacts of co-emissions from each power unit. The remainder of this section explains the data used and calculations performed at each of these steps in detail.

3.1 Electricity generation, consumption, and trade

The Central Information Transparency Platform of the European Network of Transmission System Operators for Electricity (2024), henceforth referred to as ENTSO-E, provides detailed hourly data on European electricity markets. Our analysis is conducted at the level of a generating unit which typically refers to a block with a unique ENTSO-E identifier that is part of a fossil fuel power plant. All generating units can be further matched to the power plants they belong to. ENTSO-E provides information on the unit's bidding zone b, fuel type f, installed capacity κ, outages and, importantly, the actual generation $Y_i(t)$ in any given hour t.

At the level of the bidding zone b, we use ENTSO-E's "total load", $\bar{L}$, and compute net load by subtracting from this total generation Y_g from non-fossil technologies g, also available from ENTSO-E:

$$L_b(t) = \bar{L}_b(t) - \sum_g Y_{b,g}(t) \quad g \in \{\text{wind, solar, biomass, geo-thermal, hydro, nuclear}\} \tag{2}$$

This is the residual demand that must be met by domestic fossil generation and net imports. Moreover, ENTSO-E provides day-ahead electricity prices at the hourly level. Data on

Table 1: Power Plants and Generation Units, by Country and Fuel Type

	Generating units					Power plants				
Country	**All**	Lignite	Coal	Gas	Other†	**All**	Lignite	Coal	Gas	Other*
Austria	17		2	14	1	11		2	8	1
Belgium	18		2	16		13		2	11	
Czechia	28	23	4	1		9	7	1	1	
Denmark	13		10	3		11		8	3	
Germany	141	38	53	49	1	78	10	34	28	6
France	24		5	19		21		5	16	
Great Britain[a]	112		36	76		59		11	48	
Italy	109		18	91		64		6	57	1
Netherlands	41		8	33		28		6	21	1
Poland	115	31	79	5		32	5	22	5	
Portugal	12		2	10		5		1	4	
Sweden	4		1	3		3		1	2	
Spain	98		23	72	3	48		15	30	3
Total	732	92	243	392	5	382	22	114	234	12

Notes: [a]Great Britain (England, Scotland & Wales) constitute one bidding zone. Northern Ireland is part of the Irish bidding zone and not included.
†"Other" denotes units that switch fuel during the observation period. One Austrian unit switches from coal to gas, one German unit switches from lignite to gas, and three Spanish units switch from lignite to coal.
*"Other" refers to one Italian, one Dutch, and two Spanish plants that operate both coal and gas units, and four German plants that operate both coal and gas units, plus two German plants that operate both lignite and gas units. Additionally, one Austrian, and one Spanish plant contain units that switch fuels, and are included in these totals (the "Other" units in the left part of the table).

hourly electricity trades between bidding zones is taken from ENTSO-E's dataset "Scheduled Commercial Exchange".

We fill reporting gaps in these data with information from other sources. Starting in 2021, Brexit put an end to data submission to ENTSO-E from England, Wales and Scotland. We obtain unit-level generation for British units in 2021 from Elexon, the administrator of the GB Balancing and Settlement Code, through its Insights Solution API. We obtain data on load and generation per production type for post-Brexit Britain from the UK National Energy System Operator. British electricity prices in 2021 are available from the global energy think tank Ember. Because Italy did not report scheduled flows between its domestic bidding zones to ENTSO-E until Jan 1, 2018, we download pertinent data for 2016 and 2017 from Gestore dei Mercati Energetici (GME).

We merge additional information on generator characteristics from various sources. A unit's technology type—i.e., steam turbine, gas turbine, combined cycle, combined heat-and-power (CHP)—is obtained mainly from Open Power System Data (2020), complemented with data from Jarvis et al. (2022) (for German units) and the Global Energy Monitor.[6] The JRC Open Power Plants Database (Hidalgo Gonzalez et al., 2019) provides unit-level information on minimum stable load (the minimum capacity utilization rate at which the unit can be run

[6]available online at https://www.gem.wiki/.

continuously),[7] the ramping speed at which units can power up or down,[8] net electric efficiency, a unit's geo-location, and on unit entry or exit.[9]

Our analysis sample comprises 732 generating units that belong to 382 power plants located in 13 countries, summarized in Table 1, plus 82 border points. This gives us a panel data set with 28,200,032 unit × hour observations. In studying unit-level electricity dispatch decisions, we focus on the unit's hourly *utilization rate*

$$y_i(t) \equiv \frac{Y_i(t)}{\kappa_i} \tag{3}$$

which measures, on a scale from 0 to 1, how much of the unit's capacity is used in a given hour. For example, if a unit has a capacity of 100 MW, generation of 50 MWh in one hour corresponds to a utilization rate of 0.5. Table 2 reports descriptive statistics of the utilization rate at every hour over the sample period from 2016–2021, separately by fuel type and for border points. Earlier years are not included because ENTSO-E coverage is not complete. The Russian invasion of Ukraine in February 2022 imposed a series of large structural shocks to the European energy system, including physical curtailments of natural gas supply, direct government interventions in dispatch and fleet composition (e.g., reserve coal plants returning to the market, the Iberian gas price cap, EU-wide caps on inframarginal revenues), and coordinated demand reduction targets. Since these events contradict key assumptions of our counterfactual analysis (Section 4.5), we refrain from including data from 2022 and later years in the analysis.

Lastly, note that reporting to ENTSO-E is mandatory only for units with at least 100 MW of installed capacity. This constraint means that our sample comprises about 82.3% of all fossil electricity generation (83.0% for lignite, 79.3% for hard coal, and 83.6% for natural gas).

Figure 1 shows the geographic location of all units on the map along with their capacity. To keep track of bidirectional flows between any two adjacent and connected bidding zones, those are included as 'border points' on the map. We represent every border point as two virtual units, once per direction of flow. Following Holland et al. (2026), we impute the transmission capacity of each border point as the 99th percentile of the hourly scheduled cross-border flow in each year, pooling both directions.

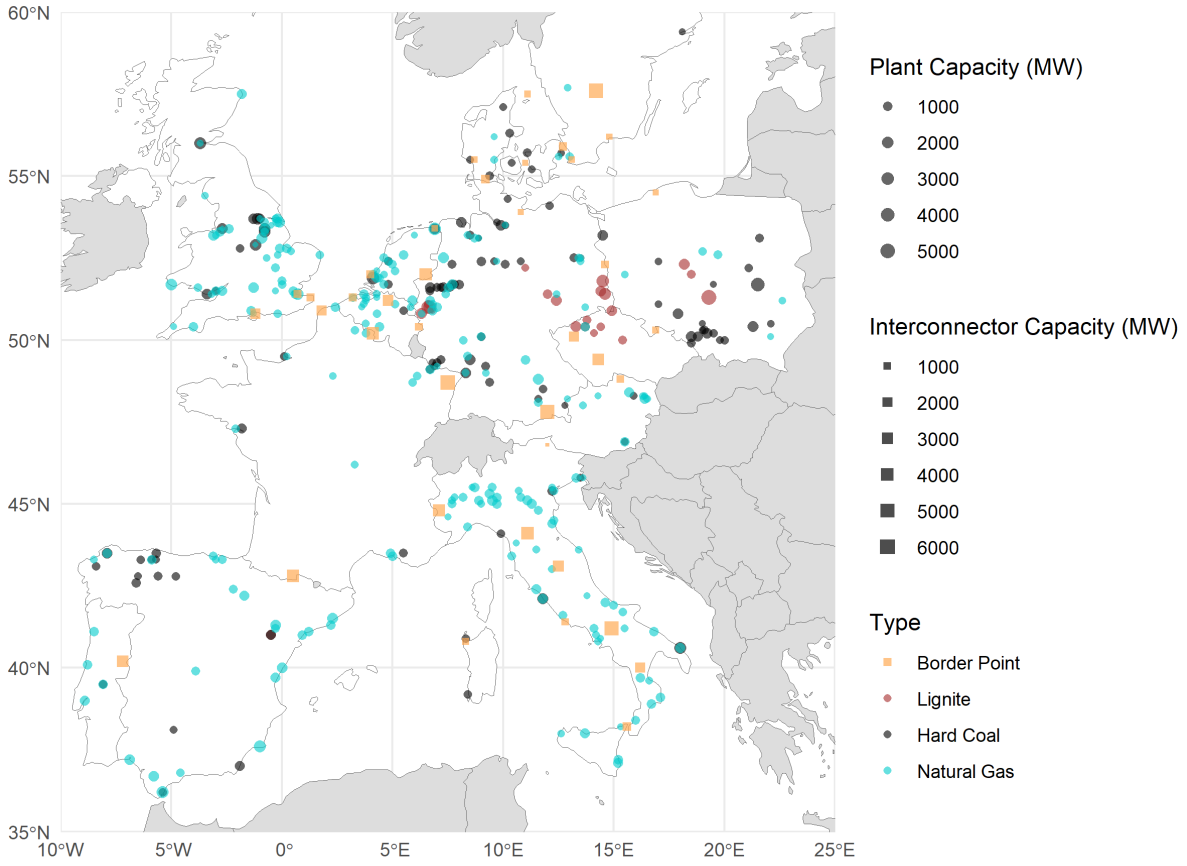
Transmission capacity between countries differs quite substantially. Germany is connected to eight neighboring countries with a combined capacity of 23.3 GW, equivalent to 45% of the aggregate capacity of its fossil units in our sample. In small countries like Belgium, Austria, and Denmark, cross-border capacity is more than two times larger than domestic fossil capacity. In contrast, the Spanish–French border, with an average capacity of 3.5 GW, is the only link between the Iberian Peninsula and the rest of our sample, equivalent to only about 10% of Iberian fossil capacity.

[7] Coal plants, for instance, need to run at a minimum load in order to ensure combustion stability and produce enough steam to maintain required temperature and pressure.

[8] Fossil fuel units generally cannot go from being shut off to maximum generation in an instant but need time to power up. In the case of lignite units, this can take several hours.

[9] Occasionally, units stay in the ENTSO-E register and continue to report zero generation even if they were already decommissioned.

Figure 1: Location of fossil-fuel fired generating units in the sample (aggregated to power plants



Notes: The map shows the location of all fossil-fuel power electricity generation units in our dataset, aggregated to the plant level for readability, as well as border points. Border points are virtual units that represent transmission capacity between bidding zones. The location is imputed as the geographic average of land-based borders. For underseas cables, the location at which the cable ends is used. Until September 2018, Czechia bordered the joint German-Austrian-Luxembourgian bidding zone, with a border point given by the geographic average of the Czech border to both Austria and Germany. Additional border points result from the split-up of the German-Austrian-Luxembourgian bidding zone into an Austrian and a German-Luxembourgian bidding zone in October 2018.

3.2 Generation cost

The private cost of electricity generation is determined by capital and operating costs. Since our analysis holds the power plant fleet constant, we consider the capital costs to be sunk and focus on operating costs only.

We assume that unit i's marginal cost MC in hour t to generate one MWh of electricity is comprised of variable operating costs and input costs:

$$MC_{i,t} = VOC_i + \frac{P^f_{i,t} + \tau_{i,t} + P^c_t \cdot \xi^f}{\eta_i}. \tag{4}$$

Estimates of variable operating cost VOC_i of steam turbines, gas turbines, and combined cycle turbines are taken from Sargent & Lundy (2019). Input costs are determined by the amount of fuel required to produce one MWh of output, multiplied by the fuel price and permit costs. To compute fuel input requirements, we invert the net electric efficiency η_i, defined as the proportion of a fuel's inherent energy that is converted into electric energy. Net electric efficiencies of conventional coal- or gas-fired power stations typically range between 30 and 40 per cent, meaning that generating one MWh of electricity requires the combustion of 2.5 to 3.3 MWh of fuel, with most of the energy difference being dissipated as waste heat. We multiply the fuel input requirement $\left(\eta_i^{-1}\right)$ by a term aggregating the fuel price $P^f_{i,t}$, transportation costs $\tau_{i,t}$, and carbon compliance costs, the latter being calculated as the product of the fuel's emissions intensity ξ_f (in tCO_2/MWh) and the carbon price P^c_t.

We obtain fuel price data from LSEG Datastream (LSEG, 2025). Natural gas prices (in €/ MWh) differ by trading hub (PSV for Italy, NBP for Great Britain, MIBGAS for Spain and Portugal, POLPX for Poland, THE for Germany, TTF for the Netherlands, and VTP for Austria and Czechia). Additionally, we impute natural gas prices for Denmark and (Southern) Sweden with data from the THE hub, and for Belgium using the TTF hub. For hard coal, we use daily price data from the API2 coal index, the standard benchmark for European steam coal used primarily in electricity generation. The hard coal market price is converted from €per metric tonne to €per MWh using conversion factors from Juhrich (2016) and Jarvis et al. (2022). We then add transportation costs in €per MWh, taken from an Open Source Electricity Model for Germany maintained by DIW German Institute for Economic Research (ELMOD-DE; cf. Egerer, 2016). We use this data to also calculate the "brown spread" as the price of coal relative to gas. A market price for lignite is not available due to the prevailing vertical integration of lignite mining and electricity generation. We follow previous research and assume a lignite price of €6/MWh (Egerer, 2016; Jarvis et al., 2022). Since lignite-fired generation occurs at the mine mouth, we assume no transportation cost.

As the relevant carbon price, we take the daily settlement price for European Union Allowances (EUAs) provided by LSEG (2025). This price is the official end-of-day price on the European Energy Exchange (EEX), a major trading venue for both spot and derivatives. To convert fuel input requirements into emissions, we use conversion factors ξ from Juhrich (2016).

Table 2 contains summary statistics on our marginal cost estimates, reported separately by fuel type and for the border points. Natural gas is on average the most expensive fossil fuel

Table 2: Summary statistics of outcome and predictors

Variable	Unit	Mean	SD	Min	p25	Median	p75	Max
Outcome								
Utilization rate	–	0.41	0.39	0.00	0.00	0.41	0.80	1.00
Lignite	–	0.71	0.32	0.00	0.56	0.85	0.96	1.00
Hard coal	–	0.44	0.38	0.00	0.00	0.51	0.79	1.00
Natural gas	–	0.36	0.39	0.00	0.00	0.19	0.77	1.00
Border points	–	0.26	0.33	0.00	0.00	0.08	0.48	1.00
Time								
Hour of day	–	11.53	6.91	0	6	12	18	23
Day of week	–	3.98	1.99	1	2	4	6	7
Month	–	6.46	3.53	1	3	6	10	12
Market conditions								
Net load	MW	15,738	12,206	−15,123	5,555	14,394	21,221	73,908
Price of hard coal	€/tonne	67.43	26.85	35.23	48.44	65.28	78.35	236.67
Price of natural gas	€/MWh	21.40	17.67	3.10	14.00	17.55	21.99	182.79
EUA price	€/tCO_2	20.77	16.53	3.91	5.94	19.93	26.29	88.88
Brown spread	–	3.65	1.03	0.71	3.08	3.47	4.08	12.31
Unit characteristics and merit order								
Net electric efficiency	–	0.49	0.20	0.23	0.35	0.49	0.54	1.00
Marginal cost	€/MWh	55.26	38.46	−24.62	37.11	46.31	58.44	926.38
Lignite	€/MWh	42.25	19.32	19.72	26.77	39.34	49.68	142.55
Hard coal	€/MWh	53.94	23.03	18.85	40.79	49.40	59.50	233.24
Natural gas	€/MWh	56.35	43.73	14.32	36.46	44.32	57.30	858.19
Border points	€/MWh	66.79	50.30	−24.62	43.25	53.58	69.57	926.38
Should run (dummy)	–	0.64	0.48	0	0	1	1	1
Capacity gap	–	−0.14	130.55	−158,845.58	−0.74	−0.28	0.39	39,644.38
Output gap	–	−0.61	18.60	−33,365.31	−0.89	−0.66	−0.41	5,994.18
Partly unavailable (dummy)	–	0.06	0.24	0	0	0	0	1
Latitude	degrees	47.79	5.30	36.18	43.31	50.13	51.65	59.35
Longitude	degrees	8.09	7.99	−8.95	2.16	8.87	14.30	22.56
Counterfactual marginal cost								
Marginal cost, €0 EUA price	€/MWh	41.50	33.40	−24.62	26.16	35.09	45.06	926.38
Lignite	€/MWh	19.10	1.40	15.99	17.93	19.36	20.46	23.93
Hard coal	€/MWh	34.54	11.78	13.61	26.77	32.93	40.38	146.09
Natural gas	€/MWh	47.61	38.55	6.46	30.59	38.25	49.32	787.83
Border points	€/MWh	51.90	43.99	−24.62	32.35	41.36	54.22	926.38
Marginal cost, €50 price floor	€/MWh	76.20	35.40	−24.62	56.25	71.00	84.21	926.38
Lignite	€/MWh	77.91	8.07	63.65	71.00	77.11	81.83	142.55
Hard coal	€/MWh	85.62	16.66	52.10	74.63	82.88	93.18	233.24
Natural gas	€/MWh	68.68	40.88	24.89	50.03	58.09	70.42	858.19
Border points	€/MWh	88.63	46.83	−24.62	67.45	81.83	97.02	926.38

Notes: Statistics are computed over 28,200,032 unit-hour observations, 2016–2021. Border points are included, with a net electric efficiency of one. Categorical predictors (production type, bidding zone, technology, CHP) are not shown. Border points are assigned a net electric efficiency of one.

at €56.35/MWh, but with substantial variation from about €14 to several hundreds of euros. Lignite has the lowest average marginal cost at €42.25/MWh and very low spread. Its marginal cost remains below €50/MWh even at the 75th percentile. Hard coal is more expensive with an average generation costs of €53.94/MWh. The spread is low relative to natural gas, with most observations falling in the range of €40/MWh to €60/MWh. Finally, border points have the highest average marginal cost at €66.79/MWh and the largest variation. Occasionally, they even turn negative when the supplying bidding zone features no active fossil fuel units and a negative spot price for electricity.

3.3 Emissions

While we observe hourly electricity output of generating units, the associated emissions need to be estimated. We exploit the direct proportionality between the amount of fuel burned and the amount of CO_2 emitted to compute hourly CO_2 emissions at unit i as

$$e_i(y_i) = \frac{y_i \kappa_i}{\eta_i} \xi_i \tag{5}$$

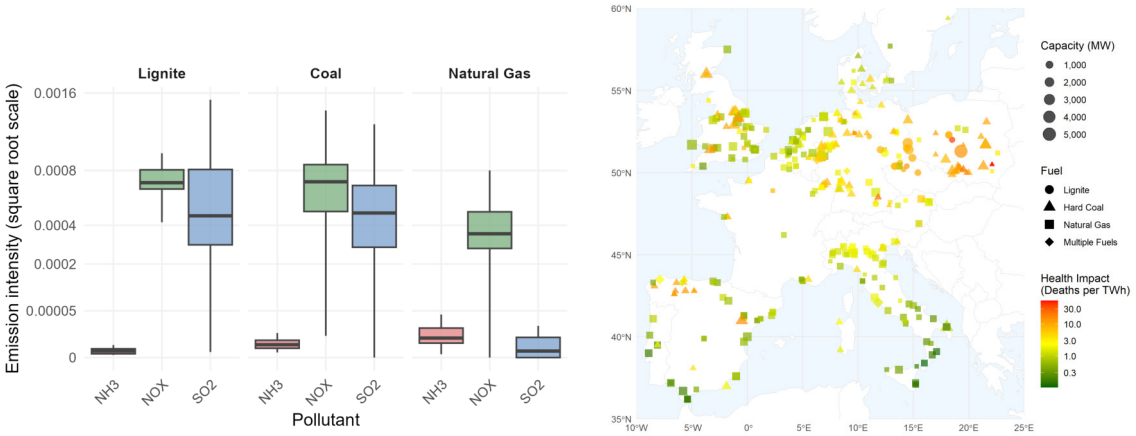
where y_i and κ_i are the hourly utilization rate and capacity, respectively, and η_i and ξ_i as defined in the previous subsection.

Similarly, we leverage proportionality of CO_2 and co-pollutant emissions to estimate unit-level emissions of NO_x, SO_2 and NH_3, the main precursors to fine dust ($PM_{2.5}$). While pollution emissions are not continuously measured in the European power sector,[10] annual data are available from the European Environment Agency (2025). Generating units with a total rated thermal input of at least 50 MW are legally obliged to report their emissions of SO_2, NO_x, and PM_{10} under the LCP and Industrial Emissions Directives (cf. Section 2 above). LCP reporting units—so-called *installation parts*—often correspond to one or a few generating units in ENTSO-E. LCP emissions are sometimes reported on the "smokestack"-level, which in some cases subsumes multiple units. In addition, data on annual releases of a broader set of air pollutants are available E-PRTR facilities, which may encompass anywhere from one to roughly a dozen installation parts. E-PRTR facilities are more aggregate and roughly correspond to power plants (not units) in ENTSO-E.

Appendix A.2 explains in detail how we match generating units in ENTSO-E to those in the air pollution database and compute annual co-pollutant intensities $\mathbf{z_{i,t}}$ for each of the generating units in our sample. Of the 732 generating units, 703 are matched to LCP installation parts. In 391 of these cases, the generating unit is the only unit linked to its installation part(s). We can then compute $z^k_{i,t}$ by dividing reported emissions of pollutant k in year t by our estimate of CO_2 emissions for that year. In the remaining 312 cases, multiple generating units are linked to the same LCP installation part. In those cases, we compute pollution intensities for aggregate generation and assign this measure to each generating unit. We apply the same procedure to the 27 units that can only be matched to LCP installations or E-PRTR facilities. Matching to E-PRTR facilities also allows us to compute co-pollutant intensities for air pollutants not reported in LCP data, notably ammonia (NH_3).

[10] In contrast, U.S. power plants are subject to a Continuous Emissions Monitoring System (CEMS).

Figure 2: Co-pollution emissions of electricity generation and mortality impacts



(a) Co-pollution intensities

(b) Mortality per TWh

Notes: Panel (a) shows emissions intensities of NH_3, NO_x and SO_2, in metric tonnes per tonne of CO_2 emitted. The boxplots represent the distribution across power units for the period 2016–2021. Emission intensities are reported on the generating unit × year level. Panel (b) displays the estimated impact of one TWh of electricity generated on premature deaths. Premature deaths per TWh of electricity generated, averaged across calendar months and power plant units, vary by two orders of magnitude, from 0.1 to 54.3. Different shapes denote different fuels used at the plant. Some plants combine units of several fuels. These are denoted as "Multiple Fuels".

Figure 2a displays the distribution of co-pollutant intensities among generating units in our sample, by fuel type. Across fuel types, we see that natural gas-fired plants are the least pollution intensive per tonne of CO_2 emitted. This is true for NO_x and particularly for SO_2 emissions. Because natural gas-fired generation is less CO_2 intensive than coal and lignite, gas-fired power plants are even cleaner per MWh of electricity produced. Moreover, measurable differences in co-emission intensities persist across plants with the same fuel type.

3.4 Mortality impacts

A well-known fact about air pollution is that the place of emission is not necessarily equal to the place of impact. Once pollutants are released from the smokestack, weather, topography, and pre-existing air pollution suspended in the local atmosphere influence their atmospheric dispersion and the formation of secondary pollutants such as ozone and $PM_{2.5}$. Therefore, mapping emissions at generating units in our sample to ambient pollution concentrations and population exposure requires us to use a chemical transport model (CTM), i.e. a computational tool capable of tracking the highly non-linear atmospheric transport and chemical reactions of air pollutants from the place of release to their final deposition.

We employ the nested GEOS-Chem Adjoint model developed by Gu et al. (2023). Based on primary pollution emissions of, among others, NO_x, SO_2, and NH_3, this CTM computes population exposure to both primary $PM_{2.5}$ (organic carbon, black carbon, and fine dust) and secondary $PM_{2.5}$ (e.g., sulfate, nitrate, ammonium) on a $0.01° \times 0.01°$ nested grid for the European domain. The CTM is calibrated to 2015 weather and emissions data, which aligns well with the beginning of our sample period.[11] Health impacts in a given grid cell are computed

[11]Since CTMs are computationally expensive, they are calibrated to a particular base year in which weather

as the number of $PM_{2.5}$-related premature deaths. Following the methodology proposed in the Global Burden of Disease Study (GBD) 2019 (Murray and *et al.*, 2020), this number is obtained by adjusting the age-stratified mortality within a cell for the $PM_{2.5}$ induced mortality among the "big six" causes of death (chronic obstructive pulmonary disease, ischemic heart disease, lung cancer, type-2 diabetes, lower respiratory infections, and stroke). We provide more detail on this calculation in Appendix B.

Importantly, the CTM's adjoint feature allows us to quantify the change in aggregate $PM_{2.5}$-related mortality induced by a power plant's change in precursor emissions. In a so-called adjoint run, the CTM computes partial derivatives of the total mortality count at each step of the computation. In a reverse application of the chain rule, these can be aggregated up to marginal effects for each pollutant and grid cell, referred to as adjoint sensitivities. To compute a unit's mortality impact $d_i = e_i \mathbf{z_i}' \boldsymbol{\lambda_i}$ we assign the vector of adjoint sensitivities $\boldsymbol{\lambda_j}$ in grid cell j to all units i located in the cell.

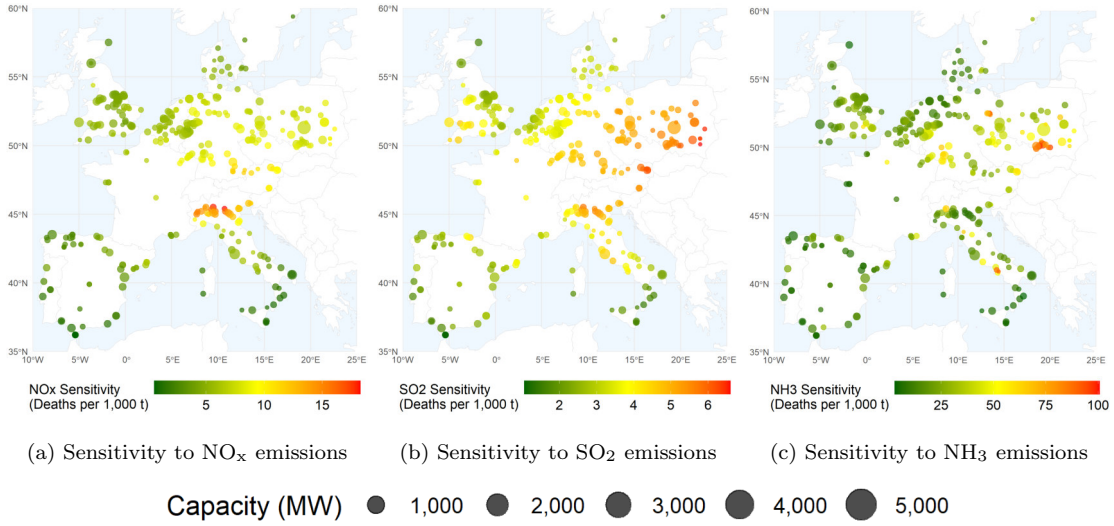
Adjoint sensitivities are extremely useful for our analysis with respect to source attribution (how many PM-related deaths can be attributed to a given power plant?), counterfactual analysis (how does PM-mortality change with the carbon price?), and computational efficiency (one adjoint model run suffices). Importantly, this convenience is not achieved by reducing the complexity of the underlying CTM, but rather by limiting the adjoint run to a scalar-valued loss function, the number of $PM_{2.5}$-related premature deaths.[12] The mortality impact of a given amount of pollution varies over the course of the year due to seasonal variation in weather and emissions inventories. Temperature, precipitation, and wind patterns affect mass transport, emissions from agriculture, transportation, and residential heating affect atmospheric chemistry. To account for this variability, we compute mortality impacts of electricity generation using adjoint sensitivities $\boldsymbol{\lambda}_i(m)$ that vary by calendar month m (Gu et al., 2023).

