



CER-ETH – Center of Economic Research at ETH Zurich

How Effective Was the UK Carbon Tax?—A Machine Learning Approach to
Policy Evaluation

J. Abrell, M. Kosch, S. Rausch

Working Paper 19/317bis
June 2020

Economics Working Paper Series



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

How Effective Was the UK Carbon Tax?—A Machine Learning Approach to Policy Evaluation

By JAN ABRELL, MIRJAM KOSCH, AND SEBASTIAN RAUSCH*

While carbon taxes are generally seen as a rational policy response to climate change, knowledge about their performance from an ex-post perspective is still limited. This paper analyzes the emissions and cost impacts of the UK CPS, a carbon tax levied on all fossil-fired power plants. To overcome the problem of a missing control group, we propose a policy evaluation approach which leverages economic theory and machine learning for counterfactual prediction. Our results indicate that in the period 2013-2016 the CPS lowered emissions by 6.2 percent at an average cost of €18 per ton. We find substantial temporal heterogeneity in tax-induced impacts which stems from variation in relative fuel prices. An important implication for climate policy is that in the short run a higher carbon tax does not necessarily lead to higher emissions reductions or higher costs. JEL Codes: C54, Q48, Q52, Q58, L94

To avoid dangerous and costly climate change, the disposal space for carbon dioxide (CO₂) in the atmosphere is “scarce” and will soon be exhausted (McGlade and Ekins, 2015; IPCC, 2018). In tackling this major 21st-century challenge, and based on an elementary understanding of how today’s market-oriented systems organize economic activity based on scarce resources, economists have long been advocating for carbon pricing as an effective and efficient policy response (Nordhaus, 1994; Goulder and Parry, 2008; Metcalf, 2009). About one quarter of global CO₂ emissions are currently regulated under some form of carbon pricing (World Bank, 2018). While a plethora of studies offers ex-ante assessments of carbon pricing using theoretical and quantitative simulation-based work,¹ there is much less research on the the ex-post effects of carbon pricing. This, however, is pivotal for designing effective and efficient climate policies.

This paper contributes by providing an ex-post evaluation of a real-world policy experiment of carbon pricing: the UK carbon tax, also known as the *Carbon Price Support* (CPS). The CPS was introduced to enhance economic incentives for

* Abrell: Centre for Energy Policy and Economics ETH Zurich and Center for Energy and Environment of ZHAW Winterthur, Zuerichbergstrasse 18, CH-8032 Zurich, Switzerland, phone: +41-44-632 06 52 (email: jabrell@ethz.ch). Kosch: Center for Energy and Environment of ZHAW Winterthur and Centre for Energy Policy and Economics ETH Zurich, Switzerland (email: mirjam.kosch@zhaw.ch). Rausch: Department of Management, Technology and Economics, ETH Zurich, Switzerland, Center for Economic Research at ETH (CER-ETH), Centre for Energy Policy and Economics ETH Zurich, and Massachusetts Institute of Technology, Joint Program on the Science and Policy of Global Change, Cambridge, USA (email: srausch@ethz.ch).

¹See, for example, Tavoni et al. (2014), Golosov et al. (2014), Liski and Gerlagh (2016), Goulder, Hafstead and Williams III (2016), Bretschger et al. (2017), and a series of papers from multi-model comparison studies carried out under the framework of the Stanford Energy Modeling Forum for the U.S. (Fawcett et al., 2014) and Europe (Weyant et al., 2013).

carbon abatement in the heavily fossil-based UK electricity sector. As the CPS affects the output and operating decisions of all fossil-fueled generation facilities in the UK electricity market, the challenge arises that no suitable control group or counterfactual exists against which the impact on treated units can be evaluated. In order to estimate the causal effects of the CPS policy intervention, it is thus not possible to use standard program evaluation methods based on comparing treated and untreated units—such as e.g. difference-in-differences (DiD) methods (Angrist and Pischke, 2008; Athey and Imbens, 2017). To overcome this problem, we develop and implement a new approach which combines economic theory and machine learning (ML) techniques to estimate the treatment effect of a policy intervention in settings with observational, high-frequency data when no control group exists. We apply our approach to analyze the environmental effectiveness and costs of the UK carbon tax. To our knowledge, this is the first paper in economics to incorporate ML methods to estimate the ex-post effects of carbon pricing.

Our proposed approach leverages economic theory on price and production decisions in electricity markets and ML techniques to estimate the treatment effect of a policy intervention when no control group exists. The main idea is to train a model, based on an equilibrium model of the wholesale electricity market and ML, which predicts outcomes under observed and counterfactual treatment, and derives a treatment effect based on these predictions. First, to train the model, we use data before the introduction of the carbon tax and after its introduction. As the UK carbon tax is adjusted annually, it provides insufficient variation in terms of tax rate changes. To overcome this, we exploit the variation in relative market prices for coal and natural gas, which influence input costs and thus power plant output decisions through the same channel as the carbon tax itself. We use high frequency (hourly) panel data of power generation at the plant level together with market information on hourly demand, production capacities by power plant, fuel and carbon prices, and factors affecting the thermal efficiency of power plants, such as temperature. Second, we use the prediction model, which has been “machine-learned” using data under observed treatment (i.e., with the carbon tax), to predict outcomes under unobserved, counterfactual intervention (i.e., no carbon tax). Third, the estimator of the treatment