Figure 3 shows heat maps of co-pollutant sensitivities for the power plant locations in our sample. Depending on the place of release, an additional 1,000 metric tons of NO_x is associated with as little as half an additional premature deaths in Southern Spain up to more than 18 in Italy's Po valley (mean 6.4). The corresponding ranges for SO_2 lie between about one in Southern Spain and above six in Southern Poland (mean 3.7). NH_3 has the widest range, from about three in Southern Italy to one hundred in Southern Poland (mean 27.6), with the caveat that actual total NH_3 emissions are several orders of magnitude below NO_x and SO_2. Cross-sectional variation in sensitivities is driven by multiple factors. First, weather and topography both affect pollutant deposition and mass transport. Second, since $PM_{2.5}$ forms via chemical reactions during atmospheric transport, the amount and composition of pre-existing air pollution affects pollution concentrations. Third, population exposure to those concentrations, and hence mortality impacts, scale with population density. Variation across the maps arises largely from differences in pollutant lifetimes. NO_x is relatively short-lived compared to SO_2, which remains in the atmosphere much longer and can thus travel farther. This explains the stronger regional correlation of SO_2 sensitivities.

conditions are typical.

[12] Linearly approximating the effect of changes in emissions works well because grid cells comprise many pollution sources, not just power plants. Thus, even large changes in a plant's emissions would entail only small perturbations of the emissions emanating from the grid cell in which it is located.

Figure 3: Sensitivity of $PM_{2.5}$-related premature deaths to precursor emissions



(a) Sensitivity to NO_x emissions (b) Sensitivity to SO_2 emissions (c) Sensitivity to NH_3 emissions

Notes: The maps display predicted mortality impacts of NO_x, SO_2 and NH_3 released at the locations of power plants in our sample, in premature deaths per 1,000 tons emitted. Values are derived from monthly adjoint sensitivities to $PM_{2.5}$ precursor emissions, computed with the GEOS-Chem model (Gu et al., 2023). We aggregate those to the annual level by taking arithmetic means across calendar months at each plant location.

Figure 2b visualizes the cross-sectional variation in the mortality impact of electricity generation in our sample. We combine the average pollution intensities of generating units of each power plant with the marginal impact for the respective pollutant, averaged across calendar months, to obtain a measure of premature deaths due to $PM_{2.5}$ per TWh of electricity generated. Impacts computed in this way are lowest for natural gas plants on the Italian and Spanish coasts (as few as 0.1 deaths per TWh) and highest for Polish coal plants (between 8 and 17 deaths per TWh, and one extreme case rising to 54).

3.5 External costs of electricity generation

We compute the mortality cost of power generation due to air pollution in our sample as the product of premature deaths due to $PM_{2.5}$ exposure and the value of a statistical life (VSL),

$$TC_{AQ} = \sum_{t=1}^{T} \sum_{i=1}^{N} e_i\left(y_i(t)\right) \cdot \mathbf{z_i}' \boldsymbol{\lambda_i}(t) \cdot VSL(t). \tag{6}$$

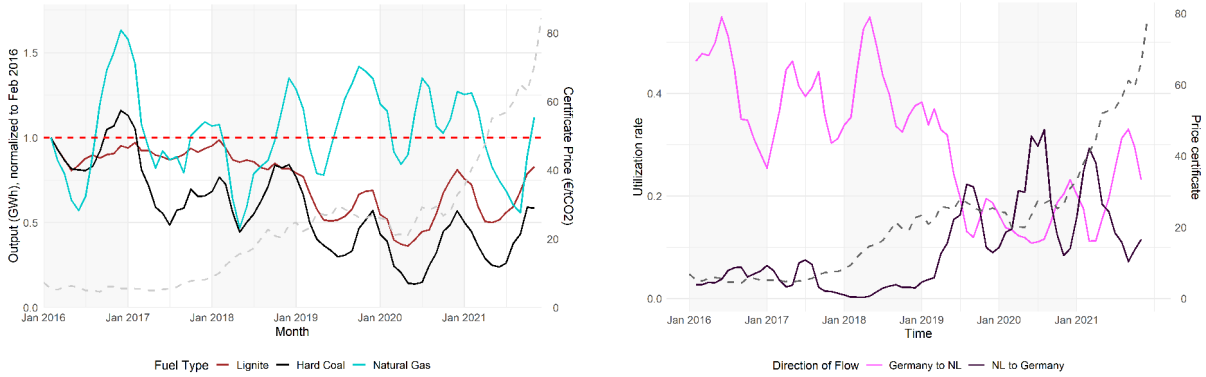
We use a conservative VSL estimate of 3.6 million in 2016 euros, in line with European Union guidelines on VSL in regulatory impact analysis (European Commission, Directorate-General for Mobility and Transport, 2019). We do not adjust VSL estimates for income differences across EU countries and regions.

We compute the total carbon cost of power generation in our sample as

$$TC_{CO2} = \sum_{t=1}^{T} \sum_{i=1}^{N} e_i\left(y_i(t)\right) \cdot SCC(t) \tag{7}$$

the product of CO_2 emitted and the social cost of carbon (SCC). We use an SCC estimate of €155 in 2016 euros (Rennert et al., 2022).

Figure 4: Changing generation patterns and carbon pricing



(a) Fuel switching within markets

(b) Relocation of generation between markets

Notes: The figure illustrates the empirical correlation between carbon prices and fuel switching within and across bidding zones. Panel a shows 3-month moving averages of generation by lignite, hard coal, and natural gas units in Germany. Panel b plots 3-month moving averages of the utilization rate of the Dutch-German border point.

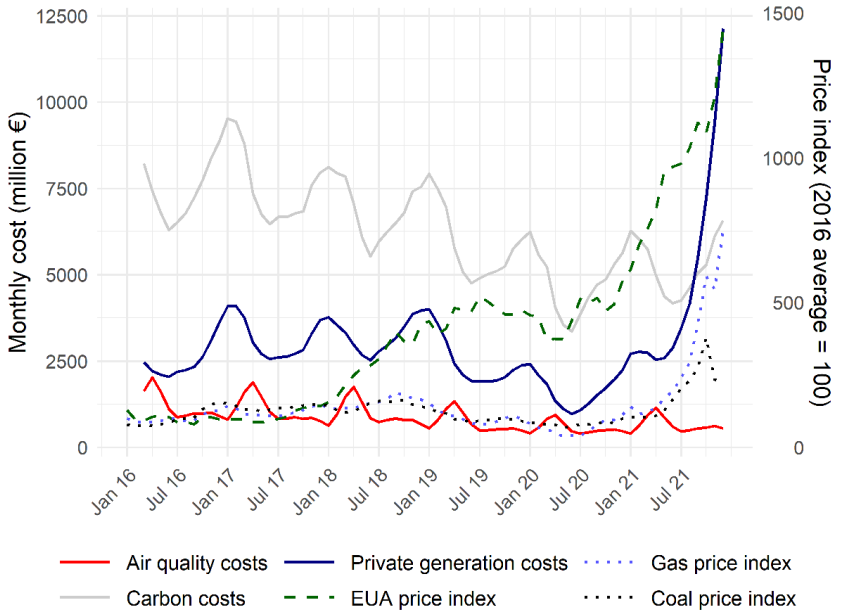
We estimate that external costs from power generation, both air quality and climate related, over the sample period amounted to €535 billion with TC_{AQ} = €63.8 billion and TC_{CO2} = €471.2 billion.

3.6 Trends in generation, prices, and externalities

The EUA price varied substantially over the sample period, rising from initial prices as low as €4 to over €80 at the end (cf. Table 2). Higher carbon prices provide an incentive to abate CO_2 by reducing the utilization rates of carbon-intensive generation in favor of less carbon-intensive generation. With integrated electricity markets, such fuel switching may occur both within and across bidding zones, as Figure 4 illustrates for the case of Germany. Panel (a) plots carbon prices over the sample period, alongside monthly output at German units by fuel type. One can see that natural gas-fired generation peaks in winter and bottoms out in summer. This pattern remains fairly stable over time, whereas generation at lignite and coal units follows a downward trend. Higher allowance prices in later years likely contribute to this trend by making carbon-intensive generation at coal-fired units more expensive relative to gas-fired generation. As a result, the share of natural-gas fired generation increases over time, consistent with coal-to-gas switching.

Fuel switching does not stop at the border, however. Within the limits given by physical transmission capacity, electricity generation can relocate between bidding zones to facilitate arbitrage based on relative fuel costs and installed capacity. As the allowance price rises, competitive advantage shifts from coal-heavy zones to gas-heavy zones and may turn the latter from net importers into net exporters of electricity. Figure 4b illustrates such reallocation effects by plotting the utilization rates of the Dutch-German border points (one for each direction) over the sample period. Early on, when the allowance price was below €20, the gas-heavy Dutch electricity zone exported very little to coal-heavy Germany but imported a lot. As the EUA price rose above €20 in the later years of the sample, Dutch exports to Germany increased markedly while imports from Germany declined.

Figure 5: Carbon prices and social costs of electricity generation, 2016-2021



Notes: The graph shows 3-month moving averages of several cost metrics. The left-hand y-axis shows monthly costs in millions of euros. The red line displays air quality costs, which is the monetary value of premature deaths from air pollutant emissions (premature deaths × VSL). The blue line displays private generation costs, which is the cost of fossil electricity generation net of carbon prices. The grey line is carbon costs, which is the monetary damage of carbon emissions (CO_2 emissions × social cost of carbon). The cost metrics are in constant 2016 Euros. The dashed line is the monthly EU ETS allowance price, with the corresponding indexed values on the right-hand y-axis. The dotted lines represent the prices of hard coal and natural gas.

Figure 5 plots the evolution of different cost metrics and average input prices (indexed to their 2016 average value) over the sample period. With respect to the EUA price, one can distinguish between three regimes; (i) a low-price period until early 2018 with EUA prices below €10; (ii) a mid-price period from early 2018 to late 2020 with prices between €20 and €30; and (iii) a high-price period in 2021 with sharply increasing EUA prices of up to almost €90. Prices of natural gas and coal both increased sharply towards the end of 2021.

Additionally, the figure plots the private and the external costs of electricity generation (in millions of €), separately for climate change TC_{CO2} and air quality TC_{AQ}. All three cost metrics spike in winter and bottom out in summer. Because renewable electricity output is lower in winter, fossil power plants step in and their CO_2 emissions drive up carbon costs. Private generation costs increase both at the intensive and the extensive margin, as some units enter the market that are too expensive to run in summer. Air quality costs also rise given as more active fossil units emit more air pollutants. In addition, adjoint sensitivities λ are higher in winter months when ambient pollution is higher (e.g., due to heating) and weather conditions aggravate air-pollution related health hazards.

Over time, the external costs of electricity generation follow a downward trend, consistent with the concurrent reduction of lignite and coal generation. The trend is particularly pronounced for carbon costs, but air quality costs are also trending downward and exhibit more muted winter peaks in later years. Private generation costs plummet in early 2020 with the outbreak of the Covid-19 pandemic. In late 2021, generation costs skyrocket as the prices

for both hard coal and natural gas quadrupled within a few months.

The apparent negative correlation between the trends in carbon prices and external costs related to air pollution and carbon emissions is consistent with a causal relationship, but not a proof thereof. Concurrent developments such as the rapid expansion of renewable electricity, regulatory policies targeting air pollution from the power sector, or mandated phase-out schedules for coal-fired generators may all have contributed to the patterns described above. To disentangle the causal impact of carbon pricing from such confounding effects, we develop a quantitative model of interconnected European electricity markets that allows us (i) to quantify how much of the observed carbon abatement and air quality improvements were driven by the EU ETS, and (ii) to estimate how those outcomes would have changed under counterfactual carbon prices or air pollution taxes.

4 Modeling electricity markets

4.1 Empirical merit-order curves

Under competitive bidding in the day-ahead market, fossil generating units offer electricity at their marginal cost. Renewable and nuclear units are dispatched first because of their near-zero short-run marginal costs. Many renewable installations have priority dispatch privileges.[13] Fossil units therefore serve net load, i.e., the part of total load not covered by these sources. Our discussion of the merit curve accordingly focuses on fossil units and net load.

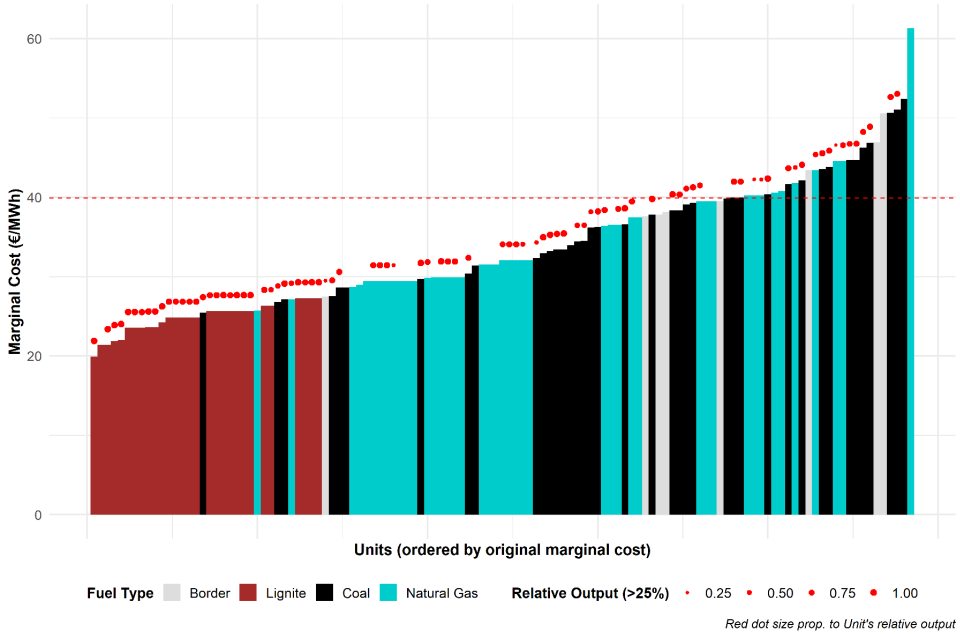
In the absence of cross-border transmission, knowledge of net load and marginal costs of all fossil units suffices to build the empirical merit curve and intersect it with net load to determine each bidding zone's autarky price as the marginal cost of the last unit activated to satisfy demand. However, differences in autarky prices between neighboring bidding zones invite price arbitrage via cross-border electricity trades. To account for this, we add border points to our sample and treat them as separate units. That is, we model electricity exports from bidding zone A to B as an additional power unit in bidding zone B and transmission from B to A as an additional power unit in bidding zone A. The marginal cost of zone A's border point is determined by the autarky price in zone B, and vice versa. If electricity in B was cheaper than in A without cross-border transmission, the border point in A would supply electricity because the market clearing mechanism would activate additional units in B to produce electricity for A and shut off some units in A that are more expensive than those in B.

Because autarky prices are unobserved, we impute them as the marginal cost of the last domestic unit needed to satisfy net load under the assumption that only domestic units are active.[14] We then reconstruct the merit curves, including the border points as virtual domestic units with their marginal cost according to the imputed autarky price in the zone they connect from. This treatment of cross-border flows closely follows what the European Market Coupling

[13]The Renewable Energy Directive (2009/28/EC) obliged system operators to grant priority dispatch to renewables, subject to system security. From 2020, Regulation (EU) 2019/943 restricted this to small and pre-existing installations. Germany's renewable energy law (EEG), for instance, guarantees priority feed-in for wind, solar, biomass, hydro and geothermal units.

[14]To be specific, the marginal unit is the one with the lowest marginal cost and a positive 'capacity gap', defined below.

Figure 6: An estimated merit-order curve for fossil power plants



Note: Estimated merit curve for September 5, 2016, at the 18th hour of the day, in the German-Austrian-Luxembourgian bidding zone. After normalizing capacity, units are ordered by their estimated marginal costs in €/MWh, plotted on the y-axis. A red dot indicates that the unit was active, with a size proportional to the utilization rate. The dashed red line indicates the equilibrium price on the day-ahead market.

Company actually calculates to determine equilibrium on the day-ahead market.

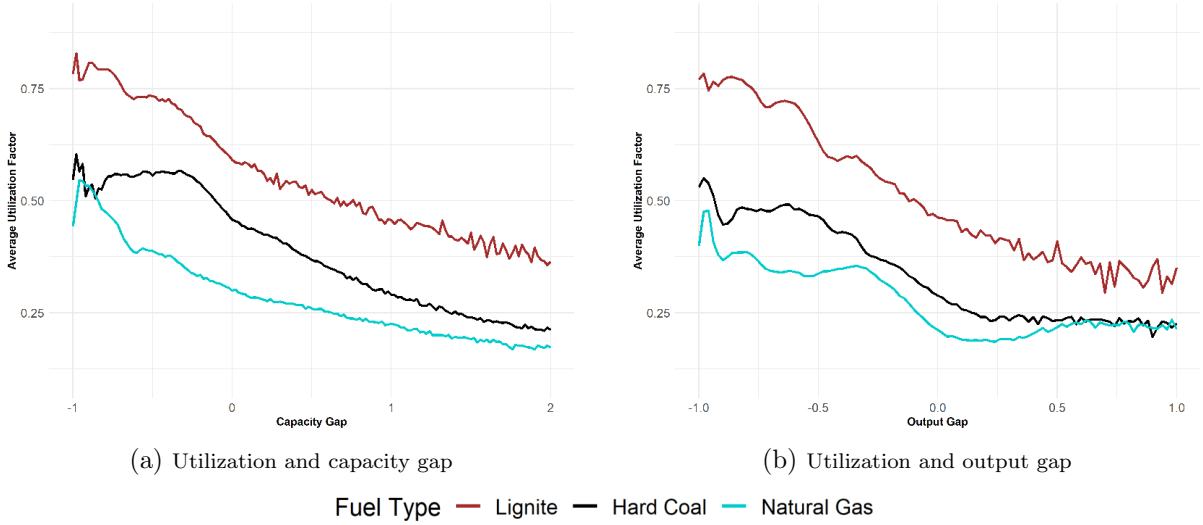
Figure 6 displays an estimated merit curve for the German-Austrian-Luxembourgian bidding zone on September 5, 2016 at the 18^{th} hour of the day for units running on either lignite, hard coal, or natural gas, as well as border points representing electricity supply from neighboring bidding zones. The units are in ascending order according to their marginal costs (in €/MWh, on the vertical axis). The size of the red dots indicate the unit's utilization rate; no dot indicates the unit was below its minimum stable load. Because the EUA price was at only €4.09 on that day, lignite units were the cheapest fossil fuel units in the market and operated near full capacity. Next, there is a block of highly efficient natural gas units at estimated marginal costs of around €30/MWh. In spite of being 'in the money', not all of those units are active. Conversely, some of the most costly generators at the upper end of the merit curve are active although their estimated marginal cost exceeds the equilibrium day-ahead price.

For a systematic assessment of how closely the merit-order principle aligns with the millions of hourly dispatch decisions in our dataset, we examine empirical relationships between a unit's utilization rate and its position in the merit-order curve. To do so, we define the *output gap* as

$$O_i^y(t) = \frac{\sum_{\{j|MC_j \leq MC_i\}} y_j(t)\kappa_j - L(t)}{L(t)} \tag{8}$$

where y_j and κ_j are the hourly utilization rate and capacity, respectively, and bidding zone subscripts have been omitted. This measures the gap in net load that a unit would have to fill in order to be on the margin, given its marginal costs, net load, and the supply of other units with lower costs. The normalization makes sure that $O_i^y(t) \geq -1 \;\; \forall t : L(t) \geq 0$. As a close

Figure 7: Utilization vs. position in the merit curve



(a) Utilization and capacity gap

(b) Utilization and output gap

Note: The figures plot the average utilization rates of generating units against their relative position in the merit curve, measured as the unit's capacity gap (a) or output gap (b), as defined in eq. (8). A negative gap indicates that net load is larger than cumulative capacity (a) or generation (b) of all units with lower marginal costs, so that the unit is profitable. A positive gap indicates that net load is not large enough for the unit to be run profitably. A zero gap indicates that the unit is at the margin.

analogue, we define the *capacity gap* $O_i^{\kappa}(t)$ by substituting $y_j(t) = 1$ in the above equation.

Both variables measure unit i's position in the merit curve relative to net load, indicating whether the unit is "in the money" and how far it is from being marginal. A unit with a negative gap should be active, as the aggregate capacity or generation of all units with lower marginal costs falls short of net load, whereas a unit with a positive gap should not be. As we move up the merit curve, both gaps decrease in absolute terms and eventually switch sign from negative to positive once we cross the intersection with net load.[15]

Based on actual output rather than installed capacity, O_i^{y} is a more flexible measure of the true merit order than O_i^{κ}. However, it is also endogenous to actual dispatch decisions which makes it more challenging to use in counterfactual analysis. In contrast, the capacity gap is determined by exogenous factors only (net load, input prices, and capacity).[16]

Figure 7 plots how the utilization rate varies with the capacity gap (a) and the output gap (b), differentiated by fuel types lignite, hard coal, and natural gas. Across all fuel types, the utilization rate is declining in either gap variable. This is in line with the merit-order principle because units with higher values of the gap are more expensive relative to other units. However, the utilization rate does not drop all the way to zero, even for very positive gaps. Such deviations from merit-order driven dispatch decisions in our data could be due to several factors.

First, the merit-order principle assumes that any unit can satisfy electricity demand at any location within the bidding zone. However, if transmission constraints matter also within bidding zones, balancing demand and supply requires deviations from merit-order dispatch.

[15] When non-dispatchable supply exceeds total load, net load becomes negative. Under these exceptional circumstances, no dispatchable unit should run.

[16] If discrepancies between capacity and actual output are approximately constant across bidding zones and across different market conditions, the capacity gap will yield an accurate ordering of units' propensity to generate in a given hour.

For example, a local demand spike may only be served by local plants if transmission from cheaper yet more distant plants is constrained. Therefore, actual generation patterns may differ substantially from generation predicted by empirical merit curves.

Second, the merit curve reflects what suppliers bid on the day-ahead market, based on expectations about load and availability of units. It ignores intraday trading and short-term re-dispatch. Realized load may exceed predictions, or suppliers may be unable to deliver on their day-ahead bids, for instance due to forced outages. In either case, operators of dispatchable units can bid extra capacity into the intraday market, allowing for last-minute adjustments to balance realized supply with realized demand.

Third, plant operation schedules must allow time for ramping up and down power plants, which may take several hours. This is particularly true for plants burning lignite, a fuel that cannot be easily put out and lit back on. Going through a full ramping cycle with a lignite unit may be less economical or infeasible compared to letting it run through a few hours of adverse market conditions. In fact, lignite units tend to have the highest utilization rates (more than 70%) and to operate near their full capacity when in the money.

To sum up, our investigation of empirical merit curves highlights the need for a unified empirical framework capable of integrating the merit-order principle and systematic deviations from it.

4.2 A random-forest model of electricity dispatch

We propose an empirical model of European electricity markets that captures both merit-order driven dispatch and the limitations to it, which derive from transmission constraints within and between bidding zones, re-dispatch, intermittent renewable energy supply, and other sources of congestion. Our approach is to embed the structural determinants in a machine learning model capable of predicting hourly electricity generation at each unit.

The exposition of the model is subdivided in three subsections. We start by presenting our model and assessing its predictive power. Next, we introduce a novel search algorithm that combines model predictions into an approximated counterfactual equilibrium that clears all markets simultaneously. Finally, we state and discuss identifying assumptions that allow us to interpret the difference between predicted outcomes for different carbon prices as causal effects.

4.2.1 Quantile Regression Random Forest

We use a Quantile Regression Random Forest (QRRF) to predict the hourly utilization rate at fossil fuel powered units as well as at border points (Meinshausen, 2006). In machine learning, a forest is a collection of individual decision trees. To illustrate the concept, suppose we have training data with outcome y that we want to predict using K predictors $\mathbf{X} = (X_1, X_2, \ldots, X_K)^T$. A decision tree takes all observed outcomes y in the training data and successively assigns them to an increasing number of ever smaller bins. This is done by repeatedly splitting up the current number of bins along the lines of the predictors in $\mathbf{X}$.[17] At

[17] As a simple example, assume that the only available predictor is demand and we can only split once. Suppose further that unit utilization rates and demand are positively correlated in the training data. A logical split groups low demand with low utilization rates in one bin and high demand with high utilization rates in the other. When

each node, the algorithm considers the candidate predictors in $\mathbf{X}$ and splits on the one that yields the largest reduction in the variance of the outcome. In other words, it picks the predictor that best explains the outcomes of the training observations in that node. The algorithm continues to implement splits according to this criterion until a pre-specified stopping criterion becomes binding, usually a lower bound on the number of observations in each terminal node of the tree.

Relying on a single decision tree is prone to high-variance predictions because small changes in the data may lead to a very different tree. We therefore construct an entire forest of prediction trees. All of these differ from one another because we randomly restrict the set of predictors available to split on at any given node. The forest significantly reduces variance without introducing bias. Using training data from our sample, we estimate a random forest with $S = 100$ trees. Each tree s is grown by recursive binary splitting, where the split at each node is chosen to minimize the within-node variance of y in the resulting child nodes.

Now suppose we observe $\mathbf{X_i}$ but not y_i, as would be the case in our counterfactuals when we want to predict generation under alternative carbon prices. To obtain a model prediction $\widehat{y}_i$ based on $\mathbf{X_i}$, we drop $\mathbf{X_i}$ down every tree s, where it falls into a terminal node $\iota_s = \iota(\mathbf{X_i}, s)$ which contains a set of observed outcomes from the training data, $\{y_{\iota_s,s}\}$. We obtain S such sets for the random forest. Their union $\Upsilon_i = \{y_{\iota_1,1}\} \uplus \{y_{\iota_2,2}\} \uplus ... \uplus \{y_{\iota_S,S}\}$ contains all training-data outcomes y from all selected terminal nodes $\iota_1, ..., \iota_S$. One can think of the elements in Υ_i as candidates for predicting $\widehat{y}_i$ given $\mathbf{X_i}$. An obvious choice is the mean value of Υ_i. However, since the algorithm recovers the entire conditional distribution of possible predictions instead of just the mean value, the QRRF allows us to calculate the predicted outcome $\widehat{y}_i(q)$ at any given percentile q as

$$\widehat{y}_i(q) = \inf \left\{ y \in \Upsilon_i : \frac{1}{|\Upsilon_i|} \sum_{z \in \Upsilon_i} I(z \leq y) \geq q \right\}. \tag{9}$$

4.2.2 Estimation setup

Estimating the QRRF requires training data and several calibration choices. Our training data comprises the hourly utilization rate y and predictors $\mathbf{X}$ for all units observed in the 13 countries. During the six-year sample period, this amounts to more than 28 million observations. We oversample at the extremes to improve the accuracy of the model predictions when units generate either near full capacity or near zero.

We include 20 predictors that can be classified into four broad categories: (i) information relevant for merit-order dispatch in the day-ahead market such as net load (as a proxy for the day-ahead forecast relevant for dispatch), estimated marginal cost, and three variables relating to a unit's position in the supply curve relative to all other units and to the intersection with demand (output and capacity gaps, and a dummy for negative capacity gap); (ii) contemporaneous fuel prices to account for the fact that units might additionally bid in the intraday market, responding to sudden overnight price changes; (iii) unit characteristics such as fuel type (lignite, coal, gas, or border point), net electric efficiency, installed technology, a

evaluating an external observation with known demand but unknown utilization rate, the tree maps it to the corresponding bin. If demand is high, the model predicts a high utilization rate because the training data shows that units are typically active under those conditions, and vice versa.

CHP dummy, longitude and latitude, and a dummy for temporarily reduced capacity, and (iv) dummies for the bidding zone, hour-of-day, day-of-week, and month-of-year.

We estimate 100 trees with the restriction that terminal nodes must not contain less than ten data points. At each node, splits are chosen to maximize the reduction in impurity, measured as the decrease in within-node variance of the outcome achieved by partitioning the data according to a given predictor and threshold.

4.2.3 Market clearing and percentile choice

For any given $\mathbf{X_i}$, the QRRF provides a set Υ_i from which the function $\widehat{y}_i(q)$ picks an element $\widehat{y}_i$ as the prediction at percentile q. Following Jarvis et al. (2022), we choose the values of q such that predicted power generation satisfies demand. Conceptually, we could implement this market clearing condition by choosing percentiles q_i^* at every unit×hour such that $\widehat{y}_i(q_i^*) = y_i$, yet this is computationally prohibitive.[18] To circumvent this issue, we calculate utilization rates $\widehat{y}_i\,(q)$ at nine deciles $q \in \{0.1, 0.2, ..., 0.9\}$ and multiply them with the unit's installed electric capacity κ_i to obtain corresponding predictions for output in MWh,

$$\widehat{Y}_i\,(q) = \widehat{y}_i\,(q) \cdot \kappa_i \quad \forall q \in \{0.1, 0.2, ..., 0.9\}. \tag{10}$$

We aggregate predicted output to the level of bidding zone b by fuel f by month-of-sample m,

$$\widehat{Y}_{b,f,m}(q) = \sum_{i \in I_{b,f,m}} \widehat{Y}_i(q) \quad \forall q \in \{0.1, 0.2, ..., 0.9\}, \tag{11}$$

and search for the percentiles that equate these predictions to observed output $Y_{b,f,m}\,\forall b, f, m$. To this end, we identify the decile predictions that enclose observed generation and interpolate between them to match predicted and actual generation. Algebraically, for each zone, fuel and month-of-sample, we find a decile $q_{b,f,m}$ and a corresponding step length $\alpha_{b,f,m} \in (0, 1)$ such that

$$\widehat{Y}_{b,f,m}(q_{(b,f,m)}) < Y_{b,f,m} < \widehat{Y}_{b,f,m}(q_{(b,f,m)} + 0.1) \quad \wedge \tag{12}$$

$$\widehat{Y}_{b,f,m}(q_{b,f,m}) \;+\; \alpha_{b,f,m} \cdot \Big[\widehat{Y}_{b,f,m}(q_{b,f,m} + 0.1) - \widehat{Y}_{b,f,m}(q_{b,f,m})\Big] = Y_{b,f,m}. \tag{13}$$

Provided that $q_{b,f,m}$ exists, linear interpolation works because predicted utilization rates monotonically increase with q. We approximate the percentiles that establish a market equilibrium as $q^*_{b,f,m} = q_{b,f,m} + 0.1\alpha_{b,f,m}$ and collect these across zones B, fuel types F, and months M in the set $\mathbf{Q}^* = \{q^*_{b,f,m} \,|\, b \in B, f \in F, m \in M\}$.[19]

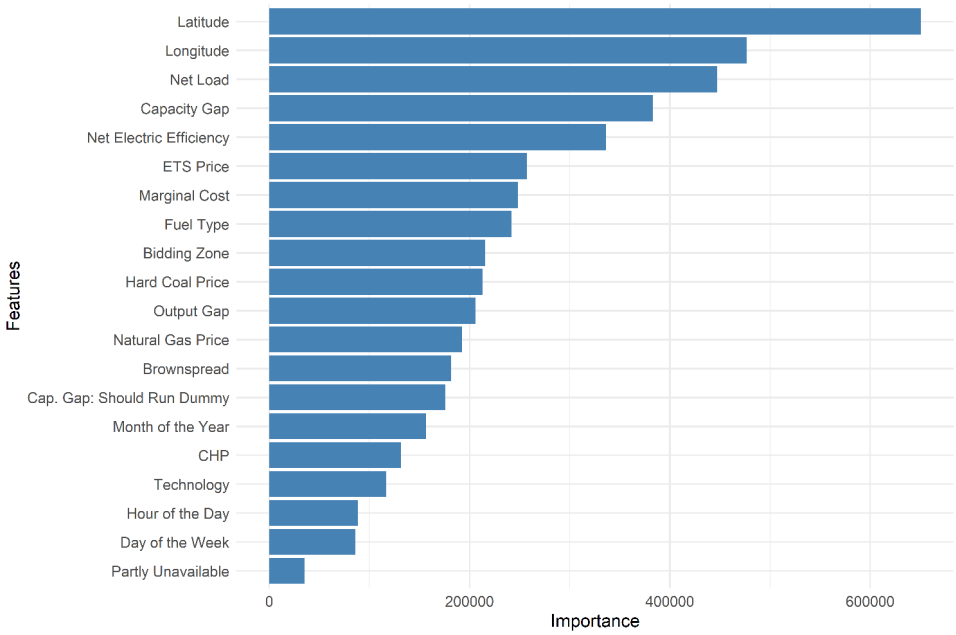
4.2.4 Predictive quality of the random forest

We evaluate the predictive quality of the QRRF at the percentiles $\mathbf{Q}^*$ chosen to establish market clearing. Figure 8 provides an overview of the predictors in our model, ranked in

[18] In our training data, equating predicted and actual utilization for every observation would require us to find and calculate more than 28 million individual values for q_i^*.

[19] Appendix Figure C.3 provides an example to illustrate how we pick the percentile that establishes the market equilibrium.

Figure 8: Relative importance score of predictors



Notes: The figure shows the importance of the predictors in the QRRF. We use the "impurity" measure for importance scores, which measures the total reduction in variance (impurity) at all splits on the respective predictor across all trees in the random forest. The impurity measure is a unitless score and makes sense only relative to other predictors. For instance, a unit's latitude is the most important single predictor with an importance score of just below 650,000. Longitude, the second-most important single predictor, has a score of about 477,000. The least important predictor, the dummy indicating that a unit is partly unavailable, has an importance score of about 35,000.

descending order by their predictive power. A unit's geographic position in terms of latitude and longitude has the highest predictive power, contributing to about 22.4% of all variance reduction in predicted utilization rates. To the extent that there is unobserved heterogeneity across units within a bidding zone, the model might load this on latitude and longitude, akin to a unit fixed effect in conventional regression models. Within-zone transmission constraints are a likely source of such heterogeneity because a congested electricity grid will require substantial deviations from merit-order dispatch. If transmission within bidding zones were as if on a copperplate, the geographic location of a unit in a market should be irrelevant.

Net load is highly predictive because higher net load predicts higher utilization, and so is net electric efficiency because efficient units can generate at a lower cost per MWh than inefficient units. The capacity gap measures a unit's position in the merit order, relative to all other units and relative to the supply curve's intersection with demand.

All counterfactual predictions in this paper are based on a QRRF trained on the full sample of more than 28 million observations. To assess out-of-sample performance, we randomly partition the 52,608 hours-of-sample (24 hours $\times$ 365 days $\times$ 6 years, plus two leap days in 2016 and 2020) into one training and one test set of equal size. Following standard practice, we re-estimate the model on the training data, and then assess its predictive quality on the out-of-sample test data.

Overall, our model fares very well in terms of goodness-of-fit measures. The MSE of the predicted utilization rate is 0.040, with an R^2 of 0.729 and a correlation coefficient of 0.864.

When partitioning by country, the model tends to perform better in smaller than in larger countries (in terms of the size of their fossil fuel fleet). When partitioning by fuel type, the model works particularly well for natural gas and border points, which are the most flexible and can be ramped up and down from one hour to the next without major constraints. Conversely, the predictions are somewhat less accurate for lignite generation, which tends to be sluggish and takes several hours to fully ramp up or down, so that the margin of error is plausibly larger in a model that independently predicts hour-by-hour generation. The interested reader is referred to Appendix C.1 for more comprehensive information on this exercise.

4.3 Counterfactual equilibrium

The counterfactual equilibria we consider differ from the observed one in terms of (i) the generation mix within every zone and (ii) the allocation of generation across zones. For example, with a high counterfactual carbon price, generation within a given zone may shift from coal-fired units towards natural-gas-fired units. In addition, the zone might increase its net imports from a neighboring zone, which in turn increases its gas-fired generation to export more electricity. Since bidding zones are interconnected, relocation effects in a counterfactual equilibrium may extend well beyond neighboring zones and thus amplify or dampen fuel switching compared to autarky.

To be able to pin down such an equilibrium, we assume that net load, i.e. residual demand, in each zone is unchanged in the counterfactual. What changes is how that residual demand is met, through a different fuel mix among domestic plants and a different split between domestic generation and net imports. Holding counterfactual net load constant rests on two implicit assumptions. First, in line with a large strand of research on electricity markets, demand is perfectly price inelastic. That is, in case the electricity price changes in a counterfactual scenario, we assume that consumers do not adjust their demand in response to this price change. Second, we assume that generation levels of non-fossil power units do not change. This departs from other research that has explicitly considered counterfactual changes to non-fossil generation, in particular from renewable sources.[20] We believe this is a tenable assumption because, unlike a structural shock to the entire grid, the EU ETS changes electricity prices only gradually via a market-based scarcity signal that is too muted and volatile to change long-term investments in the power plant fleet on a broad scale. This assumption is not obvious for nuclear power plants in France, which have been constructed to be dispatchable to a degree. In Appendix E.6, we directly assess the responsiveness of French nuclear power units to changing market conditions. We find that they respond to changes in domestic and foreign residual load, which we hold constant across our scenarios, but not to the carbon price. We therefore treat French nuclear generation as non-dispatchable and net it out of load.

[20] In their analysis of Germany's nuclear exit after 2011, Jarvis et al. (2022) assume that counterfactual generation from renewable sources would have been lower, because keeping nuclear energy in the fuel mix would have resulted in lower electricity prices, attenuating investments in renewable energies.

4.4 Computing counterfactual equilibria

Given net load $L_{b,m}$ in zone b and month m, denote by $Y_{b,f,m} \equiv \sum_j \sum_{t \in m} Y_i(t)$ the aggregate generation from all domestic fossil fuel units j running on fuel f. Market clearing requires that

$$\forall m \; \forall b: \; L_{b,m} = \sum_f Y_{b,f,m} + IM_{b,m} - EX_{b,m} \tag{14}$$

where $IM_{b,m}$ and $EX_{b,m}$ denote electricity imports and exports of zone b, respectively. When predicting generation under *actual* carbon prices, this equation holds by construction of the percentiles $q^*_{b,f,m} \in \mathbf{Q}^*$

$$\forall m \; \forall b: \; L_{b,m} = \sum_f \widehat{Y}_{b,f,m}\left(q^*_{b,f,m}\right) + \widehat{IM}_{b,m}\left(\mathbf{Q}^*\right) - \widehat{EX}_{b,m}\left(\mathbf{Q}^*\right). \tag{15}$$

With a counterfactual EUA price, this is no longer true because the marginal cost of power units and their ordering in the supply curve will change. Given a vector of counterfactual covariates $\mathbf{X_i^c}$, the QRRF model generates a set of candidate predictions Υ_i^c for counterfactual utilization y_i^c at unit i. The qth percentile is defined as

$$\widehat{y_i}^c(q) = \inf\left\{ y \in \Upsilon_i^c : \frac{1}{|\Upsilon_i^c|} \sum_{z \in \Upsilon_i^c} I(z \leq y) \geq q \right\}. \tag{16}$$

The choice of q pins down the level of counterfactual generation. In the counterfactual market equilibrium, the (a priori unknown) set of percentiles $\mathbf{Q}^{*\mathbf{c}} = \{q^{*c}_{b,f,m}\}$ induces markets to simultaneously clear across all zones b:[21]

$$\forall m \; \forall b: \quad L^c_{b,m} = L_{b,m} = \sum_f \widehat{Y}^c_{b,f,m}(q^{*\,c}_{b,f,m}) \; + \; \widehat{IM}^c_{b,m}(\mathbf{Q}^{*\mathbf{c}}) \; - \; \widehat{EX}^c_{b,m}(\mathbf{Q}^{*\mathbf{c}}) \tag{17}$$

Equation (17) makes explicit our assumption to keep net load constant across scenarios, $L^c_{b,m} \stackrel{!}{=} L_{b,m}$. It characterizes the counterfactual equilibrium as a new combination of domestic fossil generation and net imports that adds up to the same net load in each zone. Because bidding zones are interconnected, their equilibrium choices of generation and net imports must be mutually consistent.

Computing the percentiles $\mathbf{Q}^{*\mathbf{c}}$ that induce the counterfactual equilibrium amounts to finding a fixed point of eq. (17) for every month-of-sample m. To do so, we devise a numerical search algorithm that mimics the two mechanisms at play: domestic fuel switching and cross-border reallocation. Using $\mathbf{Q}^*$ as our starting values, we first predict domestic dispatchable, fossil-fuel fired generation as

$$\widehat{Y}_{b,m}\left(q^*_{b,\cdot,m}\right) = \sum_f \widehat{Y}^c_{b,f,m}\left(q^*_{b,f,m}\right)$$

[21] Technically, market clearing must hold in every hour-of-sample, but computing this is computationally infeasible, as was noted before. Our equilibrium concept tolerates hour-by-hour differences between supply and demand as long as they balance out over the course of one month.

Interpolating as in Section 4.2.3, we then identify scaling factors $\sigma_{b,m}$ for the percentiles at the zone $\times$ month level such that predicted domestic dispatchable generation matches the one observed in the data for each zone b and month m:

$$\widehat{Y}_{b,m}\left(\sigma_{b,m} \cdot q^*_{b,f,m}\right) = \sum_f \widehat{Y}^c_{b,f,m}\left(\sigma_{b,m} \cdot q^*_{b,f,m}\right) = Y_{b,m} \quad \forall (b,m)$$

This step allows for changes in the fuel mix due to counterfactual EUA prices but keeps total domestic dispatchable generation constant. We collect these percentiles in

$$\mathbf{Q^{int}} = \{q^{int}_{b,f,m} = \sigma_{b,m} \cdot q^*_{b,f,m} | \, b \in B, f \in F, m \in M\}$$

and refer to them as an intermediate solution that does not yet include equilibrium adjustments of cross-border electricity trade.

We approximate those adjustments in the next step. Analogously to residual demand L, we define residual supply as the sum of predicted domestic dispatchable generation and net imports:

$$\widehat{RS}(Q) = \widehat{Y}(Q) + \widehat{IM(Q)} - \widehat{EX}(Q) \tag{18}$$

In the counterfactual equilibrium, predicted residual supply must equal residual demand.[22] The loss function

$$\Lambda_m(Q) = \sum_b \left(L_{b,m} - \widehat{RS}_{b,m}(Q)\right)^2 \tag{19}$$

measures the distance to that equilibrium as the sum of squared deviations between L and $\widehat{RS}$ across zones in a given month m.

Our numerical search for the counterfactual equilibrium amounts to minimizing Λ. We conduct a directed grid search that starts at initial percentiles $\mathbf{Q^{int}}$ and scales them up or down by ϵ percent, one bidding zone at a time.[23] This can be repeated across all bidding zones until Λ_m approaches zero. In practice, we find that three iterations are computationally feasible and sufficient to almost entirely close the gap between L and $\widehat{RS}$. As shown in Appendix Figure C.4, predicted residual supply in the intermediate solution differs from residual demand by on average 16.9% across zone $\times$ month pairs. Three search iterations (using $\epsilon = 1$ in the first two and $\epsilon = 0.1$ in the last round) reduce this gap to on average 2.6%. The numerical iteration towards the counterfactual equilibrium works best in zones with many fossil fuel units and a high degree of interconnectedness. This is reassuring as those are also the zones that matter most for aggregate emissions and health outcomes. L and $\widehat{RS}$ differ by on average less than 1.4% in Germany, Poland, Great Britain, the Netherlands, and the Northern Italian bidding zone. In contrast, the algorithm does not perform as well in small zones with little interconnectedness, or in zones with very little fossil generation, such as Eastern Denmark, Sardinia, or Sweden.

[22] This is not true for the intermediate solution because it does not account for cross-border flows, i.e. $L_{b,m} \neq \widehat{RS}_{b,m}\left(\mathbf{Q^{int}}\right) \forall b$.

[23] For example, starting with zone b_1, we compare $\Lambda_m(q_{b,f,m})$ to $\Lambda_m((1-\epsilon/100) \cdot q_{b_1,f,m}, \{q_{b,f,m}\}_{b \neq b_1})$ and $\Lambda_m((1+\epsilon/100) \cdot q_{b_1,f,m}, \{q_{b,f,m}\}_{b \neq b_1})$ for small $\epsilon > 0$. If one of the scaled Λs is smaller, we continue looking towards that direction, so we check $1-2\epsilon/100$ next if $\Lambda_m(1-\epsilon/100) \cdot q_{b_1,f,m}, \{q_{b,f,m}\}_{b \neq b_1}) < \Lambda_m(q_{b,f,m})$, or check $1+2\epsilon/100$ next if $\Lambda_m(1+\epsilon/100) \cdot q_{b_1,f,m}, \{q_{b,f,m}\}_{b \neq b_1}) < \Lambda_m(q_{b,f,m})$.

4.5 Identifying assumptions

In this section, we summarize the identifying assumptions on which our empirical approach and model predictions rely. A more detailed discussion of potential threats to identification and how our model addresses them is relegated to Appendix E.

First, we take fleet composition of the power plant fleet to be exogenous, in line with our goal to identify short-run dispatch responses to carbon pricing. In the longer term, higher carbon prices may accelerate the shutdown of coal-fired plants and drive entry of low-carbon generating units. If this happened already during our analysis period, the ETS-induced fuel switching is larger than what we estimate and would lead us to underestimate the associated reductions in co-pollution emissions.

Second, there are no time-varying confounding regulatory policies on the state level that would fundamentally change the hour-by-hour operating decisions of the existing power plant fleet. We argue that this is plausible during our analysis period. Structural policies such as national coal phase-out trajectories affect fleet composition through plant entry and exit, but do not change hour-by-hour dispatch decisions. Plant entry and exit only influence our predictions through the merit order, which the model handles directly, and redispatch opportunities, which the model learns from the training data. In other words, we rule out policies that would alter the hour-by-hour relationship between carbon prices and dispatch decisions for units already in the fleet.

Third, we follow previous work such as Grossi et al. (2017), Abrell et al. (2022) and Jarvis et al. (2022) and take fuel prices as exogenous because these are largely determined by fixed supply frameworks and global benchmarks. For instance, oil-indexed contracts play an important role in setting natural gas prices, and oil-fired power generation is only marginal in continental Europe and Great Britain.

Fourth, in line with the previous literature, we assume that wholesale markets are competitive and power plant operators do not exercise market power. Specifically, we estimate supply curves based on marginal costs without mark-up. To support this assumption, we show in Appendix E that (i) our predictions are not systematically biased at certain carbon price levels; (ii) on the unit-level, the model prediction error is largely uncorrelated with the allowance price; and (iii) the model's MSE and R^2 are stable year-on-year across the sample, spanning periods of low, medium, and high allowance prices.

Fifth, we abstract from a possible influence of carbon prices on short-term re-dispatch and model a direct connection between bidding in the day-ahead market and actual realized dispatch and generation. We do not explicitly model market frictions such as within-zone transmission constraints. In Appendix E, we explain how our Random Forest can plausibly account for and correct these disturbances by including adequate covariates.

5 Results

This section presents predicted electricity market outcomes under different assumptions about the carbon price. Our baseline scenario assumes a zero carbon price, $\tilde{P}_t^c = 0 \, \forall t$. The ETS scenario assumes that EUA prices follow their actual path $\{P_t^c\}_{t=1}^T$. The price floor

Table 3: Generation [TWh] by type, cumulative for 2016-2021

Type	Lignite	Hard Coal	Natural Gas	Total	Imports
Zero EUA price [€ 0]	952.4	1,383.2	1,824.5	4,160.2	1,415.9
Actual EUA price $[P_t^C]$	924.5	1,115.5	2,119.5	4,159.5	1,635.1
% Change vs € 0	-2.93%	-19.35%	+16.17%	-0.02%	+15.48%
Price floor $\left[\max(P_t^C, 50)\right]$	788.1	855.1	2,516.1	4,159.3	1,750.8
% Change vs € 0	-17.25%	-38.18%	+37.91%	-0.02%	+23.65%
Pollution tax	804.4	1,274.5	2,080.4	4,159.4	1,575.4
% Change vs € 0	-15.54%	-7.86%	+14.03%	-0.02%	+11.26%
% Change vs Actual	-12.99%	+14.25%	-1.84%	-0.00%	-3.65%

Notes: Aggregate generation in terawatt-hours (TWh) by type of plant of all units in the analysis sample from 2016 to 2021. We compare allowance price counterfactuals. Row (i) contains predicted generation in a € 0 allowance price counterfactual scenario. Row (ii) contains observed generation under actual allowance prices. Row (iii) contains predicted generation in a € 50 allowance price floor counterfactual scenario. Row (iv) contains predicted generation under a Pigouvian tax on air pollutant emissions and zero EUA price.

scenario assumes that EUA prices follow their actual path but never drop below € 50, $\tilde{\tilde{P}}_t^c = \max\{P_t^c, 50\}\ \forall t$.

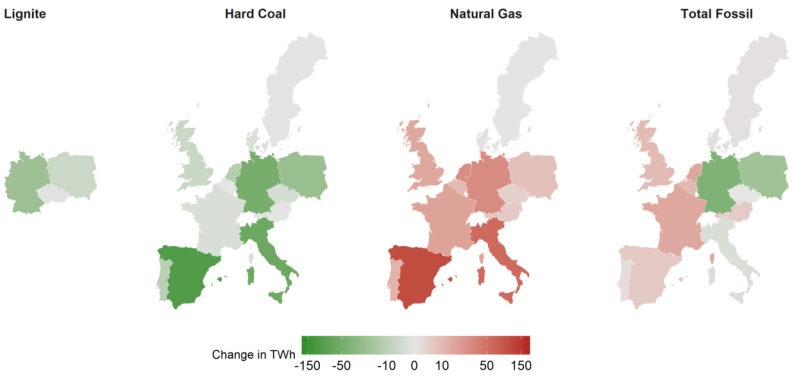
5.1 Electricity generation

In the six-year period from 2016 to 2021, the generating units in our sample generated 4,160 TWh of electricity. Holding aggregate output constant at this quantity, Table 3 reports how the allocation across fossil sources changes for different carbon pricing scenarios. With a zero EUA price, natural gas-fired units account for 43.9%, followed by hard coal and lignite with 33.2% and 22.9% of total electricity output, respectively. Electricity trades between bidding zones are significant and make for one third of total fossil generation. Unlike renewable power sources, fossil-fuel fired units are dispatchable and thus frequently generate electricity for exports.

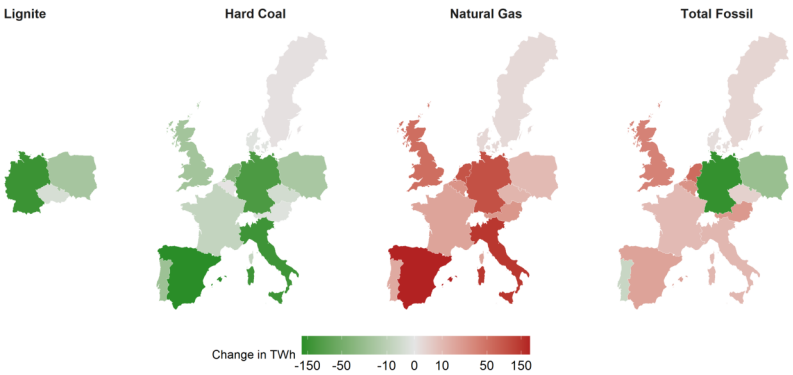
Relative to this baseline, carbon pricing induces substantial fuel switching from coal to natural gas. With actual EUA prices, this shift mainly materializes between hard coal, which drops by 19%, and natural gas, which rises by more than 16%. In contrast, lignite-fired generation decreases by less than 3%. This suggests that actual EUA prices—which fluctuated between € 5 and € 30 during most of the sample period—were sufficiently high to push efficient natural gas units ahead of hard coal units in the merit order, but not ahead of lower-cost lignite units. Only when the allowance price is raised to € 50 or more, as in the price floor scenario, does the model predict a substantial reduction in lignite generation of about 17.3% relative to the baseline of a zero EUA price. At the same time, hard coal generation further decreases to 38% below the baseline, whereas natural gas generation further grows to 38% above the baseline.

In line with the descriptive evidence shown in Figure 4b, the volume of cross-border flows also grows with the allowance price. Relative to the baseline, electricity imports increase by 15% with observed EUA prices, and by almost 24% with the € 50 price floor. We attribute this to

Figure 9: Generation changes in response to carbon pricing



(a) Actual EUA price vs. €0.



(b) EUA with €50-price floor vs. €0.

Notes: The figure shows the effect of carbon pricing on total electricity generation aggregated by country and fuel type. Panel (a) shows the effect of actual EUA prices relative to zero carbon prices. Panel (b) shows the effect of EUA prices with a price floor of €50 throughout the entire period. The underlying values are reported in Appendix Tables D.1 and D.2.

the asymmetric impact of carbon pricing on regional power-plant fleets. By disproportionately increasing the marginal costs of carbon-intensive generation, a higher allowance price widens the cost heterogeneity between bidding zones. This expands the economic scope for trade and incentivizes imports from regions with more efficient marginal units.

We break down aggregate generation changes across scenarios by country (cf. Appendix Tables D.1 and D.2) and visualize them on fuel-specific maps in Figure 9. Fuel substitution is strongest in Spain, a country with substantial gas-fired capacity installed in excess of what is usually needed. Changes in the carbon price therefore induce very large shifts in which units generate. In addition, Spain must rely a lot on domestic fuel switching because of its relatively small transmission capacity to France. The share of Spanish coal units in total fossil generation falls from almost 70% (225 TWh) at baseline to just 38.6% (126 TWh) under the ETS. At the same time, natural gas generation roughly doubles from 97 TWh (30%) to 201 TWh (61.4%).

The €50 price floor further increases the gas share by another 92.6 TWh to more than 86% of all fossil fuel generation, with coal generation dropping by another 81 TWh to 13.4%.

In Italy, substitution from coal to gas is almost one-for-one: Relative to the baseline, coal generation drops by 61 TWh with actual EUA prices while gas generation increases by 58 TWh, which is about 8% of total fossil generation, respectively. With the price floor, coal generation decreases by another 76 TWh while gas generation goes up by another 90 TWh.

Germany significantly increases electricity imports. Under actual EUA prices, combined generation from lignite and hard coal units drops by 70 TWh. It is replaced by 29 TWh of additional gas-fired generation and 43 TWh higher electricity imports. These effects are relatively small given total fossil generation of 1,150 TWh in Germany in the baseline, but they become more pronounced with the carbon price floor. Relative to the zero baseline, lignite and hard coal generation decrease by a total of 253 TWh, which is balanced by an additional 95 TWh of natural gas generation, 119 TWh of additional imports, and a reduction of electricity exports to its neighbors by 28 TWh. On aggregate, carbon pricing reallocates generation away from coal-heavy Germany (-158 TWh) and towards its gas-heavy neighbors such as the Netherlands (+57 TWh), Belgium (+25 TWh), and Austria (+22 TWh). Such increases in cross-border electricity trade allows Germany to reduce coal-fired generation and abate CO_2 emissions.

5.2 CO_2 emissions

Reallocation of electricity generation across plants in response to carbon pricing gives rise to significant reductions in CO_2 emissions, reported in Table 4. With a zero carbon price, cumulative emissions from the generating units in our sample amount to 3,105 metric megatons (Mt) of CO_2 over the six-year period. Lignite and hard coal units account for nearly 77% of CO_2 emissions despite supplying only 56% of total electricity generation. In contrast, natural gas generates 44% of electricity while contributing just 23% of emissions.

Relative to this baseline, CO_2 emissions fall by 5.7% (-176 Mt) to 2,929 Mt under actual EUA prices and by 13.9% (-431 Mt) to 2,674 Mt under the €50 price floor. The changes by fuel closely mirror the generation changes in Table 3. Total emissions fall because the generation that natural gas units take over from lignite and hard coal units is far less carbon-intensive.

Figure 10 shows how CO_2 emissions change on a per-country basis when a €50 price floor relocates generation across space. Countries in yellow and green reduce their CO_2 emissions, where yellow represents below-average mission reductions, and green above-average emission reductions. Countries in purple increase CO_2 emissions, with darker purple indicating a larger increase in emissions.

Spain and Portugal achieve the largest relative reductions in CO_2 emissions, at -39.1% (-100 Mt) and -32.5% (-12.5 Mt), respectively, as Spanish gas-fired units also displace coal-fired generation in Portugal. Italy reduces its emissions by almost 20% (-75.2 Mt). Poland and Czechia reduce their emissions by only 5.7% (-33.5 Mt) and 1.6% (-1.8 Mt), respectively, because their domestic capacities to switch to gas are limited and their strongest interconnections are to Germany, another coal-heavy country.

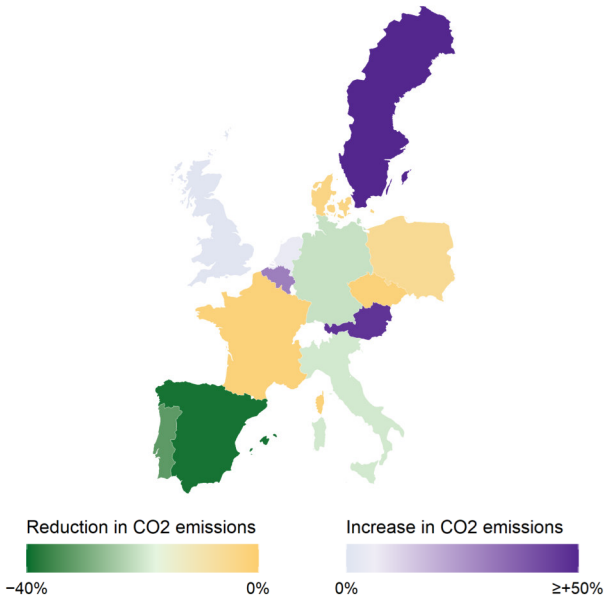
Germany achieves the largest absolute reduction, -235.8 Mt (almost -21%), the CO_2 counterpart of the decline in German lignite and hard coal generation and the rise in imports

Table 4: CO_2 Emissions [Mt] by type, aggregated 2016-2021

Type	**Lignite**	**Hard Coal**	**Natural Gas**	**Total**
Zero EUA price (€0)	1,070	1,317	718	3,105
Actual EUA price	1,038	1,059	831	2,929
% Change vs €0	-3.0%	-19.5%	+15.8%	-5.7%
EUA price floor at €50	879	810	986	2,674
% Change vs €0	-17.9%	-38.5%	+37.2%	-13.9%
Pollution tax	901	1,215	820	2,936
% Change vs €0	-15.8%	-7.7%	+14.2%	-5.4%
% Change vs Actual	-13.2%	+14.7%	-1.3%	+0.3%

Notes: Aggregate CO_2 in metric megatons (Mt), 1 Mt = 1,000,000 tons, by type of plant of all units in the analysis sample from 2016 to 2021. We compare four policy scenarios: (i) predicted emissions under €0 allowance prices, (ii) observed emissions under actual allowance prices, (iii) predicted emissions under a €50 allowance price floor, (iv) predicted emissions under a Pigouvian tax on air pollutant emissions and no carbon price.

Figure 10: Country-specific CO_2 abatement with a €50 carbon price floor, 2016-2021



Notes: The maps show how CO_2 emissions differ under the €50 carbon price floor relative to the €0 carbon price, in percent. Countries with reductions in CO_2 emissions are colored in green and yellow. Green corresponds to above-average CO_2 reductions, yellow to below-average CO_2 reductions, relative to all countries with reductions. Light lavender and purple relate to countries with CO_2 increases. Exact numbers can be found in Table D.4.

documented in the previous subsection. Part of these emissions relocate to Germany's gas-heavy neighbors: the Netherlands (+3.9%, +7.7 Mt), Belgium (+29.9%, +10.1 Mt), Austria (+47.1%, +8 Mt) and Sweden (+159.1%, +2.3 Mt) emit more CO_2 under the €50 price floor, because the additional gas-fired generation they export to Germany outweighs their own coal-to-gas switching. France and Denmark also export more to Germany, but their emissions still fall because domestic fuel switching dominates there.

Great Britain illustrates a second-round trade effect. Although British coal generation falls by 18.5 TWh (-39.3%), gas generation rises by 54.5 TWh (+10.2%), so that British CO_2 emissions increase slightly (+0.9%, +2.3 Mt). The additional output goes to France, Belgium and the Netherlands, whose net imports from Great Britain rise by a combined 31.7 TWh while their own gas-fired units export more to Germany. Our market-clearing algorithm thus captures how carbon pricing reshapes trade beyond the bilateral flows with Germany.

Table D.4 reports exact numbers for all countries and all three scenarios. Under actual EUA prices, the patterns are similar but muted, and their timing follows the merit order discussed in the previous subsection: in hard-coal-heavy Poland, almost 80% of the abatement under the €50 floor already occurs at actual EUA prices, whereas lignite-heavy Germany achieves only about a fourth of its €50-floor abatement. Likewise, the emission increases in Germany's neighbors largely materialize only under the price floor. France is an exception in that emissions rise by 4.4 Mt (+6,7%) under actual EUA prices but fall by 1.1 Mt (-1.7%) under the price floor, indicating that trade intensification precedes domestic fuel switching in France.

5.3 $PM_{2.5}$-related mortality

Table 5 reports total $PM_{2.5}$-related mortality over the six-year analysis period for different scenarios. In the baseline scenario with a carbon price of zero, we estimate that 18,790 premature deaths are due to air pollutant emissions from fossil generation in our analysis sample. Hard coal and lignite-fired units jointly account for 90% of these (47% and 43%, respectively), but contribute only 56% of electricity generation. Gas-fired units cause only 10% of all premature deaths. This is less than half of the gas share in CO_2 emissions and less than a quarter of the gas share in total electricity generation. These numbers are in line with the fuel-specific pollutant emission intensities from Figure 2a. Natural gas plants emit only about a quarter as much NO_x per MWh as lignite and hard coal plants, and effectively no SO_2.

Fuel switching induced by the actual EUA prices avoids 1,667 premature deaths, a reduction by 8.9%. Switching out of lignite and hard coal avoids 1,959 premature deaths, whereas switching into gas-fired generation only adds 292 deaths. Relative to those numbers, the €50-price floor avoids an additional 2,170 premature deaths, after netting 2,572 avoided deaths from reduced lignite and coal generation against 402 additional deaths from higher gas generation. These numbers add up to a total mortality reduction of 20.4% relative to the €0 EUA price baseline. The fact that mortality reductions are larger than the corresponding reductions in CO_2 emissions points to a potential complementarity in policies that target air pollution and climate externalities of power generation.

The sizable mortality effects are the product of complex interactions between the market-based policy instrument and heterogeneities in the regulated facilities with respect to (i) how

Table 5: Cumulative premature deaths due to co-emissions, 2016–2021

Type	Lignite	Hard Coal	Natural Gas	Total
Zero EUA price	8,071	8,837	1,882	18,790
Actual EUA price	7,742	7,207	2,174	17,123
Change vs. Zero Price [%]	−4.1	−18.4	+15.5	−8.9
EUA price floor €50	6,569	5,808	2,576	14,953
Change vs. Zero Price [%]	−18.6	−34.3	+36.9	−20.4
Air pollution tax	6,161	7,924	2,156	16,241
Change vs. Zero Price [%]	−23.7	−10.3	+14.6	−13.6
Change vs. Actual [%]	−20.4	+9.9	−0.8	−5.2

Notes: Predicted premature deaths due to co-emissions of NO_x, SO_2, and NH_3 in different policy scenarios. Numbers are calculated cumulatively for the years 2016–2021. Percentages are calculated relative to the counterfactual Zero EUA price scenario.

generating units respond to CO_2 abatement incentives, (ii) their co-pollutant intensity, (iii) their geographic location, which is decisive for atmospheric diffusion of pollutants once release and chemical reactions with the pre-existing pollution inventories and (iv) the downwind exposure of people to the resulting secondary air pollutants. To shed light on the relative importance of those channels, we decompose changes in $PM_{2.5}$-related mortality across scenarios into CO_2 abatement, fuel switching and sorting effects.

Let $E_f = \sum_{i \in f} e_i$ denote total emissions by all generating units of fuel type f, $\theta_f = \frac{E_f}{E}$ the share of fuel f in total CO_2 emissions E, and $s_i = \frac{e_i}{E_f}$ the emissions share of unit i in total emissions of its fuel type f. Rewriting aggregate mortality (eq. 1) with these definitions yields

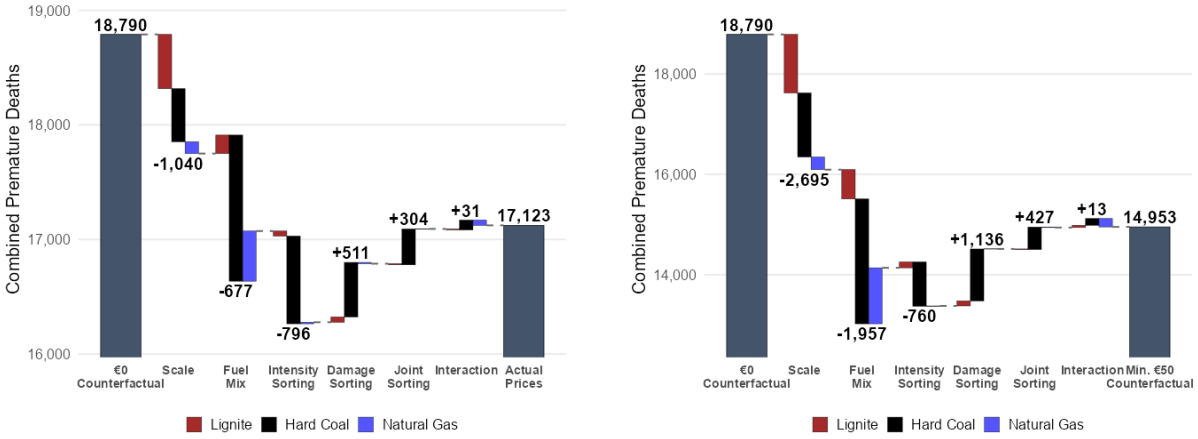
$$D = E \sum_f \theta_f \sum_{i \in f} s_i \mathbf{z}_i' \lambda_i, \tag{20}$$

which we totally differentiate to obtain

$$dD = dE(\sum_f \theta_f \sum_{i \in f} s_i \mathbf{z}_i' \boldsymbol{\lambda}_i) + E \sum_f (d\theta_f) \sum_{i \in f} s_i \mathbf{z}_i' \boldsymbol{\lambda}_i) + E \sum_f \theta_f \sum_{i \in f} (ds_i) \mathbf{z}_i' \boldsymbol{\lambda}_i. \tag{21}$$

The first two terms decompose overall CO_2 abatement into a scale effect and a fuel mix effect, respectively. They closely correspond to the scale and composition effects used in prior decomposition analyses of pollution emissions (e.g., Levinson, 2009, 2015; Shapiro and Walker, 2018). The third term reflects the health consequences of reallocating production across heterogeneous plants, holding fixed total CO_2 emissions and the fuel mix. Decomposing

Figure 11: Decomposition of changes in $PM_{2.5}$-related mortality



(a) €0 to actual EUA prices

(b) €0 to €50-floor on actual EUA prices

Notes: The graphs compare combined premature deaths due to NO_x, SO_2, and NH_3 emissions between the zero carbon price counterfactual and actual EUA prices (left) as well as between the zero carbon price counterfactual and the €50-carbon price floor counterfactual (right). The graphs decompose the total change in premature deaths into scale, fuel mix, and sorting effects, alongside an interaction term. Every partial effect is subsequently further decomposed by fuel type.

this term as follows[24]

$$E\sum_f \theta_f \sum_{i\in f}(ds_i)\mathbf{z}_i'\lambda_i) =$$

$$E\sum_f \theta_f N_f \left[\underbrace{\bar{\boldsymbol{\lambda}}'\,\mathrm{Cov}(ds_f,\mathbf{z})}_{(i)} + \underbrace{\bar{\mathbf{z}}'\,\mathrm{Cov}(ds_f,\boldsymbol{\lambda})}_{(ii)} + \underbrace{\mathbb{E}\left[ds_{i,f}\ \left(\boldsymbol{\lambda}_i-\bar{\boldsymbol{\lambda}}\right)'\left(\mathbf{z}_i-\bar{\mathbf{z}}\right)\right]}_{(iii)}\right] \quad (22)$$

allows us to distinguish between three sorting channels: (i) pollution-intensity sorting captures whether CO_2 emissions shift toward units with higher or lower co-emissions per ton of CO_2, (ii) damage sorting captures reallocation toward units where co-emissions generate higher or lower marginal damages, and (iii) joint sorting captures whether output shifts systematically toward units characterized by particular combinations of emissions intensity and marginal damages (e.g., plants that are simultaneously relatively dirty and located in high-damage areas.) Equations (21) and (22) jointly characterize the change between baseline mortality D and counterfactual mortality $\widehat{D}$, up to a residual $\widehat{\epsilon}_D = \widehat{D} - (D + dD)$, referred to as the interaction effect.

Figure 11 visualizes the decomposition of mortality changes into abatement, fuel mix, sorting effects, and interaction effects. Relative to a zero carbon price, premature deaths due to air pollutant emissions decrease by 1,667 under actual EUA prices and by 3,837 under the price floor. Across scenarios, scale, fuel mix, and pollution-intensity sorting all suppress premature deaths, whereas damage sorting and joint sorting offsets parts of these reductions.[25]

[24] We use the fact that $\sum_{i\in f} ds_i = 0\ \forall f$, and hence $\sum_{i\in f}(ds_i)\mathbf{z}_i'\lambda_i) = N_f\mathrm{Cov}(ds_f,\mathbf{z}'\lambda)$.

[25] An overview of which country × fuel pairs contribute how much to every sorting effect can be found in Table D.6.

The scale effect captures that carbon pricing lowers total CO_2 emissions. Holding the fuel mix and the allocation of emissions across plants fixed, eeach fuel contributes to the scale effect in proportion to its baseline premature deaths, which is why lignite and hard coal account for most of it. The fuel mix effect holds total CO_2 emissions fixed and captures that carbon pricing shifts CO_2 emissions between the fuels, largely away from lignite and hard coal towards natural gas. The gas component of the fuel mix effect is positive because the rising gas share adds premature deaths, but this increase is much smaller than the reduction from the falling shares of lignite and hard coal, so the fuel mix effect as a whole lowers premature deaths.

Intensity sorting looks at how CO_2 emissions move between plants of the same fuel and measures how deaths would change if all plants had the fuel's average damage per tonne of pollutant emissions. The beneficial effect of intensity sorting we find reflects a reallocation of generation from less to more efficient units of the same fuel, potentially reinforced by a likely positive correlation between net electric efficiency and pollution control technology across generator vintages. This effect is almost entirely driven by hard coal plants. The dirtiest coal plants are located in Spain, and they cut back emissions more than those in the other countries. At the same time, Netherlands, Germany and Poland gain hard-coal fired generation share with relatively clean plants. The only outlier is Italy, which features relatively clean coal plants but loses generation share. Still, across all countries in our sample, generation moves to cleaner coal plants on average. Intensity sorting among gas plants is almost zero because plants hardly differ in their air pollutant emissions per tonne of CO_2. For lignite, we even measure a slightly positive effect of intensity sorting under the €50 price floor, because lignite generation (in relative terms) moves from Germany to Poland and Czechia, where lignite plants tend to be dirtier.

Damage sorting is the main countervailing effect. It captures how premature deaths change when generation shifts between plants of the same fuel that differ only in their location, i.e., in the damages caused by each ton of pollutant they emit. For hard coal, this effect is large and positive because CO_2 emissions shift away from Spanish plants, located in places where emissions cause relatively little harm, towards plants in Poland, where damages per ton of pollutant are considerably higher. Lignite contributes much less, largely because lignite plants sit in locations with similar damages. For natural gas, damage sorting is close to zero. Damage sorting only captures changes in plants' shares within their fuel. Gas-fired generation increases across the board, and this increase is captured by the fuel mix effect; the remaining shifts within the gas fleet are not tilted towards either high- or low-damage locations.

Lastly, we estimate a positive joint sorting effect. This implies that production shifts towards plants that are dirtier than average and located in high-damage areas, or away from plants that are dirtier than average but located in low-damage areas. To a large degree, this is driven by the shrinking share of Spanish units in total hard-coal fired generation. Intensity sorting values relatively high emissions of those units at average damage whereas they are actually located where emissions cause relatively little harm. Joint sorting corrects for this overstatement, yielding a positive value.

While scale, fuel mix and sorting effects have the same sign for observed EUA prices (Figure 11a) and a price floor (Figure 11b), the key to interpreting the significant quantitative differences

across scenarios lies in the changing sources of abatement. As was discussed above, abatement from lignite-fired units becomes significant only with the higher EUA price, which drives up the contributions of the scale and fuel mix effects and leaves less scope for mortality reductions via intensity sorting within fuel types. The flip side is that premature deaths due to damage sorting double compared to actual EUA prices. As the carbon price rises, hard coal generation shifts further away from Spanish plants, located where emissions cause relatively little harm, towards Polish plants, where damages per ton of pollutant are considerably higher.

6 Cost-benefit analysis

As shown by our results, carbon pricing in the European electricity sector has not only reduced CO_2 emissions but also yielded sizable health co-benefits by reducing $PM_{2.5}$-related mortality. Abatement is realized by switching—both within and across bidding zones—from lignite and coal-fired generation towards natural gas-fired generation, which is cleaner yet more expensive. To assess the overall efficiency effects, we compute the net monetized health co-benefits and abatement costs, summarized in Table 6. We compare the resulting net co-benefits to a counterfactual where air pollution externalities are regulated with a Pigouvian tax. In what follows, we explain the steps involved in this calculation and discuss the results.

6.1 Abatement costs

To estimate total generation cost in a given scenario s, we compute unit-level generation cost as the product of marginal cost and output, and aggregate this across all fossil-fuel units and all hours-of-sample:

$$TC_Y(s) = \sum_{t=1}^{T} \sum_{i=1}^{N} \hat{y}_i(t,s)\kappa_i \cdot MC_{i,t}(\cdot|P_t^c = 0). \tag{23}$$

Although plant operators pay the EUA price, we do not include this as an economic cost in eq. (23) because permit revenues are rebated to firms (new entrants) and national governments. Abatement costs arise because carbon pricing distorts the otherwise cost-minimizing order of power units in the supply curve. In line with the literature, we assume that fixed costs and capital cost are already sunk when plant operators bid in the day-ahead market. Bids are based on avoidable costs only (variable operating costs, fuel, and carbon), allowing plants to recoup at least part of their capital costs as long as the market price exceeds its short-run marginal costs. The same assumption applies to our counterfactual scenarios as the fleet of power plants is held constant.

Under the assumption that fixed and capital costs do not vary across scenarios, total abatement cost of a positive carbon price scenario s, $\Delta TC_Y(s)$ are given by the difference in costs of generating output under positive carbon prices s and a zero carbon price. We compute aggregate emissions changes $\Delta E(s)$ analogously and compute average costs per ton of CO_2 avoided as

$$AAC(s) = \frac{\Delta TC_Y(s)}{\Delta E(s)}. \tag{24}$$

We calculate that the EU ETS increased total generation costs by €1.14 billion over the six-

Table 6: Abatement costs and health co-benefits of carbon pricing

	(1) CO_2	(2) Total Costs	(3) Health Benefits	(4) Net Benefits	(5) Deaths
A. Aggregate Change	[Mt]	[m€]	[m€]	[m€]	
Actual EUA Price	-176	1,141	6,315	5,174	-1,666
EUA price floor € 50	-430	4,127	14,270	10,143	-3,836
Δ Increment	-254	2,986	7,955	4,969	-2,170
Pollution tax	-169	1,093	9,504	8,411	-2,549
B. Per tonne CO_2 abated		[€]	[€]	[€]	
Actual EUA Price		6.48	35.88	29.40	
EUA price floor € 50		9.60	33.19	23.59	
Δ Increment		11.76	31.32	19.56	
Pollution tax		6.47	56.24	49.77	

Notes: The table reports the change in aggregate outcomes for the years 2016-2021 when going from zero carbon prices to actual EUA prices (row 1) plus an EUA price floor (row 2), the increment between the two (row 3), and when adopting a Pigouvian tax on air pollution. Negative CO_2 emissions changes (in millions of tonnes) are interpreted as abatement, and total generation cost changes correspond to abatement costs (in millions of euros). Mortality reductions are expressed as negative changes to premature deaths due to PM2.5 exposure. To compute health benefits of mortality reductions, we multiply avoided deaths with a conservative value of statistical life (VSL) estimate of €3.6 million in 2016 euros.

year observation period. With an associated abatement of 176 $MtCO_2$, the average abatement cost amounted to €6.48. Had EUA prices been subject to a price floor of €50, total abatement cost would have been €4.13 billion for 430 $MtCO_2$ abated, averaging at €9.60 per ton. Hence abatement becomes more expensive on a per-tonne basis as the carbon price rises. Consistent with economic theory, units that can contribute to abatement at lowest cost will do so first in a market-based mechanism such as the EU ETS where allowance trading is permitted. Raising EUA prices to €50 and beyond justifies more costly emissions reduction efforts on the part of firms.

6.2 Health co-benefits

We estimate health co-benefits of carbon pricing at €14.3 billion with a price floor and €6.3 billion at observed EUA prices. Importantly, those quantities are gross co-benefits and hence do not necessarily correspond to a welfare improvement. To see this, suppose that all external effects of air pollution from power plants were already internalized by a set of Pigouvian regulations that target individual plants and pollutant species. This is highly unlikely to hold in practice, but it helps to build intuition for the insight that additional mortality reductions brought about by carbon pricing need not enhance efficiency. In our hypothetical example, carbon pricing would reduce $PM_{2.5}$ concentrations to levels *below* the social optimum.

We thus turn to evaluating whether carbon pricing yields positive *net* co-benefits, computed as the difference between gross co-benefits and total abatement costs. We find positive net co-

benefits of €5.2 billion for actual EUA prices and €10.1 billion with a €50 price floor.[26] Per tonne of CO_2 abated, actual EUA prices yields an average gross co-benefit of €35.88. Incremental abatement induced by the €50 price floor yields smaller gross co-benefits of €31.32/tCO_2 and also comes at a higher cost of €11.76/tCO_2. This leads to average net co-benefits falling from €29.40/tCO_2 under actual carbon prices to €19.56 for incremental CO_2 abatement efforts induced by the price floor.

Taken together, these results suggest that carbon pricing in the European power sector yielded substantial health co-benefits between 2016 and 2021. Short-run responses to carbon pricing, such as fuel switching within and across bidding zones, improved total welfare. Our results also indicate that marginal net co-benefits are diminishing with higher carbon prices because increases in marginal abatement costs in response to higher prices are not matched by similar increases in marginal health co-benefits.

6.3 Pigouvian air pollution taxes as an efficiency benchmark

The net co-benefits we have estimated provide a sufficient statistic for positive welfare effects of co-pollution changes. Alternatively, we can use our model also to assess efficiency relative to optimal regulation of air pollution. To this end, we compute counterfactual outcomes assuming that each generating unit faces a Pigouvian tax in the amount of the health damages created by its emissions of NO_x, SO_2 and NH_3. The marginal cost function becomes

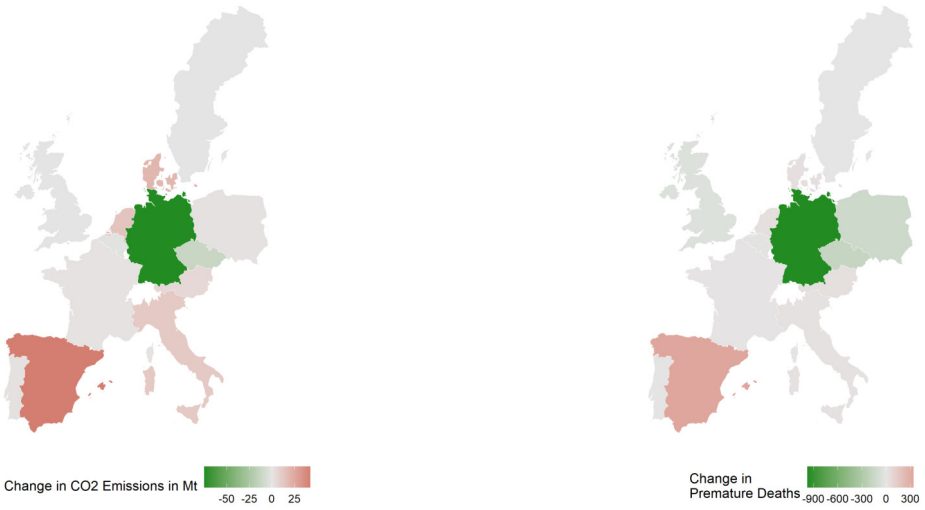
$$MC_{i,t}^{\text{Pigou}} = VOC_i + \frac{P_{i,t}^f + \tau_{i,t}}{\eta_i} + \underbrace{\frac{\xi_i}{\eta_i}\mathbf{z}_i'\boldsymbol{\lambda}_i(t)\cdot VSL(t)}_{T_i(t)}, \tag{25}$$

where the carbon price is set to zero, and $T_i(t)$ is a unit-specific Pigouvian tax per MWh of electricity. It equals the unit's CO_2 emissions per MWh (ξ_i/η_i) times its co-pollution intensities $\mathbf{z}_i$, the month-specific damage sensitivities $\boldsymbol{\lambda}_i(t)$ at its location, and the VSL. Such a tax gives firms an incentive to internalize the external health costs of those pollutants when taking their production decisions. Note that this Pigouvian tax is specific to the power plant and reflects the variability in the pollution intensities across plants and in health impacts across locations.

We evaluate the counterfactual scenario in which only the Pigouvian air pollutant emissions tax is in place and no carbon price. As shown in Table 3 above, carbon pricing causes fuel switching away from hard coal towards natural gas, with only moderate effects on lignite generation at actual EUA prices. In contrast, the air pollution tax reduces lignite-fired generation more strongly (by 15.5% vs. 3%) than hard coal generation (by 7.9% vs. 19.3%), relative to the baseline. The impact on natural gas differs only a little (+16.2% under actual EUA prices vs. +14% under the air pollutant tax). The extent of between-zone relocation is smaller under the pollution tax (+11.3%) than with carbon pricing (+15.5%). On aggregate, these differences do not affect CO_2 abatement in major ways: The pollution tax leads to

[26] Net benefits might be higher because the VSL estimate we use is at the lower end; yet they might also be lower due to omitted elements of marginal costs that are related to air pollution control or due to omitted components of variable operating costs that do not cancel out when we take the difference between generation costs with and without carbon prices. To err on the side of caution, we maintain the low VSL estimate which allows us to conclude by a comfortable margin that the ETS yielded positive net co-benefits.

Figure 12: CO_2 emissions and premature deaths, 2016-2021, replacing the EU ETS with an air pollution tax



(a) CO_2 emissions

(b) Premature deaths

Notes: The maps show the effects if we switched from the EU ETS to the Pigouvian air pollution tax. Map 12a on the left shows how CO_2 emissions in every country would change if this switch occurred. Map 12b on the right shows how premature deaths would change if this switch occurred. Importantly, the map shows premature deaths due to emissions from that country, not premature deaths in that country.

abatement by 169 Mt (-5.4%) relative to the baseline, only 7.4 Mt less than carbon pricing.

Fundamentally, the pollution tax penalizes generation by plants that (i) emit high quantities of air pollutants and (ii) are located where co-emissions cause large damages. In contrast, carbon pricing only imperfectly proxies the first factor, and ignores the second. Consequently, the air pollution tax reduces premature deaths substantially more than carbon pricing, namely by about 2,550 (-13.6%) and 1,670 (-8.9%), respectively. The principal reason for this is a better targeting of highly polluting generation from lignite fired power plants. While the degree of targeting would likely prove prohibitive in practice, the Pigou scenarios serves as a useful benchmark to assess how closely the EUA price proxies for Pigouvian taxes on air pollutants.

Given that both policies achieve very similar levels carbon abatement (169 to 176 Mt) at virtually the same average cost (€6.5/tCO_2), an approximate measure of efficiency can be gleaned directly from differences in mortality, where the EU ETS achieved almost two thirds (65.4%) of the Pareto-efficient mortality reduction. Adjusting for the minor difference in overall CO_2 abatement by expressing health benefits per tonne of CO_2 abated, the EU ETS achieves 63.8% of the optimal co-benefit of €56.2. Additionally correcting for differences in abatement costs leaves the optimal net co-benefit at €49.8, of which abatement induced by the EU ETS achieved 59.1%.

Alternatively, we assess the cost effectiveness of both instruments at reducing mortality impacts of air pollution. We find that the cost of one premature death avoided is about 37% lower under pollution taxes (€429k per avoided death) than under carbon pricing (€685k per avoided death).

Comparing aggregate policy effects masks substantial geographic heterogeneity. Figure 12

shows country-level changes in CO_2 emissions and premature deaths when carbon pricing at actual EUA prices is replaced with the Pigouvian air pollution tax. The different responses to such a policy change can be classified according to four broad patterns. In a first group of countries, the pollution tax leads to less abatement and higher health damages. A case in point is Spain where carbon pricing induces coal-to-gas switching at a large scale. At the same time, air pollutant emissions from Spanish gas-fired power plants cause the lowest downwind health impacts, relative to the other European countries in our sample (cf. Figure 2b). This is why pollution abatement with EUA prices is inefficiently high from the Pigouvian point-of-view. The pollution tax thus lowers the financial burden on Spanish fossil fuel plants compared to carbon pricing, with the result that Spanish CO_2 emissions go up by 42 Mt, and the concurrent increase in co-pollution causes 339 additional premature deaths.

The opposite pattern can be observed in the second group of countries where targeted air pollution taxes reduce emissions and premature deaths by more than carbon pricing. For example, in Germany and Czechia, the Pigouvian tax puts a relatively high burden on fossil power plants because of their high downwind health impacts. Consequently, German CO_2 emissions are 75 Mt lower, and 977 fewer premature deaths occur due to German emissions. The excess abatement is driven almost exclusively by lignite-fired plants located either in the very densely-populated Rhein-Ruhr area or in the far east where predominant western winds blow pollution into densely populated parts of Poland. In Czechia, where carbon pricing barely changes the dominance of lignite plants in electricity generation, air pollution taxes are high enough to induce significant output reductions by particularly harmful power plants. Accordingly, Czech CO_2 emissions are 12.6 Mt lower under air pollutant pricing than under the EU ETS, with 178 fewer premature deaths from Czech emissions.

A third group of countries have substantially higher CO_2 emissions under the Pigouvian tax than under the EU ETS, such as Italy (+11.5 Mt), the Netherlands (+13.4 Mt), and Denmark (+19.1 Mt), whereas mortality due to co-emissions hardly increases. In Italy, coal-fired generation is higher and gas-fired generation lower under the air pollution tax. The reason lies in very high pollution taxes on gas-fired plants clustered in Northern Italy, where SO_2 and especially NO_x emissions have higher mortality impacts than anywhere else in Europe. Generation shifts to coal plants located along the Italian coast where co-emissions cause less harm. Accordingly, reverting the coal-to-gas shift associated with carbon pricing leads to only a moderate increase by 21 premature deaths in Italy. A similar mechanism is at work in the Netherlands and Denmark whose power plants have below-average downwind health impacts. In addition, both countries increase electricity exports to Germany over and above what is required under carbon pricing (by 14.5 TWh for the Netherlands and 18.1 TWh for Denmark) when the Pigouvian tax replaces the ETS. Again, the net effect is a slight increase in premature deaths caused by air pollution from power plants in Denmark (+28) and the Netherlands (+31).

A fourth case is illustrated by Poland where lignite plants emit substantially higher amounts of air pollutants than hard coal plants. Taxing air pollution instead of CO_2 induces substitution of 35 TWh from lignite to hard coal, which prevents 133 premature deaths but keeps CO_2 emissions practically unchanged (+1.8 Mt / +0.3% under the EU ETS).

6.4 Discussion

Our results, in particular those obtained by comparing outcomes with actual EUA prices to those without carbon prices, are informative about the causal effects of carbon trading under the EU ETS on the power sector. In what follows, we state and discuss our main conclusions in this respect.

First, the EU ETS has been an effective tool to mitigate climate-change related externalities of power generation in Europe. We attribute a 5.7%-reduction in CO_2 emissions to the carbon price, driven by fuel switching among fossil power generating units. We do not attempt to quantify abatement due to alternative channels affected by carbon pricing, such as efficiency and capacity upgrades to the fossil power plant fleet, increased supply from non-fossil power plants, or reduced electricity demand due to carbon pricing. All of these channels work in the direction of larger CO_2 abatement. Thus, we consider our abatement estimate a lower bound.

Second, the EU ETS has given rise to significant economic co-benefits by reducing air pollution and $PM_{2.5}$-related premature mortality. Our monetary estimate of the net-benefits is about 4.5 times higher than average abatement costs. Additional co-benefits may arise due to air quality benefits other than mortality changes arising from reduced $PM_{2.5}$ concentrations, such as morbidity reductions from $PM_{2.5}$ and its precursor emissions. While we deem the mortality impacts estimates in this paper to be of first-order importance, we consider those estimates to be conservative in the sense that they do not account for air quality co-benefits in their entirety. Similarly, our monetary co-benefits estimates would scale up when using a less conservative VSL estimate than we do. When taken at face value, our estimate of EUR 29/ tCO_2 suggests that it pays for the EU to abate CO_2 in the power sector, irrespective of the global climate benefits that accrue on top.

Third, the EU ETS has internalized almost 60% of air pollution externalities from power generation when compared to a hypothetical Pigouvian tax instrument on air pollution. Given the bluntness of the carbon price vs. the sophistication of taxes differentiated at the pollutant-by-location-by-calendar-month level, carbon pricing can be regarded as a striking success in terms of substituting for the lack of efficient air pollution taxes in Europe. A caveat is that our comparison assumes that both carbon pricing and pollution taxes act exclusively via fuel switching. Taxes may yield even higher environmental gains than what we calculate if they induce plants to improve pollution control on site rather than just curbing utilization.[27]

Finally, regulatory overlap of the EU ETS with air pollution policies set by local regulators at plant locations might interfere with our benefit calculation. If such local policies were chosen

[27] This observation draws attention to a fundamental difference between CO_2 and air pollutant emissions. The former increase linearly with fuel input. Doubling the amount of fuel burned will double CO_2 emissions as plant operators have very little influence over this relationship. Carbon Capture and Storage (CCS) technology would alter the relationship between fuel burned and CO_2 emitted. However, CCS was not applied at any power plant in Europe during the late 2010s and early 2020s. Europe's first, rather experimental, post-combustion CCS plant opened in Italy in late 2024. In contrast, air pollutant emissions depend not only on the fuel burned, but also significantly on plant-specific emissions control technology, such as filters or exhaust fume recirculation. Variable operating costs of such technologies, primarily reagents and auxiliary power, typically remain below €1 MWh. This is small relative to variable compliance costs with the EU ETS. Furthermore, as pollution control technologies are 'end-of-pipe', they do not influence the carbon-driven merit order shift. This is because the costs of emission control technology are sunk and rather than a variable decision lever. This explains our finding that gross co-pollution benefits do not significantly vary over different magnitudes of carbon prices, in panel B of Table 6.

optimally, power plants would already operate in ways that internalize air pollution externalities. Accordingly, the ETS might be "over-correcting" at that specific location. However, we believe that this is unlikely to drive our results. First, local policymakers are heavily constrained by harmonized European regulations that already target these power plants. Second, any effects on the fleet composition of such regulations are held fixed across scenarios. Third, our estimates reveal that marginal health benefits of additional abatement would significantly exceed the marginal costs, implying that existing regulations remain sub-optimal. In this case, the carbon price acts as a corrective co-benefit, moving the system toward a social optimum that the existing mix of local and European regulation has failed to achieve.

7 Conclusion

In this paper, we have developed an empirical framework for quantifying emissions abatement and ancillary health co-benefits due to carbon pricing in the electricity sector. Drawing on machine learning tools to model electricity dispatch, rich data on unit-level co-pollutant emission intensities, and a chemical transport model linking emissions to population exposure, we estimate that the EU ETS caused changes in air pollutant emissions that significantly reduced $PM_{2.5}$-related mortality over the period from 2016 to 2021. Driven by policy-induced fuel switching from coal to natural gas fired generation, the monetized co-benefits—net of moderate abatement costs—are unambiguously positive.

A novel feature of our empirical framework is the modeling of cross-border electricity flows to predict how generation relocates across zones under different carbon prices. This is key to understanding decentralized abatement decisions in interconnected electricity grids. Germany offers the most striking illustration because it is a lignite- and coal-heavy country surrounded by several gas-heavy neighbors. As the carbon price rises, Germany reduces generation at hard coal and—at higher carbon prices—also lignite fired power units. To compensate for these reductions, the country not only ramps up gas-fired generation but imports increasing amounts of gas-fired electricity from neighboring countries like the Netherlands which, in turn, import more electricity from their neighbors such as Great Britain. Our granular analysis at the unit level, paired with a chemical transport model allows us to trace how the EU ETS has redistributed air pollution across Europe.

Our findings carry several implications for climate and air quality policy. First, they confirm that the EU ETS, in addition to reducing carbon emissions, has delivered local health dividends by internalizing air pollution externalities that pre-existing regulation had failed to fully address. The magnitude of these co-benefits suggests that evaluations of this policy instrument should also place significant weight on ancillary health benefits, particularly in coal-heavy electricity systems. Second, our results caution against assuming that these co-benefits will automatically scale with carbon prices in the long run. As marginal abatement costs rise and the remaining scope for low-cost fuel switching narrows, the realization of further health gains will increasingly depend on broader structural changes to the energy system that lie beyond the scope of our analysis of the short-run dispatch margin.

Several caveats deserve emphasis. Our analysis focuses on the intensive margin of electricity

dispatch and abstracts from longer-run investment responses, which may have amplified the effects we document. Moreover, while we find that the health co-benefits of the EU ETS substantially exceed its abatement costs, this result rests on the assumption that pre-existing air quality regulations are insufficient to close the gap between private and social costs, which we can verify only indirectly. A promising avenue for expanding our work will be to explore the distributional dimension of our findings in more detail, examining how the health benefits from carbon pricing may differ across communities, and how these differences relate to socioeconomic variables.

References

Abrell, J., Kosch, M., and Rausch, S. (2022). How effective is carbon pricing?—a machine learning approach to policy evaluation. *Journal of Environmental Economics and Management*, 112:102589.

Aunan, K., Berntsen, T., O'Connor, Persson, T. H., Vennemo, H., and Zhai, F. (2007). Benefits and costs to China of a climate policy. *Environment and Development Economics*, 12:471–497.

Basaglia, P., Grunau, J., and Drupp, M. A. (2024). The european union emissions trading system might yield large co-benefits from pollution reduction. *Proceedings of the National Academy of Sciences*, 121(28):e2319908121.

Bayer, P. and Aklin, M. (2020). The European Union Emissions Trading System reduced CO_2 emissions despite low prices. *Proceedings of the National Academy of Sciences U.S.A.*, 117:8804–8812.

Borenstein, S. and Bushnell, J. B. (2022). Do two electricity pricing wrongs make a right? cost recovery, externalities, and efficiency. *American Economic Journal: Economic Policy*, 14(4):80–110.

Burtraw, D., Krupnick, A., Palmer, K., Paul, A., Toman, M., and Bloyd, C. (2003). Ancillary benefits of reduced air pollution in the US from moderate greenhouse gas mitigation policies in the electricity sector. *Journal of Environmental Economics and Management*, 45(3):650–673.

Butters, R. A., Dorsey, J., and Gowrisankaran, G. (2025). Soaking Up the Sun: Battery Investment, Renewable Energy, and Market Equilibrium. *Econometrica*, 93(3):891–927.

Carlson, C., Burtraw, D., Cropper, M., and Palmer, K. (2000). Sulfur dioxide control by electric utilities: What are the gains from trade? *Journal of Political Economy*, 108(6):1292–1326.

Center For International Earth Science Information Network-CIESIN (2018). Gridded population of the world, version 4 (gpwv4): Population count, revision 11 (version 4.11) [data set]. Technical report, Columbia University, Palisades, NY. NASA Socioeconomic Data and Applications Center (SEDAC).

Colmer, J., Martin, R., Muûls, M., and Wagner, U. (2024). Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence From the European Union Emissions Trading System. *Review of Economic Studies*, 92(3):1625–1660.

Cushing, L., Blaustein-Rejto, D., Wander, M., Pastor, M., Sadd, J., Zhu, A., and Morello-Frosch, R. (2018). Carbon trading, co-pollutants, and environmental equity: Evidence from california's cap-and-trade program (2011–2015). *PLOS Medicine*, 15(7):1–20.

Davis, D., Krupnick, A., and McGlynn, G. (2000). Ancillary benefits and costs of greenhouse gas mitigation - an overview. In OECD, editor, *Ancillary Benefits and Costs of Greenhouse Gas Mitigation*, pages 9–49, Paris.

Deschenes, O. and Weber, P. (2025). Benefits, costs, and distributional outcomes in an environmental market. UC Berkeley: Energy Institute Working Paper No. 355. Retrieved from https://escholarship.org/uc/item/2162289d.

Driscoll, C. T., Buonocore, J. J., Levy, J. I., Lambert, K. F., Burtraw, D., Reid, S. B., Fakhraei, H., and Schwartz, J. (2015). Us power plant carbon standards and clean air and health co-benefits. *Nature Climate Change*, 5:535–540.

Egerer, J. (2016). Open source electricity model for germany (elmod-de). Working Paper No. 83, German Institute for Economic Research.

Ekin, P. (1996). The secondary benefits of CO2 abatement: How much emission reduction do they justify? *Ecological Economics*, 16(1):13–24.

Ekins, P. (1996). How large a carbon tax is justified by the secondary benefits of CO2 abatement? *Resource and Energy Economics*, 18(2):161–187.

Ellerman, D., Marcantonini, C., and Zaklan, A. (2016). The european union emissions trading system: Ten years and counting. *Review of Environmental Economics and Policy*, 10(1):89–107.

European Commission, Directorate-General for Mobility and Transport (2019). Handbook on the external costs of transport: Version 2019. Technical report, Publications Office of the European Union, Luxembourg. Prepared by CE Delft, INFRAS, TRT and Ricardo.

European Environment Agency (2025). Industrial Reporting under the Industrial Emissions Directive 2010/75/EU and European Pollutant Release and Transfer Register Regulation (EC) No 166/2006. Dataset, version 14.0. March 2025 release. Reporting period: 2007–2023.

European Network of Transmission System Operators for Electricity (2024). Entso-e transparency platform: Actual generation per production type [16.1.b&c]. Online electricity market database. Brussels: European Network of Transmission System Operators for Electricity. Accessed: 2024-07-18.

Fowlie, M., Holland, S. P., and Mansur, E. (2012). What do emissions markets deliver and to whom? Evidence from southern california's nox trading program. *American Economic Review*, 102(2):965–993.

Grainger, C. and Ruangmas, T. (2018). Who Wins from Emissions Trading? Evidence from California. *Environmental & Resource Economics*, 71(3):703–727.

Grossi, L., Heim, S., and Waterson, M. (2017). The impact of the German response to the Fukushima earthquake. *Energy Economics*, 66:450–465.

Gu, Y., Henze, D. K., Nawaz, M. O., Cao, H., and Wagner, U. J. (2023). Sources of pm2.5-associated health risks in europe and corresponding emission-induced changes during 2005-2015. *GeoHealth*, 7(3):e2022GH000767.

Henze, D. K., Hakami, A., and Seinfeld, J. H. (2007). Development of the adjoint of geos-chem. *Atmospheric Chemistry and Physics*, 7(9):2413–2433.

Hernandez-Cortes, D. and Meng, K. C. (2023). Do environmental markets cause environmental injustice? evidence from california's carbon market. *Journal of Public Economics*, 217:104786.

Hidalgo Gonzalez, I., Kanellopoulos, K., De Felice, M., and Bocin, A. (2019). JRC Open Power Plants Database (JRC-PPDB-OPEN). Available at: https://data.jrc.ec.europa.eu/dataset/9810feeb-f062-49cd-8e76-8d8cfd488a05. [Accessed September 26, 2026].

Hirth, L., Muehlenpfort, J., and Bulkeley, M. (2018). The entso-e transparency platform – a review of europe's most ambitious electricity data platform. *Applied Energy*, 225.

Holland, S. P., Kotchen, M. J., Mansur, E. T., and Yates, A. J. (2026). On the Feasibility, Costs, and Benefits of an Immediate Phasedown of Coal for US Electricity Generation. *Journal of the Association of Environmental and Resource Economists*, 13(1):77–114.

Holland, S. P. and Mansur, E. T. (2008). Is real-time pricing green? the environmental impacts of electricity demand variance. *The Review of Economics and Statistics*, 90(3):550–561.

Holland, S. P., Mansur, E. T., Muller, N. Z., and Yates, A. J. (2016). Are there environmental benefits from driving electric vehicles? The importance of local factors. *American Economic Review*, 106(12):3700–3729.

International Energy Agency (2026). Electricity 2026: Analysis and forecast to 2030. Technical report, International Energy Agency (IEA), Paris, France.

IPCC (2023). *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. IPCC, Geneva, Switzerland.

Jarvis, S., Deschenes, O., and Jha, A. (2022). The private and external costs of Germany's nuclear phase-out. *Journal of the European Economic Association*, 20(3):1311—-1346.

Juhrich, K. (2016). CO2 Emission Factors for Fossil Fuels. Climate Change 28/2016, German Environment Agency (Umweltbundesamt).

Levinson, A. (2009). Technology, International Trade, and Pollution from US Manufacturing. *American Economic Review*, 99(5):2177–2192.

Levinson, A. (2015). A Direct Estimate of the Technique Effect: Changes in the Pollution Intensity of US Manufacturing, 1990–2008. *Journal of the Association of Environmental and Resource Economists*, 2(1):43–56.

LSEG (2025). LSEG Datastream. Database, accessed via LSEG Workspace. Daily API2 coal prices, EUA settlement prices and natural gas hub prices, 2016–2021. Retrieved [06 05 2025].

Martin, R., Muûls, M., and Wagner, U. J. (2016). The impact of the european union emissions trading scheme on regulated firms: What is the evidence after ten years? *Review of Environmental Economics and Policy*, 10(1):129–148.

Mauzerall, D. L., Sultan, B., Kim, N., and Bradford, D. F. (2005). NOx emissions from large point sources: variability in ozone production, resulting health damages and economic costs. *Atmospheric Environment*, 39(16):2851–2866.

Meinshausen, N. (2006). Quantile regression forests. *Journal of Machine Learning Research*, 7(35):983–999.

Muller, N. Z. and Mendelsohn, R. (2009). Efficient Pollution Regulation: Getting the Prices Right. *American Economic Review*, 99(5):1714–1739.

Murray, C. and *et al.* (2020). Global burden of 87 risk factors in 204 countries and territories, 1990–2019: a systematic analysis for the global burden of disease study 2019. *The Lancet*, 396(10258):1223–1249.

Murray, C. J. L., Aravkin, A. Y., Zheng, P., Abbafati, C., Abbas, K. M., and Abbasi-Kangevari, M., e. (2020). Global burden of 87 risk factors in 204 countries and territories, 1990–2019: A systematic analysis for the global burden of disease study 2019. *The Lancet*, (10258):1223 – 1249.

Nawaz, M. O. and Henze, D. K. (2020). Premature deaths in Brazil associated with long-term exposure to PM2.5 from Amazon fires between 2016 and 2019. *GeoHealth*, 4(8):e2020GH000268.

Open Power System Data (2020). Data package conventional power plants. version 2020-10-01.

Papke, L. E. and Wooldridge, J. M. (1996). Econometric Methods for Fractional Response Variables with Application to 401(k) Plan Participation Rates. *Journal of Applied Econometrics*, 11(6):619–632.

Pittel, K. and Rübbelke, D. T. G. (2008). Climate policy and ancillary benefits: A survey and integration into the modelling of international negotiations on climate change. *Ecological Economics*, 68(1-2):210–220.

Rennert, K., Errickson, F., Prest, B., and *et al.* (2022). Comprehensive evidence implies a higher social cost of CO_2. *Nature*, 610:687–692.

Rypdal, K., Rive, N., Åström, S., Karvosenoja, N., Aunan, K., Bak, J. L., Kupiainen, K., and Kukkonen, J. (2007). Nordic air quality co-benefits from European post-2012 climate policies. *Energy Policy*, 35(12):6309–6322.

Sargent & Lundy (2019). Generating Unit Annual Capital and Life Extension Costs Analysis. Technical report, U.S. Energy Information Administration. Prepared for the U.S. Energy Information Administration, Contract No. DE-EI0003250.

Shapiro, J. S. and Walker, R. (2018). Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade. *American Economic Review*, 108(12):3814–3854.

van Donkelaar, A., Hammer, M. S., Bindle, L., Brauer, M., Brook, J. R., Garay, M. J., Hsu, N. C., Kalashnikova, O. V., Kahn, R. A., Lee, C., Levy, R. C., Lyapustin, A., Sayer, A. M., and Martin, R. V. (2021). Monthly global estimates of fine particulate matter and their uncertainty. *Environmental Science & Technology*, 55(22):15287–15300.

van Vuuren, D. P., Cofala, J., Eerens, H. E., Oostenrijk, R., Heyes, C., Klimont, Z., den Elzen, M. G. J., and Amann, M. (2006). Exploring the ancillary benefits of the Kyoto Protocol for air pollution in Europe. *Energy Policy*, 34(4):444–460.

Zaklan, A. (2023). Coase and Cap-and-Trade: Evidence on the Independence Property from the European Carbon Market. *American Economic Journal: Economic Policy*, 15(2):526–58.

Zhang, S. and Almond, D. (2026). Too little too late? Nationwide carbon trading in China. *Journal of Public Economics*, 263:105798.

Appendix - For Online Publication Only

A Data Appendix

A.1 Electricity generation

The main data source for this study is the Central Information Transparency Platform of the European Network of Transmission System Operators for Electricity (2024). The ENTSO-E Transparency Platform offers free access to essential data and insights on electricity generation, transmission, and consumption across the pan-European wholesale energy market. ENTSO-E provides hourly data on the electricity market for all countries of the European Economic Area (EEA) plus Albania, Bosnia and Herzegovina, North Macedonia, Montenegro, Serbia, Switzerland, and Ukraine. Turkey and Moldova have observer status. The United Kingdom was a member until June 15, 2021. Because Northern Ireland is part of the Single Electricity Market (SEM) on the island of Ireland, the Northern Irish system operator has remained an ENTSO-E member and continues to submit data on electricity generation to ENTSO-E.

Actual Generation per Generation Unit This data set [16.1.A] is the centerpiece of this project. It contains hourly data on electricity generation on the unit level. A unit refers to an individual unit in a fossil fuel power plant. The data set contains the unit's bidding zone, a unit identifier, the unit's production type (lignite, hard coal, natural gas, hydro power, wind, etc.), the installed capacity, and the actual generation in the given hour. The units can be further matched to the plants they belong to, allowing us to identify sets of units within one plant. EU regulation requires all operators to report a unit's hourly electricity generation to ENTSO-E if the unit has at least 100 MW of installed capacity. However, this threshold appears to be applied somewhat inconsistently. Hirth et al. (2018) point out that some units with less than 100 MW are included while other generating units are completely missing. In the case of solar parks and wind parks, some operators apply the threshold to the individual PV panels or turbines and therefore do not report any unit-level data (e.g. German operators) whereas others consider their entire park to be one unit and therefore do report (e.g. British operators). Data submission for England, Wales and Scotland largely ceases during the early months of 2021 due to Brexit. We therefore obtain unit-level generation for British units in 2021 from Elexon, the administrator of the GB Balancing and Settlement Code, through its Insights Solution API.

Actual Total Load and Actual Generation per Production Type The first data set [6.1.A] contains gross electricity demand on the bidding zone level. Gross electricity demand is called "Total Load". The second data set [16.1.B] contains total electricity generation by production type at the level of the bidding zone. In both data sets, most bidding zones provide data at the hourly level. Some bidding zones (Hungary, Belgium, the Netherlands, Germany, Austria, and Luxembourg) report total load and generation per type in 15-minute intervals. The British and Irish bidding zones report in half-hourly intervals. The information from both data sets is combined to obtain an estimate for the share of total load that is relevant for fossil

fuel units.[28] We obtain data on load and generation per production type for post-Brexit Britain from the UK National Energy System Operator.

Further ENTSO-E data sets ENTSO-E's data set "Scheduled Commercial Exchange [12.1.F]" provides hourly data on electricity trade between bidding zones. Because Italy does not report scheduled flows between its domestic bidding zones before Jan 1, 2018, we download publicly available data on scheduled flows between Italian bidding zones in 2016 and 2017 from Gestore dei Mercati Energetici (GME). Moreover, ENTSO-E provides day-ahead electricity prices in "Day-ahead Prices [12.1.D]". These are the prices that result from the merit order principle using day-ahead forecasts of total load and that generally determine dispatch in the energy market, with the exception of short-term dispatch that responds to unexpectedly high demand. These day-ahead prices vary at the hourly level as the merit curve is established separately for every hour of the day. As there is no 2021 data for Great Britain due to Brexit, we obtain data on British electricity prices in 2021 from the global energy think tank Ember. Lastly, ENTSO-E also makes available data on unavailability of units due to maintenance work or due to forced outages in Planned Unavailability of Generation Units [15.1.A] and Changes in Actual Availability of Generation Units [15.1.B].

Input prices Input price data come from LSEG Datastream, formerly known as Refinitiv Datastream. This source provides daily data on the API2 coal index. This is the standard benchmark for European steam coal used primarily in electricity generation. According to S&P global, this price "reflects the tradable, repeatable spot market value" of hard coal. LSEG Datastream also provides data on the daily settlement price of EU ETS emission allowances. This price is the official end-of-day price on the European Energy Exchange (EEX), a major trading venue for both spot and derivatives. There is also data on natural gas prices by natural gas trading hub. So far, these include the PSV hub for Italy, MIGBAS for Spain and Portugal, POLPX for Poland, THE for Germany, TTF for the Netherlands, and VTP for Austria and Czechia. Additionally, we impute natural gas prices for Denmark and (Southern) Sweden with the THE hub, and for Belgium using the TTF hub.

Generator characteristics Several other data sources expand the unit-level data from ENTSO-E. The JRC Open Power Plants Database from the Joint Research Centre of the European Commission provides unit-level information on minimum stable load, ramping capabilities, and net electric efficiency. This data is valuable to calculate marginal costs. "Minimum stable load" is the lowest fraction of its maximum capacity at which it can run continuously. Coal plants, for instance, need to run at a minimum load in order to ensure combustion stability and produce enough steam to maintain required temperature and pressure. The "ramping capability" is the speed at which units can power up or down. Fossil fuel units generally cannot go from being shut off to maximum generation in an instant but need time

[28] Assessing an early version of the ENTSO-E data in comparison to other sources, Hirth et al. (2018) find that reporting gap was relatively stable over time.

to power up. In the case of lignite units, this can take several hours. Lastly, "net electric efficiency" denotes the proportion of a fuel's inherent energy that is converted into electricity. Conventional coal- or gas-fired power stations typically exhibit net electric efficiencies of only 30–40 percent, indicating that merely 30–40 percent of the combustion energy is converted into electricity, while the remainder is dissipated as waste heat. The Open Power Plants Database also provides latitude and longitude information for units as well as information on possible decommissioning.[29]

Data on transportation cost for hard coal in Germany comes from the Open Source Electricity Model for Germany ELMOD-DE, maintained by DIW German Institute for Economic Research (Egerer, 2016). Information on a unit's technology type (steam turbine, gas turbine, or combined cycle) as well as on whether it is a combined heat-and-power (CHP) unit comes from three sources: First, Jarvis et al. (2022) provide this information for German power units. Second, Open Power System Data (2020) provides data on the technology type and CHP status of fossil fuel power units in 15 EEA countries plus the United Kingdom and Switzerland. Third, gaps in the data have been manually filled by using Global Energy Monitor's GEM.wiki.

A.2 Emissions data

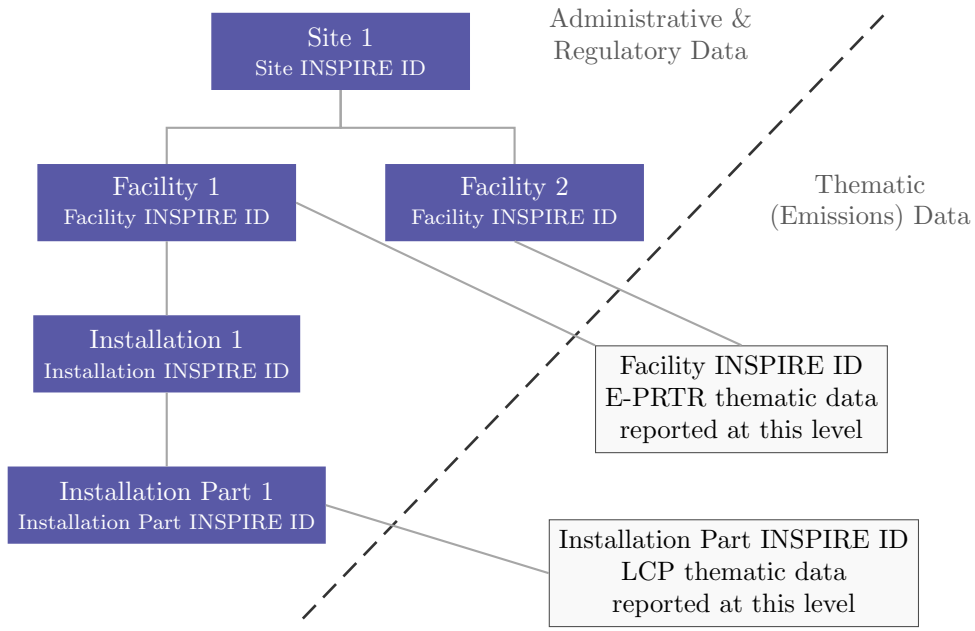
Data sources Annual reporting of pollution emissions has been mandatory for large industrial facilities under several EU regulations. The European Pollutant Release and Transfer Register Regulation (EC) No. 166/2006 has created the European Pollutant Release and Transfer Register (E-PRTR). The Large Combustion Plant (LCP) Directive 2001/80/EC, followed by the Industrial Emissions Directive (IED) 2010/75/EU, impose additional reporting requirements for large combustion plants. All data collected for the E-PRTR, LCP, and IED are publicly available as a combined database covering the years 2007-2023 from the European Environment Agency (2025). This database is organized at the level of sites, facilities, installations, and installation parts. A single site may contain multiple facilities, and each facility may in turn comprise multiple installations and installation parts. Figure A.1 illustrates this hierarchical structure.

Reporting in the E-PRTR is at the facility level and includes information on CO_2, NO_x, SO_2, and NH_3. LCP data are available at the installation-part level and include annual reported emissions of SO_2, NO_x, and dust. In addition, the database contains information on energy inputs and thermal capacity (LCP) as well as on the main fuel type (E-PRTR).

Entity linking We first match ENTSO-E power plants to E-PRTR facilities based on the name and geographic location. Within each matched plant-facility pair, we match individual ENTSO-E generating units to LCP installation parts by comparing unit names in ENTSO-E with feature names in LCP, along with generating capacity and thermal input information. Since the latter is not directly reported in ENTSO-E, we impute thermal input by dividing the

[29] Occasionally, units stay in the ENTSO-E register and continue to report zero generation even if they were already decommissioned.

Figure A.1: Hierarchy of administrative and thematic components in the data model



unit's electricity output observed in ENTSO-E by its net electrical efficiency, defined as the ratio of electricity output to energy input, available from Hidalgo Gonzalez et al. (2019). By contrast, the LCP data reports the total rated thermal input of each installation part.

We use information on thermal input to identify exact matches between individual ENTSO-E units and LCP installation parts whenever this is possible. For cases in which multiple ENTSO-E units potentially correspond to a single installation part, we aggregate the input of the relevant ENTSO-E units and compare it with the thermal input reported for the installation part. If the aggregated input is consistent with the reported input of the installation part, we classify the match as many-to-one. Conversely, when a single ENTSO-E unit potentially corresponds to multiple installation parts, we aggregate the input of the relevant installation parts and compare it with the input of the ENTSO-E unit. If the aggregated input is consistent with the unit-level input, we classify the match as one-to-many. When a power plant is comprised of multiple ENTSO-E units and the corresponding E-PRTR facility is comprised of multiple installation parts, and no unique correspondence can be established based on the input information and other available characteristics, we classify the unit-level match as undetermined and work with the plant-level match instead. To resolve ambiguous cases, we use additional plant and unit level information from the Global Energy Monitor (GEM).[30] Generating units that cannot be matched to any E-PRTR facility or LCP installation are classified as unmatched.

Co-emission intensities for NO_x, and SO_2 We define co-pollution intensities as the ratio of annual pollutant emissions to estimated annual CO_2 emissions.

[30]For example, Bełchatów Power Station has four installation parts in the LCP dataset but thirteen generating units in ENTSO-E, making a direct one-to-one match difficult. By comparing the commissioning years reported by GEM with those in the LCP dataset, we can identify a specific match between a generating unit and an installation part. This additional information helps narrow down the remaining possible matches, which we then establish by comparing the aggregated estimated thermal inputs of selected generating units with the reported thermal inputs of the remaining three LCP installation parts.

Table A.1: Fuel-Specific CO_2 Emission Factors

Fuel type	CO_2 emission factor (tCO_2/TJ)	(tCO_2/MWh)
Other Gases	152	0.5472
Biomass	112	0.4032
Lignite	**109.1**	**0.3928**
Other Solid Fuels	107	0.3852
Peat	106	0.3816
Coal	**93.6**	**0.3370**
Liquid Fuels	75.7	0.2725
Natural Gas	**55.9**	**0.2012**

Notes: CO_2 emission factors expressed in tonnes of CO_2 per terajoule (TJ) and per megawatt-hour (MWh) of thermal energy input. Emission factors in column 1 are based on Juhrich (2016). Conversion uses 1 TJ = 277.78 MWh, equivalent to multiplying the original emission factors by 0.0036.

For each installation part l in the LCP dataset, we use annual data on emissions of pollutant $p \in NO_x, SO_2$, as well as on energy inputs x_f of fuel type f. Using fuel-specific CO_2 emissions factors ξ_f reported in Table A.1, we estimate CO_2 emissions in year t as

$$\widehat{CO2}_l(t) = \sum_f x_{l,f}(t) \cdot \xi_f. \tag{A.1}$$

We then use the matching between LCP installation parts, indexed by l, and ENTSO-E generating units, indexed by i, to construct unit-level co-pollution intensities. We distinguish between the following cases.

All generating units $i \in U(l)$ that are matched to the same LCP installation part l are assigned the same co-pollution intensity in year t:

$$z_{i,p}(t) = \frac{EMISSIONS_{l,p}(t)}{\widehat{CO2}_l(t)} \qquad \forall i \in U(l). \tag{A.2}$$

A one-to-one match is the special case where $|U(l)| = 1$.

If multiple LCP installation parts $l \in M(i)$ are matched to the same generating unit i, we aggregate across l before calculating the intensity:

$$z_{i,p}(t) = \frac{\sum_{l \in M(i)} EMISSIONS_{l,p}(t)}{\sum_{l \in M(i)} \widehat{CO2}_l(t)}, \tag{A.3}$$

Many-to-many matches arise if multiple generating units can be linked to multiple LCP installation parts, but the exact matches between these entities cannot be established. Let g index the matching groups that are formed by such many-to-many matches. Let $M(g)$ be the set of LCP installation parts in group g, and $U(g)$ is the set of ENTSO-E generating units in that group. We aggregate emissions across all installation parts that belong to an installation identified by $M(g)$ to compute a common co-pollution intensity which is assigned

to all generating units in $U(g)$,

$$z_{i,p}(t) = \frac{\sum_{l \in M(g)} EMISSIONS_{l,p}(t)}{\sum_{l \in M(g)} \widehat{CO2}_l(t)}, \qquad \forall i \in U(g), \tag{A.4}$$

To generating units without a usable match we assign the average co-pollution intensity computed among successfully matched units with the same primary fuel in the same country and year:

$$z_{i,p}(t) = \frac{\sum_{j \in J(c,f,t)} EMISSIONS_{j,p}(t)}{\sum_{j \in J(c,f,t)} \widehat{CO2}_j(t)}, \tag{A.5}$$

where c and f are the country and primary fuel type of unit i, and $J(c, f, t)$ denotes the set of matched generating units in that country with that primary fuel type that report both pollutant p and fuel inputs in year t, and $EMISSIONS_{j,p}(t)$ and $\widehat{CO2}_j(t)$ are the values of the LCP installation parts linked to unit j. Missing years are filled from the nearest year of the same country and fuel type; where a country has no data for a fuel type, we use the EU-wide average for that fuel and year. We apply these averages to units without any LCP data and to five gas-fired units whose LCP installation part does not reflect their own emissions. Four are linked to parts dominated by coal, oil or steelworks gases, and one to the part of a refinery that also burns refinery oil. We exclude these five units from the averages.

We correct outliers in the resulting intensities without dropping observations. Intensities are unreliable when a unit's estimated CO_2 emissions are very low. If they fall below 1,000 tonnes in a year, we use the unit's intensity from the previous year or, failing that, its median across years in which it is observed. If emissions fall below 10% of the unit's median annual CO_2 emissions, we use the unit's median intensity over its remaining years. We then winsorize intensities at the 99th percentile, separately by pollutant. Finally, we fill missing years with the unit's nearest available year.

Co-emission intensities for NH_3 Information on NH_3 emissions is available only at the facility level from E-PRTR data, which also contains CO_2 emissions. We compute the NH_3 intensity of a facility w directly from the data (i.e., without estimating CO_2 emissions) as

$$z_{w,p}(t) = \frac{EMISSIONS_{w,p}(t)}{CO2_w(t)}. \tag{A.6}$$

Let $U(w)$ denote the set of ENTSO-E units linked to E-PRTR facility w. We assign

$$z_{i,p,t} = \frac{EMISSIONS_{w,p,t}}{CO2_{w,t}} \quad \forall i \in U(w) \wedge p = NH_3 \tag{A.7}$$

E-PRTR sets a reporting threshold for air releases of NH_3 at 10,000 kg. Facilities near that threshold can report in some years and not in others. When a facility × year pair has a missing value of z for lack of reported NH_3 emissions, we first impute missing values of z from the facility's nearest observed year. We do so only if the implied $NH_3 = CO_2 \cdot z$ is below the reporting threshold of 10,000 kg; otherwise, we set NH_3 and z back to missing because it is unlikely to be missing due to the reporting threshold. Next, if z is missing in year t but not in

Table A.2: Matching Statistics of ENTSO-E Generating Units

Matching level	Number of units	Percentage (%)
Installation part	703	96.04
Installation only	21	2.87
Facility only	6	0.82
Unmatched	2	0.27
Total	732	100.00

Table A.3: Summary statistics of annual co-pollutant emissions

Pollutant	Units	N	Mean	SD	Median	Min	Max
NO_x	732	3,963	458.078	707.836	190.436	0.000	6,570.217
SO_2	732	3,963	287.038	634.035	3.370	0.000	7,222.565
NH_3	732	3,963	2.928	4.618	1.774	0.000	71.262

Notes: The table reports summary statistics of estimated annual pollutant emissions at the ENTSO-E generating-unit level. Emissions are measured in tonnes per year. Units denotes the number of distinct generating units, N the number of unit-year observations. SD denotes the standard deviation.

$t-1$ and $t+1$, we linearly interpolate using those two values. If there are gaps before the first or after the last value, we carry backward or forward the respective nearest value.

Some facilities never report NH_3 in any year. In those cases, we assume NH_3 emissions of 5,000 kg (i.e., half the threshold value) and calculate the resulting z based on reported CO_2 emissions. We then cap those z's at the highest intensity obtained in any of the steps described before, which can be regarded as a plausible maximum.

Descriptive statistics Table A.2 summarizes the matching results across different levels. Of the 732 ENTSO-E generating units in our sample, 703 units (96%) are matched at the LCP installation-part level, 21 units (2.9%) at the installation level only, and 6 units (0.9%) at the facility level only. The remaining two units (0.3%) could not be matched at any level. Overall, 99.7% of generating units are successfully matched to at least an E-PRTR facility.

Table A.3 reports summary statistics for annual emissions of $PM_{2.5}$ precursors at the generating units in our analysis sample. Mean NO_x emissions are 458.1 tonnes, with a median of only 190.4 tonnes, indicating a highly skewed distribution driven by a small number of very large emitters. Similarly, SO_2 emissions have a mean of 287 tonnes but a median of just 3.4 tonnes, suggesting that most plants emit relatively little while a few account for disproportionately high emissions. The mean NH_3 emissions (including imputed values as described above) are 2.9 tonnes, with a maximum of 71.3 tonnes. Compared with the other pollutants, the distribution of NH_3 emissions is less right-skewed.

Table A.4 reports estimated co-emission intensities in kg per tonne of CO_2 emitted. Estimated CO_2 exhibits substantial variation, with a mean of 739,000 tonnes and a median of 466,515 tonnes, and a long upper tail reaching approximately 7,470,000 tonnes. Median intensities for NO_x and NH_3 are smaller than the mean intensity (0.5232 and 0.0054 vs. 0.6569 and 0.0087, respectively), but there is considerable variability across units and over time, with some units being much more polluting. The distribution of SO_2 intensities is particularly heterogeneous, co-emission intensities ranging between zero and 3.1925, as well as a mean

Table A.4: Summary statistics of annual co-emission intensities

Pollutant	Units	N	Mean	SD	Median	Min	Max
NO_x	732	4,392	0.6569	0.5177	0.5232	0.0039	2.9786
SO_2	732	4,392	0.3480	0.6139	0.0326	0.0000	3.1925
NH_3	732	4,392	0.0087	0.0115	0.0054	0.0002	0.1377

Notes: Annual emission intensities are measured in kg of pollutant per tonne of estimated CO_2 emissions. 2016–2021. Units denotes the number of distinct power units, N the number of unit-year observations.

(0.3480) that exceeds the median (0.0326) by an order of magnitude. This pattern highlights the scope for carbon trading to give rise to implicit co-pollution trades that are far from one-for-one.

B Atmospheric Dispersion and Chemical Transport

We employ the chemical transport model (CTM) developed and applied by Gu et al. (2023) to quantify the sources of $PM_{2.5}$-related health risks in Europe. CTMs simulate the emission, chemical transformation, transport, and removal of atmospheric pollutants using meteorological inputs and emission inventories. Gu et al. (2023) adapt the global atmospheric chemistry model GEOS-Chem with its adjoint sensitivity model and high-resolution emission data to link emissions from specific locations and sectors in Europe to health impacts. In this way, the authors can attribute $PM_{2.5}$ exposure and associated premature deaths in Europe to specific emission sources, including countries, sectors, and precursor species. The model framework comprises several steps.

1. Forward GEOS-Chem Simulation The forward CTM simulates atmospheric processes driving $PM_{2.5}$ concentrations across Europe. Starting with emissions of precursor pollutants (NO_x, NH_3, SO_2, organic carbon, black carbon, etc.), the forward simulation comprises the advection and diffusion of pollutants by meteorology, the atmospheric chemical reactions forming secondary aerosols, as well as deposition processes that remove pollutants from the atmosphere.

The model is run on a nested European grid ($0.25^\circ \times 0.3125^\circ$ resolution) using the HTAPv3 emissions inventory to capture spatial variability in emissions and pollution. The raw simulation output consists of gridded $PM_{2.5}$ concentrations that, combined with population density, results in gridded population exposure to fine dust.

To reduce model biases, the framework integrates satellite-derived $PM_{2.5}$ concentrations which combine aerosol optical depth measurements with models to estimate surface $PM_{2.5}$ concentrations (van Donkelaar et al., 2021). This helps to correct systematic model biases across regions. A downscaling component is used to improve the model's capability to capture sub-grid variability, providing more spatial information than the model can resolve for, thus mapping surface $PM_{2.5}$ concentrations on a $0.01^\circ \times 0.01^\circ$ resolution (Nawaz and Henze, 2020). This step significantly improves the representation of fine spatial variability in population exposure when combined with fine resolution (1 km) gridded population data from the Center For International Earth Science Information Network-CIESIN (2018).

2. Health Impact Assessment The model converts simulated $PM_{2.5}$ exposure into premature mortality estimates using concentration–response relationships from epidemiological studies. Specifically, Gu et al. (2023) combine population-weighted $PM_{2.5}$ exposure from the CTM with age-stratified exposure-response functions for Global Burden of Disease (GBD) mortality covering chronic obstructive pulmonary disorder (COPD), ischemic heart disease (IHD), lower respiratory illnesses (LRI), lung cancer (LC), type-II diabetes (T2) and stroke (Murray et al., 2020). Aggregate mortality impacts across Europe are summarized in the cost function

$$J_{PM_{2.5}} = \sum_{L} \sum_{A} \sum_{k \in D} \sum_{(I,J) \in k} (POP_{I,J,A} \times MOR_{I,J,A,L} \times AF_{I,J,A,L}) \tag{B.1}$$

where $AF_{I,J,A,L} = \frac{RR_{I,J,A,L}-1}{RR_{I,J,A,L}}$ denotes the attributable fraction, i.e., the excess mortality within a grid cell (I, J) and age bracket A that is due to disease $L \in$\{COPD, IHD, LRI, LC, T2D, stroke\}

3. Adjoint Sensitivity Model A central feature of the framework is the GEOS-Chem adjoint (Henze et al., 2007), which calculates the sensitivity of a cost function (here, $PM_{2.5}$-related premature deaths) to changes in emissions. Instead of running many forward simulations with perturbed emissions, the adjoint CTM offers a much more efficient way of computing sensitivities of the cost function (B.1) to all emissions simultaneously. The goal is to determine how emissions influence this quantity.

To explain this in a nutshell, consider the forward CTM simulation. The chemical transport model simulates pollutant emissions, transport, chemistry, and deposition:

$$\mathbf{x}_{t+1} = \mathcal{M}_t(\mathbf{x}_t, \mathbf{E}_t)$$

where $\mathbf{x}_t$ denotes pollution concentrations, $\mathbf{E}_t$) denotes emissions and $\mathcal{M}_t$ is the model operator describing transport and chemistry.

The object of interest is the gradient of the objective function with respect to emissions:

$$\frac{\partial J}{\partial E_i}$$

At the end of the forward simulation, the adjoint variables are initialized based on how the objective depends on the model state:

$$\boldsymbol{\lambda}_T = \frac{\partial J}{\partial \mathbf{x}_T}$$

The adjoint model propagates sensitivities backward in time through the CTM processes:

$$\boldsymbol{\lambda}_t = \left(\frac{\partial \mathcal{M}_t}{\partial \mathbf{x}_t}\right)^T \boldsymbol{\lambda}_{t+1}$$

This step accounts for how chemistry, transport, and deposition link earlier emissions to later concentrations. During the backward propagation, the sensitivity of the objective to emissions is accumulated:

$$\frac{\partial J}{\partial \mathbf{E}_t} = \left(\frac{\partial \mathcal{M}_t}{\partial \mathbf{E}_t}\right)^T \boldsymbol{\lambda}_{t+1} \tag{B.2}$$

For small emission changes, the impact on the objective can be approximated using a linear expansion:

$$\Delta J \approx \sum_i \frac{\partial J}{\partial E_i} \Delta E_i$$

The formula allows researchers to attribute $PM_{2.5}$-related deaths to specific emission sources (countries, sectors, species). Doing this via the backward model run is computationally much more efficient (more than 11,500 times faster) than forward-modeling based calculations whereby emissions perturbations at different sources are simulated, one by one, in the forward simulation

CTM.

4. Mortality and source attribution at baseline Gu et al. (2023) estimate about 449,813 $PM_{2.5}$-related premature deaths in Europe in 2015, relative to a total population of 598,970,000 in the domain.This amounts to a 46% reduction compared to 2005 when the EU ETS was launched. Roughly 59% of those premature deaths are attributable to domestic anthropogenic emissions, and 265,328 deaths (59%) are due to anthropogenic NO_x, NH_3, SO_2, OC, BC, SOAP emissions. The energy sector accounts for 21,186 (11%) of $PM_{2.5}$-related premature deaths.

C The random forest model

C.1 Goodness of fit

To assess the predictive quality of our model, we conduct out-of-sample validation. We randomly partition the 52,608 hours-of-sample (24 hours × 365 days × six years, plus two leap days in 2016 and 2020) fifty-fifty into equal-sized training and test sets. Following standard practice, we build the model on the training data, and then assess its predictive quality on the out-of-sample test data.

Table C.1: Goodness-Of-Fit Measures By Fuel Type

	All Units	Lignite	Hard Coal	Natural Gas	Border Points
	1. Weighted Percentile Prediction				
MSE	0.040	0.045	0.046	0.037	0.028
R^2	0.729	0.554	0.672	0.734	0.730
Corr	0.864	0.785	0.835	0.867	0.860
	2. Mean Prediction				
MSE	0.073	0.063	0.091	0.066	0.064
R^2	0.506	0.379	0.344	0.531	0.383
Corr	0.714	0.625	0.626	0.738	0.622

Notes: The table shows three metrics for quantifying prediction quality of the random forest model: 1. Mean Squared Error, 2. R-squared, and 3. the Correlation Coefficient. This table shows all three metrics for all units, and then for the four unit types (lignite, hard coal, natural gas, and border points). We compare the prediction at the weighted percentile to maintain a market equilibrium with the mean prediction. Numbers relate to one-fold cross validation, using 50% of hours-of-sample to train the model and assessing it using the other 50% (test data).

Table C.1 shows the performance of the QRF algorithm to predict utilization rates in the test set. The upper panel presents goodness-of-fit measures using our weighted percentile predictions that keep aggregate generation constant. The MSE of predicted relative output is 0.040, with an R^2 of 0.729, and a correlation coefficient of 0.864. These metrics highlight the good predictive quality of our model, especially when considering that the utilization rate is continuous and bounded within the $[0, 1]$ interval with many zero and near-one observations. This in turn requires a great deal of non-linearity for the model to correctly predict utilization rates.

The lower panel shows the goodness-of-fit measures if we simply used the model's mean predictions, disposing of the model's Quantile Regression part and disregarding equilibrium violations. Clearly, the weighted percentile predictions outperform the mean predictions: The MSE of the latter is almost twice as large, with a particularly pronounced difference for hard coal units and border points, for which the weighted percentile predictions work particularly well.

Comparing performance measures across types of units, the model works particularly well

Table C.2: Goodness-Of-Fit Measures By Country (sorted by MSE, ascending)

	AT	SE	DE	BE	GB	DK	IT	FR	NL	ES	CZ	PL	PT
	Weighted Percentile Prediction												
MSE	0.016	0.017	0.023	0.026	0.026	0.029	0.037	0.038	0.043	0.044	0.047	0.067	0.074
R^2	0.879	0.831	0.852	0.830	0.827	0.758	0.715	0.761	0.710	0.627	0.649	0.441	0.556
Corr	0.940	0.915	0.926	0.915	0.915	0.878	0.857	0.880	0.857	0.810	0.826	0.720	0.777
	Mean Prediction												
MSE	0.061	0.051	0.060	0.066	0.064	0.087	0.073	0.066	0.075	0.088	0.073	0.076	0.088
R^2	0.531	0.492	0.618	0.574	0.583	0.272	0.439	0.585	0.498	0.249	0.451	0.363	0.461
Corr	0.773	0.710	0.794	0.761	0.759	0.554	0.672	0.778	0.712	0.517	0.683	0.615	0.595

Notes: The table shows three metrics for quantifying prediction quality of the model: 1. Mean Squared Error, 2. R-squared, and 3. the Correlation Coefficient, separately for units in a given country. We compare the prediction at the weighted percentile to maintain a market equilibrium with the mean prediction. Numbers relate to one-fold cross validation, using 50% of hours-of-sample to train the model and assessing it using the other 50% (test data).

for natural gas units and border points, whereas it struggles a bit more with hard coal and lignite units. The latter have a substantially low R-squared of 0.554, which may reflect that lignite units are relatively sluggish and therefore tend to run either near full capacity or not at all. This would suppress R^2 because actual utilization rates vary relatively little with predictors. The relatively more flexible hard coal units may also be hard to predict, yet R^2 will be higher if the model still captures general generation patterns reasonably well.

We further document substantial differences in model performance across countries. Table C.2 shows performance measures by country, by MSE in ascending order. One can see that the model mostly performs very well for small countries with high-capacity interconnections to other countries. Units in Austria, Sweden,[31] Belgium and Denmark have small MSEs of below 0.03.

The model also works well for Germany, a country characterized by its decentralized geography, high degree of interconnectedness with other countries, and a very diverse energy mix. Therefore, price fluctuations plausibly induce large changes in generation patterns from which the model can learn. It also works well for Great Britain despite their less diverse energy mix.

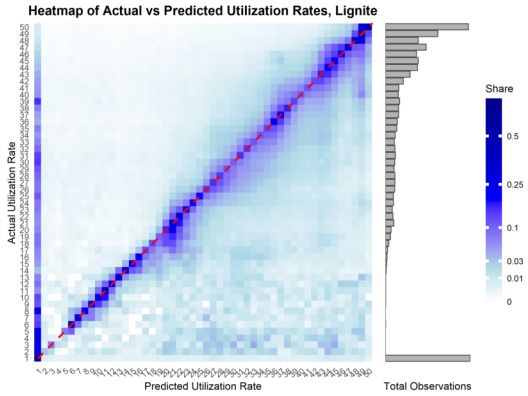
Conversely, the model struggles more in countries that rely on fewer fuel types. The Netherlands, albeit a small and interconnected country, relies heavily on natural gas to serve base load, in contrast to most other countries where natural gas is a peak demand fuel.

In Spain, the highly uneven population density coupled with insufficient transmission capacities plausibly causes significant friction in the Spanish energy market, weakening the predictive power of the merit-order principle. Lastly, Poland and Czechia depend on lignite and hard coal to a much higher degree than any other country in the sample. Accordingly, the marginal units in both Poland and Czechia are mostly coal units even under near-zero allowance prices, which is a phenomenon unique to these countries that the model might struggle to fully internalize. The model also struggles relatively heavily with Portugal.

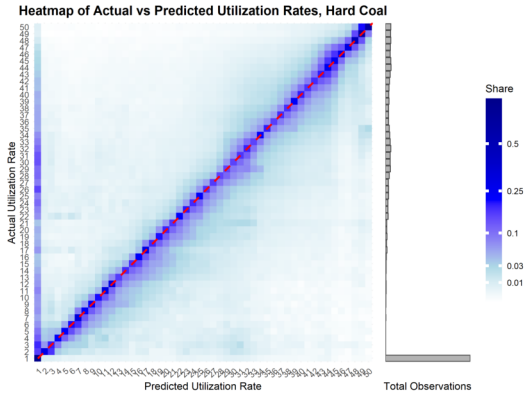
Next, we compare the predictions of units' utilization rates with its actual levels. Figure C.1

[31] Sweden is considered "small" in this context because only its two southern-most bidding zones contain fossil fuel power units.

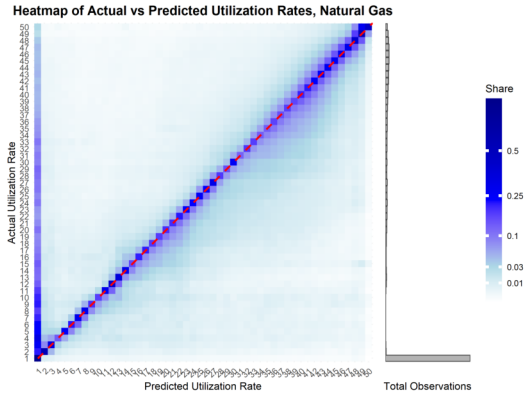
Figure C.1: Actual vs. Predicted Utilization Rates, Weighted Percentile Predictions



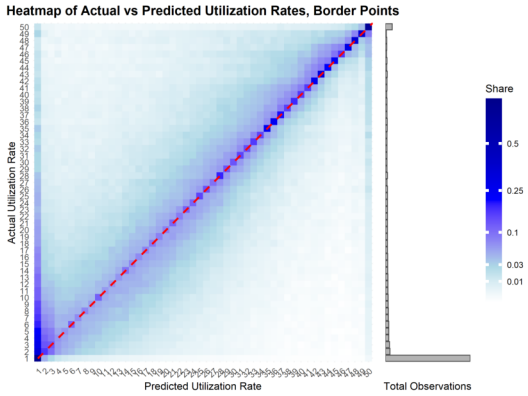
(a) Actual (y-axis) vs. Predicted (x-axis) Utilization Rates, Lignite Units



(b) Actual (y-axis) vs. Predicted (x-axis) Utilization Rates, Hard Coal Units



(c) Actual (y-axis) vs. Predicted (x-axis) Utilization Rates, Natural Gas Units



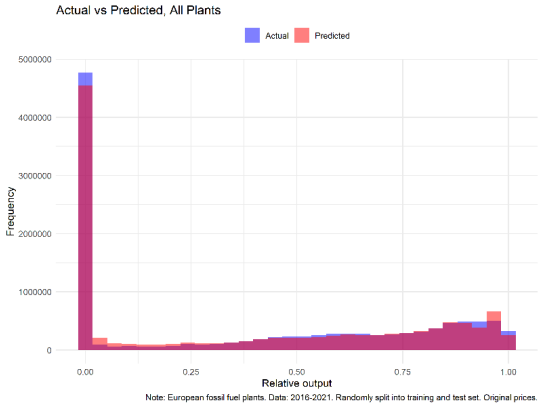
(d) Actual (y-axis) vs. Predicted (x-axis) Utilization Rates, Border In Points

Notes: The graphs show a heatmap of actual utilization rates (y-axis) and predicted utilization rates (x-axis) using the weighted percentile predictions. Utilization rates are split into fifty bins in steps of two percentage points. The heatmaps are such that one row adds up to one. In other words, we ask "what is the distribution of predicted utilization rates when actual utilization rates are in bin x". Darker colors indicate a higher density of predictions within the respective row. The bars to the right of the heatmap show the number of observations that contribute to every row. The graphs relate to one-fold cross validation, using 50% of hours-of-sample to train the model and assessing it using the other 50% (test data).

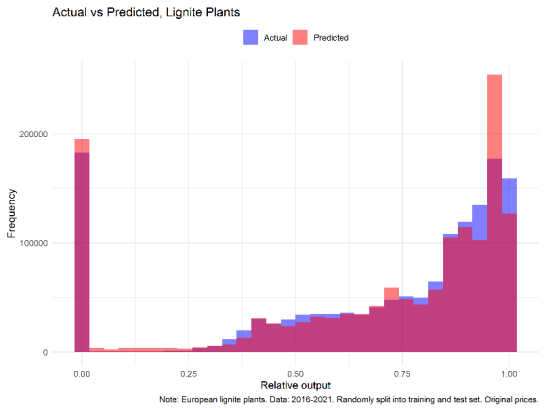
presents heatmaps of actual (y-axis) vs. predicted (x-axis) rates in which we bin the outcome in bins of 0.02. The first bin therefore corresponds to utilization rates between zero and two percent. Every row shows the distribution of predicted utilization rates within the respective bin of actual rates, so that the graph should be read row-to-row. Figures C.1a to C.1d show the heatmap for lignite units (C.1a), hard coal units (C.1b), natural gas units (C.1c) and border inflows (C.1d). The heatmaps provide us with some important insights into the model's performance. First, the model predicts utilization rates with a high degree of accuracy as most observations lie on or very close to the 45 degree line. Second, there is a tendency to predict near-zero utilization rates among fossil fuel units even when actual rates are much higher. This happens along upper-medium rates of 50 to 80% for lignite and coal units, and along the entire range of utilization rates for natural gas units, whereas this pattern does not occur for border points. Beyond that, we rarely observe the opposite that predicted rates are far from zero while actual rates are near-zero.

Third, the model performs very well when actual utilization rates are near zero or near one. The model correctly predicts utilization rates of below 10 percent 76.4% of the time for lignite units, 84.1% for hard coal units, 91.4% for natural gas units, and 87.2% for border points. For actual utilization rates north of 90 percent, the model achieves a correct prediction 82.0% of the time for lignite units, 75.1% for hard coal units, 70.4% for natural gas units, and 66.9% for border points. Fourth, overall model precision is higher for natural gas units and border points than for lignite and hard coal units. Predictions are scattered more closely around the 45 degree line for gas units and border points, whereas they spread out a bit more for lignite and hard coal units. In addition to that, we have a notable quantity of cases in which predicted rates are substantially above actual rates for lignite units.

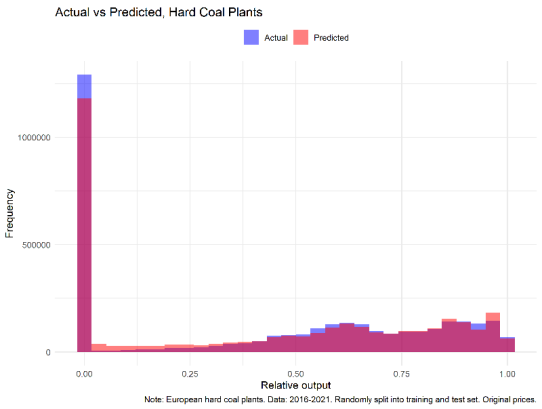
We also compare the histograms of actual and predicted utilization rates in Figure C.2. The alignment between the actual and predicted histograms suggests that the model effectively recovers the underlying empirical distribution of utilization rates. Individual prediction errors appear to be non-systematic and tend to cancel out in the aggregate, ensuring that the predicted frequency of utilization rates remains consistent with observed market behavior.



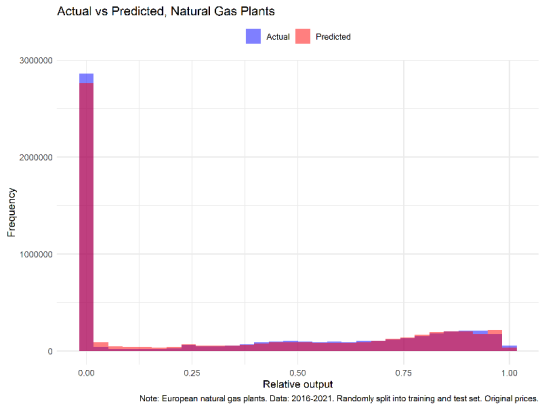
(a) Actual (y-axis) vs. Predicted (x-axis) Utilization Rates, Weighted Percentile Prediction, All Units



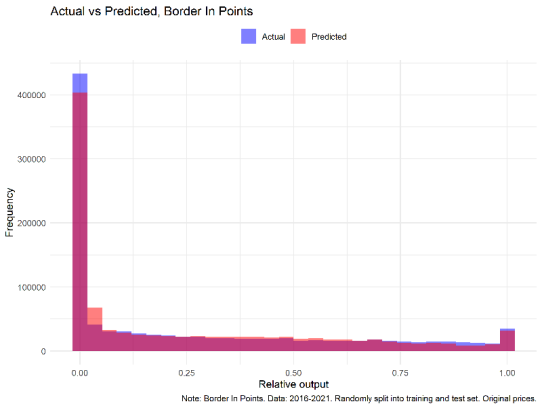
(b) Actual (y-axis) vs. Predicted (x-axis) Utilization Rates, Weighted Percentile Prediction, Lignite Units



(c) Actual (y-axis) vs. Predicted (x-axis) Utilization Rates, Weighted Percentile Prediction, Hard Coal Units



(d) Actual (y-axis) vs. Predicted (x-axis) Utilization Rates, Weighted Percentile Prediction, Natural Gas Units



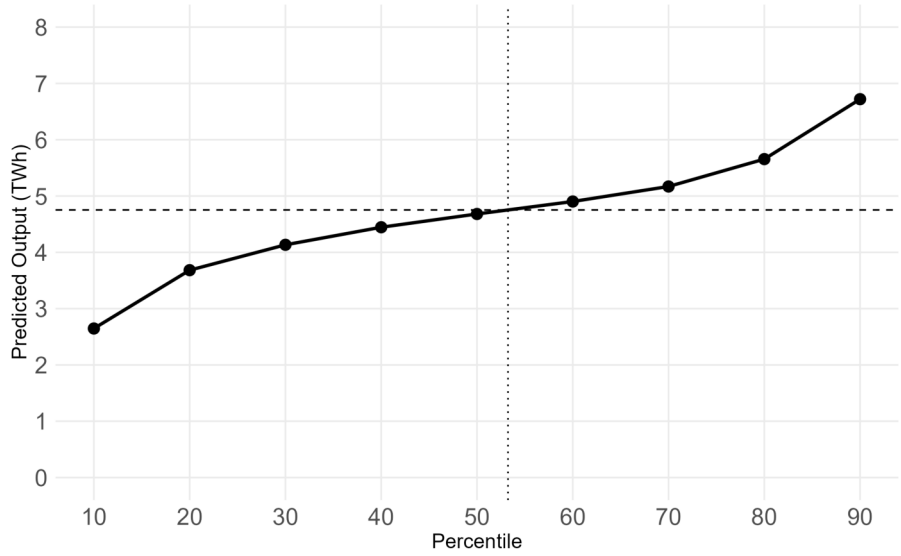
(e) Actual (y-axis) vs. Predicted (x-axis) Utilization Rates, Weighted Percentile Prediction, Border In Points

Figure C.2: Actual vs. Predicted Relative Output, Weighted Percentile Prediction

Notes: The graphs show histograms of actual relative output (blue) and predicted relative output (red) using the weighted percentile predictions for all units (a) as well as for the four unit types lignite (b), hard coal (c), natural gas (d), and border in points (e).

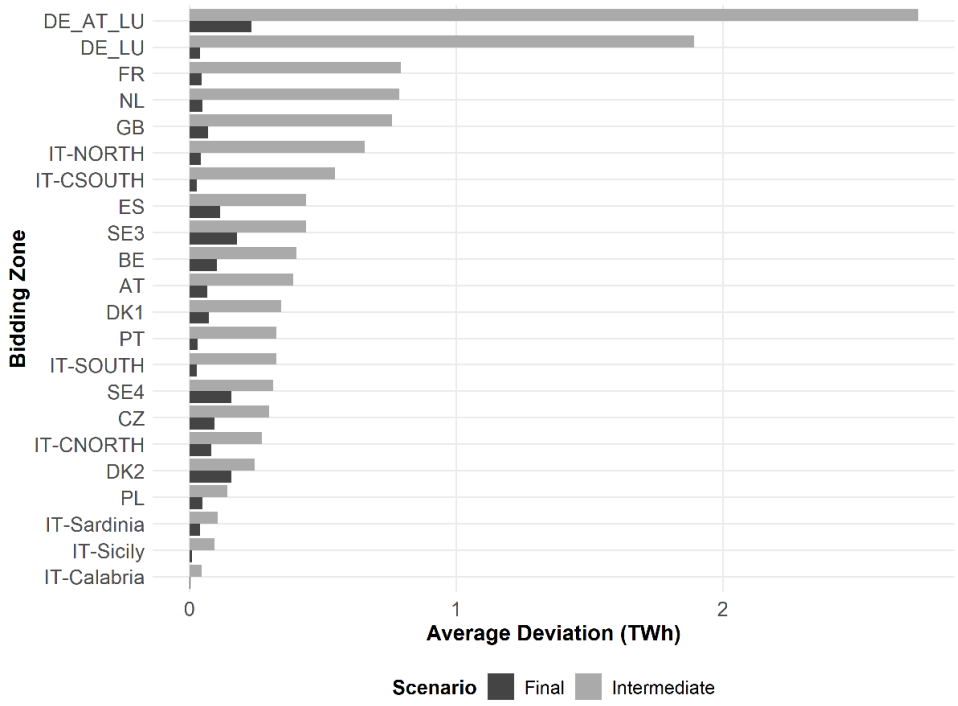
C.2 Additional tables and figures

Figure C.3: Predicted Output of Hard Coal Power Units in the German-Austrian-Luxembourgian Bidding Zone in December 2017



Notes: Hard-coal fired power units contained in the test data and located in the German-Austrian-Luxembourgian bidding zone jointly produced 4,752,709 MWh during December 2017, as indicated by the dashed horizontal line. The dots indicate deciles of predicted total output for this period, obtained from the quantile regression random forest model. Equilibrium requires that predicted generation by all units add up to total observed output. We use linear interpolation between deciles to satisfy this constraint. Total output lies between the 5th and 6th decile. The 50th percentile under-predicts output (∼4,680,000 MWh) whereas the 60th percentile over-predicts it (∼4,900,000 MWh). Algebraically, $\widehat{Y}_{b,f,m}(0.5) < Y_{b,f,m}$ and $\widehat{Y}_{b,f,m}(0.6) > Y_{b,f,m}$. Therefore, choosing $\alpha_{b,f,m} = 0.3239$ yields a weighted average of the 50th and 60th percentile predictions that brings us exactly to $Y_{b,f,m} = 4{,}752{,}709$ MWh. We also take this as the linearly interpolated prediction at the 53.239th percentile.

Figure C.4: Average absolute deviation between predicted and actual fossil demand



Notes: The figure shows, for every zone, by how much monthly fossil demand deviates on average from predicted fossil supply in the high price counterfactual with a €50 carbon price floor. The graph compares (i) the "Intermediate" solution in which we keep domestic fossil generation constant and ignore the cross-border flow predictions to (ii) the "Final" solution in which we scale percentiles to converge to the actual counterfactual equilibrium in which the gap between predicted fossil supply and fossil demand across all zones is minimized. It shows that in the German-Austrian-Luxembourgian bidding zone, for instance, predicted fossil supply and fossil demand deviate from each other by on average 2.73 TWh per month in the intermediate solution, but only by 0.23 TWh in the final solution. Relative to on average 17 TWh of monthly fossil demand, the final solution therefore gets predicted fossil supply wrong by a marginal 1.35% on average.

Table C.3: Percentiles per bidding zone and fuel type, averaged across month-of-sample

Bidding zone	**B**	**L**	**C**	**G**
Austria	55.1		53.0	65.6
Belgium	66.7		64.0	54.6
Czechia	58.4	52.4	70.4	61.1
DE-AT-LU	67.7	39.4	52.1	63.1
DE-LU	66.9	51.2	58.7	61.1
Denmark 1	67.4		56.1	66.5
Denmark 2	53.7		50.8	53.9
Spain	49.2	66.6	67.0	77.3
France	67.9		51.7	52.1
Great Britain	48.4		47.2	55.0
IT C North	52.4			54.1
IT C South	50.1		56.9	75.8
IT Calabria	72.9			66.7
IT North	52.2		61.8	65.7
IT Sardinia	74.1		61.3	
IT Sicily	40.2			67.9
IT South	68.1		76.7	65.8
Netherlands	65.0		42.6	65.5
Poland	54.1	57.0	71.0	49.9
Portugal	62.4		57.0	73.2
Sweden 3	64.7		45.1	63.9
Sweden 4	51.4			53.7

Notes: Averaged across months-of-sample, the table shows the (interpolated) percentile in every bidding zone that we use (on the zone-fuel-month level) under actual prices to equate actual total output with predicted total output. B = Border Point. L = Lignite. C = Hard Coal. G = Natural Gas.

D Additional results

Table D.1: Change in Electricity Generation and Flows (TWh) due to the EU ETS (Actual Prices)

Country	Lignite	Hard Coal	Natural Gas	Total Fossil	Imports	Exports
Germany	-21.5	-48.7	29.0	-41.2	43.0	-1.5
Poland	-5.8	-23.0	8.1	-20.6	22.4	-1.0
Italy	—	-61.0	58.4	-2.6	-0.9	0.5
Denmark	—	-2.6	0.9	-1.7	26.1	24.9
Czechia	-0.6	-3.6	4.3	0.1	5.2	3.7
Sweden	—	0.5	0.4	0.9	11.3	23.8
Portugal	—	-10.2	12.0	1.7	15.3	17.2
Austria	—	-0.4	5.8	5.4	4.2	6.2
Spain	0.0	-98.2	104.3	6.1	21.1	24.1
Great Britain	—	-6.3	15.7	9.5	-0.2	5.6
Belgium	—	-0.1	10.2	10.0	9.7	20.7
France	—	-2.5	18.1	15.6	4.2	23.2
Netherlands	—	-11.6	27.7	16.1	2.0	15.8

Notes: The numbers denote the changes in TWh when going from the €0 allowance price scenario to actual generation under actual allowance prices. Negative values in Imports/Exports indicate a decrease compared to the €0 scenario. Changes in total fossil generation plus net imports do not exactly sum to zero due to approximation error in the counterfactual equilibrium.

Table D.2: Change in Electricity Generation and Flows (TWh): €0 to €50 Scenario

Country	**Lignite**	**Hard Coal**	**Natural Gas**	**Total Fossil**	**Imports**	**Exports**
Germany	-144.0	-109.3	95.2	-158.1	118.6	-27.4
Poland	-17.2	-16.0	9.8	-23.5	23.6	-2.4
Portugal	—	-21.3	14.8	-6.5	12.7	5.2
Denmark	—	-1.0	2.7	1.7	55.4	60.9
Sweden	—	0.9	2.4	3.3	23.8	59.8
Czechia	-3.2	-4.0	11.4	4.1	10.7	6.7
France	—	-7.7	17.3	9.6	37.5	48.1
Italy	—	-137.5	148.0	10.5	-3.7	3.1
Spain	0.1	-179.5	196.9	17.5	17.8	33.6
Austria	—	-1.4	23.3	21.9	-0.9	3.7
Belgium	—	0.5	24.7	25.2	21.8	42.8
Great Britain	—	-18.5	54.5	35.9	-23.0	8.6
Netherlands	—	-33.2	90.5	57.4	-20.1	32.2

Notes: The numbers denote the changes in TWh when comparing a €0 carbon price scenario to a hypothetical €50 allowance price scenario. Changes in total fossil generation plus net imports do not exactly sum to zero due to compounding approximation errors in the counterfactual equilibria.

Table D.3: Electricity Generation by Fuel and Scenario (TWh)

	Lignite			**Hard Coal**			**Natural Gas**		
Country	€0	Actual	≥€50	€0	Actual	≥€50	€0	Actual	≥€50
Germany	626.7	605.1	482.7	372.2	323.5	262.9	155.2	184.2	250.3
Italy	—	—	—	183.5	122.5	46.0	539.4	597.8	687.4
Poland	228.3	222.5	211.1	324.6	301.7	308.6	29.8	38.0	39.6
Netherlands	—	—	—	144.3	132.7	111.1	185.9	213.6	276.4
Spain	—	—	—	224.4	126.2	45.0	96.9	201.2	293.8
Czechia	97.3	96.6	94.1	8.3	4.7	4.2	8.2	12.5	19.6
Belgium	—	—	—	0.5	0.3	1.0	83.3	93.4	108.0
France	—	—	—	24.2	21.7	16.5	107.7	125.9	125.1
Austria	—	—	—	4.8	4.3	3.4	29.1	34.9	52.4
Denmark	—	—	—	26.0	23.4	24.9	1.9	2.8	4.6
Sweden	—	—	—	0.4	1.0	1.4	2.0	2.4	4.4
Great Britain	—	—	—	47.2	40.9	28.6	533.7	549.4	588.1
Portugal	—	—	—	22.9	12.7	1.6	51.5	63.4	66.3

Notes: Numbers denote generated TWh of electricity by fuel across the countries in our sample, and across the three scenarios (€0 allowance price, actual allowance prices, €50 allowance price floor).

Table D.4: CO_2 Emissions by Scenario and Absolute Change (Mt)

	Scenarios (Mt)			Change vs. €0	
Country	**€0**	**Actual**	**≥€50**	**ΔActual**	**Δ≥€50**
Germany	1127.4	1067.1	891.6	-60.3	-235.8
Poland	591.1	564.9	557.7	-26.2	-33.5
Italy	382.1	345.4	307.0	-36.8	-75.2
Spain	256.4	199.2	156.3	-57.2	-100.1
Great Britain	253.8	253.9	256.0	0.2	2.3
Netherlands	196.2	197.8	203.8	1.6	7.7
Czechia	115.2	113.6	113.4	-1.6	-1.8
France	65.3	69.8	64.2	4.4	-1.1
Portugal	38.4	34.2	25.9	-4.2	-12.5
Belgium	33.8	37.7	43.9	3.9	10.1
Denmark	26.8	24.0	25.9	-2.7	-0.9
Austria	17.0	18.9	24.9	2.0	8.0
Sweden	1.5	2.4	3.8	0.9	2.3
Total	**3104.9**	**2928.9**	**2674.5**	**-176.0**	**-430.4**

Notes: Δ columns represent the absolute change in Mt relative to the €0 baseline. Negative values represent emission reductions.

Table D.5: Aggregate Emissions (Metric Tonnes) and % Changes relative to €0, 2016–2021

Scenario	NO_x	SO_2	NH_3
€0 Scenario	2,015,999	1,378,434	11,794
Actual Prices	1,809,271	1,138,830	11,265
	-10.3%	-17.4%	-4.5%
€50 Price Floor	1,557,153	896,095	11,967
	-22.8%	-35.0%	+1.5%

Table D.6: Sorting effects by country and fuel, 2016–2021

	€0 to actual EUA prices					€0 to €50-floor				
	Shift [Mt]	Intensity (i)	Damage (ii)	Joint (iii)	Sorting total	Shift [Mt]	Intensity (i)	Damage (ii)	Joint (iii)	Sorting total
Lignite										
Czechia	4.7	−11	5	−1	−7	20.6	−28	12	−2	−18
Germany	−7.6	16	42	−7	51	−52.8	71	73	−16	128
Poland	2.9	−49	0	−2	−52	31.9	74	15	5	94
Spain	0.0	0	0	0	0	0.2	0	0	0	0
Subtotal lignite	*0.0*	*−44*	*46*	*−10*	*−8*	*0.0*	*117*	*100*	*−12*	*205*
Hard coal										
Austria	0.8	−5	6	−3	−2	0.9	−6	4	−2	−4
Belgium	−0.1	1	0	0	1	0.9	−6	0	0	−5
Czechia	−2.9	0	−7	0	−7	−1.6	−1	−7	0	−7
Denmark	1.7	−29	−9	11	−26	12.2	−113	−41	38	−116
France	0.5	−1	−4	1	−4	0.7	7	−8	2	1
Germany	25.9	−104	15	7	−82	42.4	−174	7	20	−148
Great Britain	−0.4	22	−5	−6	11	−1.8	61	−6	−21	34
Italy	−29.8	159	−17	−7	135	−102.0	603	−57	−6	540
Netherlands	16.2	−96	−34	26	−104	29.5	−178	−62	48	−192
Poland	69.7	−87	145	−13	44	178.9	−157	413	−47	210
Portugal	−6.5	8	32	−5	35	−17.1	32	86	−19	99
Spain	−75.9	−628	359	299	30	−145.5	−932	720	424	212
Sweden	0.9	−6	−5	4	−8	2.5	−19	−15	11	−23
Subtotal hard coal	*0.0*	*−766*	*476*	*313*	*23*	*0.0*	*−882*	*1,035*	*448*	*601*
Natural gas										
Austria	0.3	1	1	1	2	3.1	−2	5	−1	2
Belgium	−1.1	4	1	0	6	−2.2	6	2	0	8
Czechia	0.8	−1	1	−1	−1	2.3	−3	4	−2	−1
Denmark	0.3	1	0	0	0	0.6	1	0	0	1
France	0.9	2	−2	1	1	−6.0	2	−4	0	−1
Germany	0.3	−4	2	−3	−5	8.8	−9	5	−5	−9
Great Britain	−20.8	3	12	1	15	−40.0	8	18	1	28
Italy	−8.9	11	15	3	28	−14.9	18	29	−2	45
Netherlands	−1.6	1	1	0	2	4.8	−6	3	0	−2
Poland	0.7	−1	0	0	−1	−0.4	−3	−3	−2	−7
Portugal	1.2	−2	−4	2	−5	−1.1	−1	0	1	0
Spain	27.9	1	−39	−2	−40	44.6	−5	−58	−1	−64
Sweden	0.1	0	0	0	0	0.5	−1	−1	0	−1
Subtotal natural gas	*0.0*	*14*	*−11*	*1*	*3*	*0.0*	*6*	*1*	*−10*	*−3*
Total	**0.0**	**−796**	**511**	**304**	**19**	**0.0**	**−760**	**1,136**	**427**	**803**

Notes: Contributions of each country's generating units to the sorting effects in eq. (22), in premature deaths, by fuel type. (i) pollution-intensity sorting, (ii) damage sorting, (iii) joint sorting. Shift is the CO_2 emissions reallocated to (+) or away from (−) a country's units within the same fuel, evaluated at the fuel's total CO_2 emissions under a zero carbon price; it reflects changes in within-fuel shares, not actual emission changes, and sums to zero within each fuel. Contributions are computed month by month at the unit level and summed. The country split of each sorting effect depends on the reference values $\bar{\mathbf{z}}$ and $\bar{\boldsymbol{\lambda}}$ (simple means across units of a fuel); the totals do not.

E Robustness & Threats to Identification

E.1 Fleet Composition and Entry/Exit

In our counterfactual analysis, we maintain a constant fleet of fossil fuel power units based on the composition observed during the 2016–2021 period. Thereby, we focus on the short-term impact of carbon pricing on electricity generation, in line with other literature on the electricity sector (Abrell et al., 2022; Butters et al., 2025).

The exit of fossil fuel power units is heavily influenced by regulatory frameworks independent of the carbon price mechanism. For instance, several coal plants were decommissioned because of their noncompliance with stringent limits on pollutant emissions in the EU Industrial Emissions Directive. Likewise, national coal phase-out trajectories are often politically mandated and negotiated years in advance, tied to fixed decommissioning dates, and designed to ensure long-run grid stability. This renders the exit of coal units inherently insensitive to the secondary market signals provided by the EU ETS within our observation window.

Similarly, the entry of new generation units is constrained. Due to the same coal phase-out trajectories, there has been no scope to construct new lignite or hard coal power units even under a zero carbon price. At the same time, the build-up of new natural gas units is predominantly a policy-mandated transition to provide system flexibility as the share of renewable energy rises, rather than a market-led expansion of the fossil fuel fleet. Beyond these policy constraints, the economic incentive for large-scale entry of new natural gas units appears limited by the significant latent capacity within the existing fleet. Our model predicts that even under a €50 allowance price floor, the average utilization rate for natural gas units across the sample only reaches 40%, up from 27% in the zero-price scenario. This indicates that the existing infrastructure remains far from a capacity threshold on average, suggesting that the system can accommodate the predicted fuel substitution. The only exception to this are Poland and Czechia in three and two years of our sample, respectively, where we observe predicted utilization rates of the natural gas fleet of above 70%, approaching the "practical maximum" (100% minus planned maintenance + forced outages + operating constraints). Still, local entry would only strengthen our results given that it would lead to broader fuel substitution and therefore larger effects. If anything this renders our estimates conservative.

Second, the decision to remain active in the energy market is primarily driven by short-run marginal costs. In the day-ahead wholesale market, operators place bids based on costs for fuel, carbon allowances, and variable operating expenditures since their capital and investment costs are already sunk. This implies that a power unit will continue to operate as long as the market clearing price exceeds its short-run marginal cost. An infamous example of this is the German hard coal unit Datteln IV. By the time it opened in 2020 with a nine-year delay, the German coal phase-out was already in progress, yet the new unit still opened because the operator (Uniper) argued that the sunk costs (approx. €1.5 billion) were too massive to write off, a decision that was clearly decoupled from the 2020 allowance price.

Third, the technical nature of existing fossil fuel infrastructure reinforces the stability of the fleet across different carbon price scenarios. The power plants themselves are high-value assets designed for multi-decade lifespans. By maintaining a constant fleet in our counterfactuals, we

isolate the transformative impact of the EU ETS on short-term dispatch optimization and fuel substitution without introducing the speculative bias of long-term investment modeling. Lastly, at least in the high allowance price counterfactual, the model should be able to capture exits of lignite and coal units to an extent given that it can consistently predict zero generation at these units.

E.2 Confounding Regulatory Policies

In addition to the EU ETS, there are various national policies in place in the countries in our sample. As we argue, fortunately, our observation period from January 2016 to December 2021 is characterized by exceptional stability in national "top-up" policies that allow us to isolate the effect of the ETS allowance price on electricity generation in our data. The UK's Carbon Price Support (CPS), a national carbon price on top of the EU ETS, has remained stable at £18 per tonne of CO2 throughout entire observation period, creating a time-invariant mark-up on marginal cost for UK power units. Sweden and Denmark have a national carbon tax, yet installations are generally exempt from this tax if they are already covered by the EU ETS.

E.3 Assumption of constant fuel prices

A potential concern in our counterfactual analysis is the assumption of constant fuel prices across different allowance price scenarios. One might argue that a significant increase in the EU-wide carbon price shifts demand sufficiently to put pressure on global fuel prices. However, we follow the approach of previous work such as Grossi et al. (2017), Abrell et al. (2022) and Jarvis et al. (2022) by treating fuel prices as exogenous to the electricity sector. This assumption is grounded in the structural reality that, during our 2016–2021 sample period, European fuel prices were largely determined by global benchmarks and long-term supply frameworks rather than local carbon-driven demand shifts.

During that time, European natural gas prices remained heavily influenced by global oil benchmarks through long-term, oil-indexed contracts. Since the oil market is the most liquid and globalized commodity market in the world, and the European fleet of oil-fired power plants is small, the EU ETS plausibly did not move the global oil equilibrium. Moreover, the supply side was characterized by significant structural rigidity; the physical nature of pipeline infrastructure and the prevalence of long-term, fixed contracts created a hold-up problem that insulated base fuel prices from fluctuations in the EU ETS.

E.4 Strategic Bidding and Market Power

A further potential threat to identification involves the assumption that power units always bid their true short-run marginal costs. We construct the supply curve and derive predictors for a unit's utilization rate from the supply curve based on our marginal cost estimates that reflect variable operating costs. In a perfectly competitive electricity market, this approach mirrors the true market dynamics since the merit-order principle dictates that operators bid their marginal costs of generation to ensure dispatch. However, in the presence of high market concentration, large utility firms may strategically submit bid above their true short-run marginal costs to

exercise market power. Although this implies that we misidentify a unit's position in the merit order based on marginal costs only, the Random Forest can account for this by internalizing strategic markups by training on realized dispatch.

A deeper threat exists if these markups are systematically correlated with allowance price movements. It might be the case that firms would have adjusted their exercise of market power in the counterfactual high allowance price scenario, for instance by increasing markups on gas units as coal becomes less competitive. In this case, our counterfactual predictions fail to capture this behavioral shift. However, given that our model focuses on the intensive margin of a highly integrated European grid, where cross-border competition often disciplines local market power, we treat these bidding strategies as a structural feature of the training data that remains reasonably stable across our counterfactual price ranges.

Likewise, the model is capable of "learning" the magnitude of systematic deviations where specific units consistently operate differently than their calculated marginal costs would suggest. By training on realized market behavior, the Random Forest endogenizes the dynamics of strategic markups. The identifying assumption is that the relationship between allowance prices and these markups is time-invariant. Under this assumption, the model leverages observed behavior from high-allowance-price periods to inform its predictions for counterfactual high-price scenarios during low-allowance-price periods. As long as the internalized markup dynamics during high-price periods are informative of how firms would have set markups in other times if the allowance price had been high as well, our counterfactual estimates provide a robust reflection of how the fossil fleet responds to carbon prices, rather than a stylized simulation based on the assumption of perfect competition.

We can draw on the relationship between the model's prediction error and the allowance price to assess whether our model systematically misses out strategic bidding. To consistently predict what power units would do if the carbon price was different (€0 or €50 price floor), we must assume that the model's prediction error is independent of our treatment, that is the carbon price itself and the derived metrics that we feed into the model (Abrell et al., 2022). If the model, for instance due to strategic bidding depending on the allowance price, systematically over-predicts output under low allowance prices and under-predicts under high allowance prices (or vice versa), our counterfactual results would be fundamentally biased.

We assess the assumption of the independence of the prediction errors with regard to the simulated carbon price in two ways. In Figure E.1a, we plot the prediction error (y-axis), that is the predicted utilization rate minus the actual utilization rate, against the ETS allowance price (x-axis). The dark red line shows how the average prediction error varies over different allowance prices, while the light red shaded area shows the standard deviation of the prediction error. Both metrics support the aforementioned assumption. The average prediction error fluctuates closely and in a seemingly random manner around zero across all values of the ETS allowance price. The predictive uncertainty also appears approximately constant: The standard deviation of the prediction error does not systematically grow or shrink with the allowance price.

In Figure E.1b, similar to how Abrell et al. (2022) assess this assumption in their analysis of the UK carbon tax, we calculate the Pearson correlation coefficient between the prediction error and the ETS allowance price for every power unit in our sample, and then plot the histogram of

these correlation coefficients. The graph shows that almost 82% of power units have correlation coefficients of no more than 0.1 in absolute size, which is a negligible amount of correlation. In other words, the carbon price is practically uncorrelated with the prediction error for the vast majority of the fossil fleet, so that the model's predictive precision does not change with the ETS allowance price. In addition to that, the symmetry of the histogram indicates that there is no global systematic bias in how the model handles carbon prices. Overall, the graphs in Figure E.1 show supportive evidence that the model's prediction error is not systematically related to the allowance price.

Beyond that, we also calculate the model's MSE and R^2 separately for every year. We can broadly partition the years of observations into periods of low prices (until early 2018), medium prices (mid-2018 to mid-2020), and high prices (since late 2020). If the predictive power of our model significantly changes between the years of observation, this would indicate that our model might fail to understand how power providers bid into the market and eventually provide electricity. As Table E.1 shows, this is not the case. Both the model's MSE as well as its R^2 are remarkably stable. The fact that neither metric changes substantially towards the end of the sample suggests that the model's predictive framework does not "break down" once the observed carbon prices approach unprecedented levels.

Table E.1: Random Forest Predictive Performance by Year

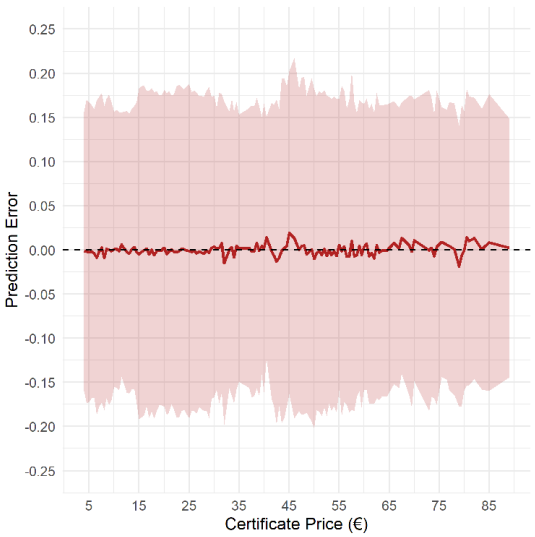
Year	**MSE**	R^2
2016	0.0268	0.8329
2017	0.0308	0.8098
2018	0.0315	0.8066
2019	0.0350	0.7756
2020	0.0303	0.7917
2021	0.0294	0.8067

Notes: The table reports the predictive performance of the Random Forest model by year. The R^2 and Mean Squared Error (MSE) are calculated based on actual vs. predicted unit-level utilization rates.

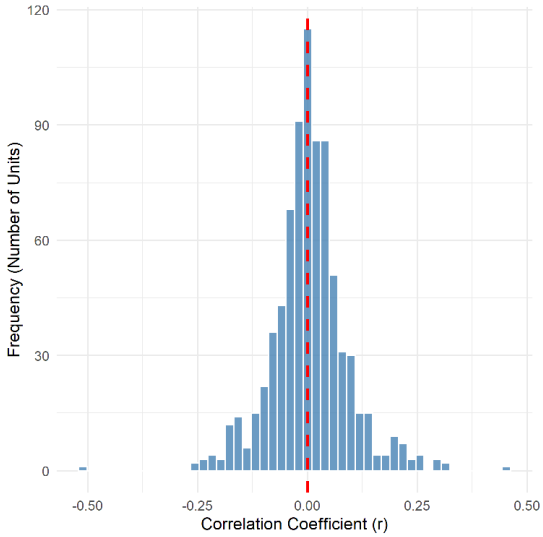
E.5 Merit Order vs. Actual Dispatch Order

Our empirical strategy feeds the Random Forest with information on the supply curve of fossil power units in the electricity market. Based on this information, the model aims to predict hourly utilization rates of power units. We derive a variety of metrics to describe a unit's relative position in the merit order, such as the capacity gap and the output gap as well as a unit's marginal cost and the zone's net load. In addition to that, time-constant unit-level information, such as their net electric efficiency, provides further information of a unit's relative position in the merit order.

The merit order establishes a market equilibrium on the day-ahead market. Unit operators submit supply bids specifying the volume of electricity they are prepared to generate at various price points for each hour of the following day. These bids are aggregated into a supply curve, and the market clearing mechanism subsequently allocates generation to units in ascending



(a) Average prediction error by ETS allowance price



(b) Correlation coefficient between prediction error and ETS allowance price, by power unit

Figure E.1: Correlation between prediction error & ETS allowance price

order of their bid prices until residual demand is met. However, the dispatch and generation volumes implied by the merit order is not the same as actual unit activity, which is the outcome that we observe in the data and what our model is supposed to predict.

Transmission constraints are an important contributing factor to why realized unit-level generation differs from generation as implied by the merit order. The day-ahead market implicitly assumes a copperplate, that is a frictionless electricity grid in which any unit can satisfy any demand that arises anywhere in the zone. This assumption is unrealistic, in particular in large bidding zones. When there is a local demand peak, it may be that only geographically close units can be dispatched to satisfy that local demand peak due to within-zone transmission constraints. Similarly, when a bottleneck is about to occur somewhere in the grid, upstream power plants are instructed to reduce their output, while downstream plants must increase their output. This short-term redispatch occurs independently of the merit order and is therefore not captured by our merit-order-based predictors.

We do not observe within-zone transmission lines, and neither do we observe the distribution of demand within a zone. To approximate both, we provide the algorithm with information on a unit's geographical position (latitude and longitude). Transmission constraints are systematic physical properties of a bidding zone's grid that can be treated as approximately constant during our six-year observation window. By providing the Random Forest with information on a unit's geographic position, our aspiration is that the algorithm understands that units in a given area might run more often during peak hours than its marginal cost suggests because they are located downstream from a bottleneck. Similarly, we hope that local demand peaks appear with sufficient regularity such that the algorithm learns its generation patterns based on a unit's geographic position plus time dummies (hour-of-day, day-of-week, month-of-year) for when these local demand peaks tend to occur. In fact, as shown in Figure 8, latitude and longitude are the model's two most important predictors and substantially improve predictive power. These geographic features are especially critical for modeling dispatch in Spain, the bidding zone with

the second-largest geographic area in our sample. Spanish electricity demand is concentrated in a few distant metropolitan centers, creating a system prone to internal transmission constraints.

Because the model predicts realized generation, the model learns the merit order, but also deviations from it. This predictive trade-off is a central feature of our empirical strategy. The alternative would be to model the "pure" day-ahead merit order using bid data. However, since we eventually convert the predicted generation pattern into pollutant emissions, such an approach would fail to capture the total air pollution burden that is determined in part by the physical constraints in the grid. To quantify realized emissions, we must train the model on actual dispatch data to include the "out-of-merit" generation required for grid stability and reliability. By utilizing observed generation, the algorithm is able to internalize the complex interactions between market signals and physical constraints that determine the actual utilization rates of fossil units.

Furthermore, training on ex-post dispatch data allows the model to account for intraday market activity and real-time balancing actions. Because we provide the Random Forest with contemporaneous fuel prices, ETS allowance prices, and actual realized net load, the algorithm can learn short-term dispatch corrections that a theoretical day-ahead model would necessarily miss. This ensures that our counterfactual predictions reflect not only shifts in the primary merit order, but also the secondary operational adjustments essential for maintaining a continuous equilibrium.

E.6 Does nuclear energy respond to carbon prices?

We model only the fossil part of the merit curve. Thereby, we assume that all other electricity generation is *non-dispatchable*, i.e. that it always generates at its available capacity irrespective of market conditions. This is certainly true for renewable energy. Solar panels and wind turbines generate electricity and feed it into the grid at virtually zero short-run marginal cost whenever the sun shines or the wind blows. However, ruling out dispatch for other forms of non-fossil energy might be less realistic. This is particularly true for nuclear power in France, which accounts for a large share of total generation. Many French reactors were designed for load-following, using special control rods that allow rapid changes in output. Given the size of the French nuclear fleet, the marginal plant may frequently be a nuclear unit, which then has to ramp up or down flexibly. Figure E.2 shows average utilization rates of French nuclear power units over hours of the day, separately for weekdays and weekends, when they were reported to be fully available. Two things become apparent. First, nuclear power units almost always run at or near full capacity. On average, they operate at 96.2% capacity on weekdays and at 94.1% capacity on weekends. Second, these units seem to follow load to at least a small degree. Utilization on weekends is lower than on weekdays, and there are visible peaks during high-demand morning and evening hours, with valleys during low-demand midday and night hours.

To analyze whether nuclear generation responds to changes in both the carbon price as well as load, we construct a panel of French nuclear power units based on data from ENTSO-E. We regress the utilization rate (output relative to available capacity) on the carbon price, the natural gas price, the coal price, as well as on domestic (French) load that nuclear power

plants face (total load minus renewable generation[32]) and the loads net of renewables of all neighboring bidding zones. Because the utilization rate is bounded between zero and one and heavily concentrated near one, we estimate a fractional logit model by quasi-maximum likelihood (Papke and Wooldridge, 1996). We specify the conditional mean of the utilization rate as a logistic function of the covariates and estimate it via quasi-log-likelihood (Bernoulli). This estimator is consistent as long as the conditional mean is correctly specified, and does not require any assumptions about the full conditional distribution of the utilization rate (Papke and Wooldridge, 1996). We report standard errors robust to this misspecification, clustered at the monthly (column (1)) and weekly (columns (2)-(4)) level, respectively.

Table E.2 reports the regression results as average marginal effects (AMEs) on the utilization rate, in percentage points. All specifications include unit and year fixed effects, so that the effect of the carbon price is identified from variation within a year. Column (1) uses the unit $\times$ month panel with calendar-month fixed effects. A €10/t increase in the carbon price is associated with a 0.27 pp higher utilization rate, but the estimate is statistically insignificant. Column (2) uses data at the unit $\times$ day level, the frequency at which the (lagged) carbon price varies, and additionally includes weekday fixed effects. The point estimate falls to 0.12 pp and remains insignificant. Columns (3) and (4) use unit $\times$ hour data with hour-of-day $\times$ calendar-month and weekday fixed effects, and column (4) additionally interacts the carbon price with French residual load. The estimated effects are again small (0.12 and 0.20 pp) and statistically insignificant. Across all four specifications, the point estimates are not statistically significant and do not exceed 0.27 pp per €10/t. Hence, short-term fluctuations in the carbon price do not significantly affect the utilization of French nuclear capacity.

The coefficients on domestic residual load[33] indicate that French nuclear units respond to demand. A 1 GW increase in domestic residual load, whether measured as a monthly, daily or hourly average, raises the utilization rate by a statistically significant 0.18 to 0.23 pp. This effect is economically meaningful. After removing the variation absorbed by the fixed effects, the standard deviation of residual load is 2.8 GW (monthly), 5.2 GW (daily) and 5.6 GW (hourly), so that a one-standard-deviation increase in residual load raises utilization by about 0.7 to 1.0 pp. Without removing the variation absorbed by the fixed effects, the standard deviation of residual load is about 10 to 12 GW, which would correspond to roughly 1.9 to 2.2 pp. In fact, we expect nuclear units to be more likely to be marginal, and therefore to be ramped down, when residual load is low, i.e., when demand is low or renewable generation is high. The positive and highly significant coefficients on residual load are consistent with this: French nuclear units run closer to their available capacity when residual load is high.

In column (4), we additionally interact the carbon price with our continuous measure of residual load and present the implied effect of the carbon price averaged within each of the five quintiles of residual load. Intuitively, a higher carbon price should attenuate the relationship between residual load and nuclear output. When carbon prices are low, coal is competitive, so

[32] We subtract from total load these types of renewable energy: Wind, Solar, Geothermal, Marine, and Hydro Run-of-River.

[33] "Residual" here takes the perspective of nuclear power plants: residual load is total load net of generation from wind, solar, run-of-river hydro, marine and geothermal, i.e., net of only those units that always precede nuclear units in the merit order. It thus differs from residual demand in our main model, which additionally nets out nuclear generation.

that nuclear power plants are pushed to the margin more often. They would then respond more strongly to changes in residual load given that they are marginal more often. When carbon prices are high, in contrast, nuclear energy has a cost advantage over coal and should be marginal less often, thereby running more continuously with lower sensitivity to changes in residual load. However, the estimated AMEs of the interaction of carbon prices and residual load are not significantly different from each other at the five percent level across residual load quintiles, although the point estimates themselves slightly tend towards zero as we go up the quintiles. Yet, it remains that the carbon price does not significantly affect either nuclear output or how nuclear output responds to residual load. Lastly, we also estimate the hourly specification, but replace the continuous measure of residual load with explicit residual load quintiles. Table E.3 contains the results. Again, we find that the carbon price has no effect on nuclear activity, neither on average nor for specific intervals of residual load, across specifications.

Furthermore, French nuclear units also respond to residual load in neighboring bidding zones, consistent with France's role as a large net exporter of electricity. In the daily and hourly specifications (columns 2–4), a 1 GW increase in residual load in Belgium raises French nuclear utilization by about 0.22 to 0.29 pp, an effect of similar size as that of a 1 GW increase in French residual load. This is plausible given that Belgium is a small country whose electricity grid is closely integrated with France. Residual load in Germany, Britain and Spain also has positive and significant, but considerably smaller, effects of about 0.05 pp, 0.025 pp and 0.02 pp per GW, respectively. Since those countries are large and have substantial domestic dispatchable capacity, their residual load moves French nuclear generation less strongly, yet still significantly. Residual load in Italy has no significant effect. This is consistent with limited cross-border capacity and frequently congested interconnectors towards Italy, so that additional residual load in Italy does not affect French nuclear activity.

Overall, Tables E.2 and E.3 support our modeling approach. First, the activity of nuclear power plants responds strongly to residual demand (both domestic and foreign). Therefore, it is essential to include it as a predictor in the machine learning model and to keep it constant in the counterfactual scenarios. Second, the activity of nuclear power plants does not significantly respond to changes in the carbon price. Our point estimates do not go beyond an effect of 0.27 percentage points per €10/t and do not indicate statistical significance. Accordingly, this supports our approach to treat French nuclear output as constant across our counterfactual scenarios with differing carbon prices.

Table E.2: Carbon price and the utilization (0–100) of available French nuclear capacity: fractional logit, excluding Swiss residual load

	Monthly (reactor-month)	Daily (reactor-day)	Hourly (reactor-hour)	
	(1)	(2)	(3)	(4)
CO_2 price (per EUR 10/t), average effect	0.272	0.118	0.115	0.200
	(0.256)	(0.155)	(0.155)	(0.159)
CO_2 effect implied by the continuous CO_2 × residual load interaction, averaged over hours in each quintile of French residual load:				
Hours in quintile 1 (lowest)				0.245
				(0.276)
Hours in quintile 2				0.227
				(0.197)
Hours in quintile 3				0.213
				(0.159)
Hours in quintile 4				0.198
				(0.129)
Hours in quintile 5 (highest)				0.137*
				(0.078)
Gas price (per EUR 10/MWh)	−0.271*	−0.135**	−0.142**	−0.244***
	(0.145)	(0.066)	(0.067)	(0.087)
Coal price (per EUR 10/t)	0.036	0.093**	0.104***	0.149***
	(0.062)	(0.038)	(0.038)	(0.048)
Domestic residual load (GW)	0.232***	0.178***	0.187***	0.188***
	(0.041)	(0.016)	(0.013)	(0.013)
Residual load BE (GW)	1.077**	0.288***	0.226***	0.218***
	(0.475)	(0.075)	(0.048)	(0.047)
Residual load DE (GW)	−0.038	0.045***	0.047***	0.048***
	(0.031)	(0.007)	(0.006)	(0.006)
Residual load ES (GW)	0.041	0.037**	0.022*	0.022*
	(0.076)	(0.014)	(0.011)	(0.012)
Residual load GB (GW)	−0.106**	0.015	0.024***	0.026***
	(0.053)	(0.014)	(0.009)	(0.009)
Residual load IT (GW)	−0.266***	−0.036	−0.010	−0.011
	(0.094)	(0.034)	(0.023)	(0.023)
Foreign residual loads jointly zero (p-value)	0.006	<0.001	<0.001	<0.001
Holiday controls	Yes	Yes	Yes	Yes
Unit FE	Yes	Yes	Yes	Yes
Calendar-month FE	Yes	Yes		
Hour-of-day × calendar-month FE			Yes	Yes
Weekday FE		Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	3,346	86,935	2,053,853	2,053,853
Squared correlation	0.201	0.160	0.129	0.129

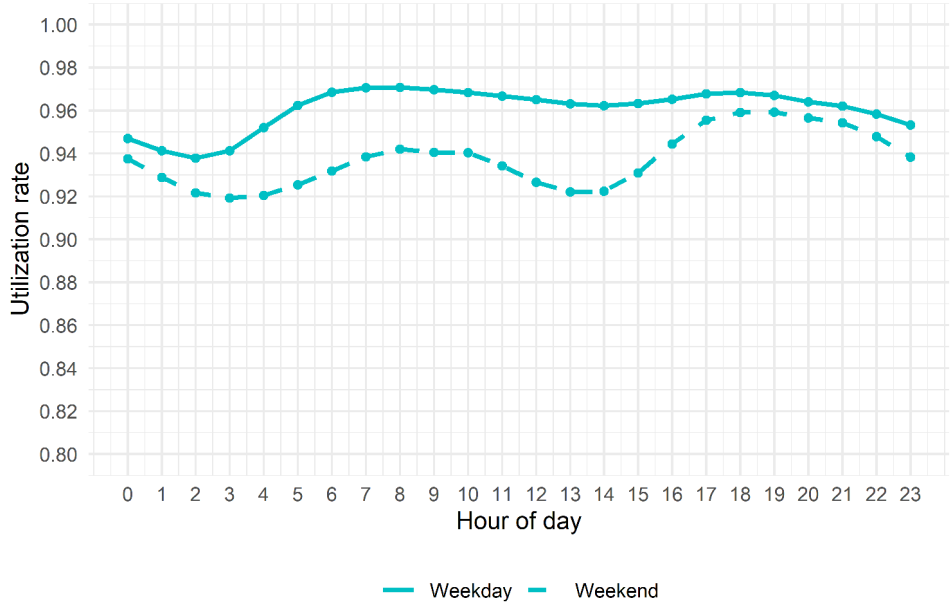
Notes: Same specifications and estimation samples as in Table E.2, but excluding Swiss residual load from the set of controls. Because run-of-river generation is largely missing from the ENTSO-E data for Switzerland, Swiss residual load effectively measures total Swiss load and is therefore not comparable to the residual loads of the other neighbouring zones. Fractional logit (quasi-binomial, logit link) estimated by quasi-maximum likelihood. Reported are average marginal effects (AMEs) in percentage points, weighted by available capacity; standard errors treat the average slope of the logistic function as fixed. The dependent variable is the utilization of available capacity (output divided by available capacity, with output capped at available capacity), 2016–2021. In columns (1) and (2), utilization is output over available energy in the remaining hours of each reactor-month and reactor-day, respectively, and regressors are monthly and daily means. Column (4) adds a single continuous interaction of the CO_2 price with French residual load; the rows by quintile are computed after estimation as described in the notes to Table E.2. The p-value refers to a Wald test of the joint null hypothesis that the coefficients on the residual loads of Belgium, Germany, Spain, Great Britain and Italy are zero. Prices are lagged by 24 hours. Holiday controls: number of weekday holidays and bridge days and share of Christmas-week hours (column 1); holiday, bridge-day and Christmas-week dummies (columns 2–4). Standard errors in parentheses, clustered by month of sample (column 1) and by week of sample (columns 2–4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.3: Carbon price and French nuclear utilization by residual-load quintile: robustness

	Pooled quintiles		Hour-specific quintiles
	(1)	(2)	(3)
CO_2 price (per EUR 10/t)	0.008		
	(0.160)		
CO_2 effect by French residual-load quintile:			
Quintile 1 (lowest)		0.064	0.105
		(0.291)	(0.271)
Quintile 2		0.041	0.015
		(0.197)	(0.191)
Quintile 3		0.092	0.126
		(0.167)	(0.168)
Quintile 4		0.167	0.188
		(0.128)	(0.127)
Quintile 5 (highest)		0.048	0.095
		(0.085)	(0.087)
Gas and coal prices	Yes	Yes	Yes
Residual load BE, CH, DE, ES, GB, IT	Yes	Yes	Yes
Holiday controls	Yes	Yes	Yes
Unit FE	Yes	Yes	Yes
French residual-load quintile FE	Yes	Yes	Yes
Hour-of-day × calendar-month FE	Yes	Yes	Yes
Weekday FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	2,053,853	2,053,853	2,053,853
Squared correlation	0.127	0.127	0.127

Notes: Fractional logit as in Table E.2, unit of observation reactor-hour. French residual load enters only through quintile fixed effects (level) and quintile-specific CO_2 coefficients (columns 2–3); there is no continuous French residual-load control. In columns (1)–(2), quintiles are defined on the pooled distribution of French hourly residual load over 2016–2021; in column (3), separately for each hour of the day, so that quintile 1 contains hours with unusually low residual load for that time of day. Reported are average marginal effects in percentage points per EUR 10/t; each quintile's effect uses the average slope of the logistic function within that quintile. Standard errors, clustered by week of sample, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure E.2: Utilization rate of French nuclear generation units



Notes: The graph shows the average utilization rate (output relative to capacity) of French nuclear power generating units between 2016 and 2021. Averages are computed by hour-of-the-day, separately for weekdays (Mo-Fr, bold) and weekends (Sa&Su, dashed), and excluding maintenance downtime (nuclear units are regularly unavailable while undergoing maintenance).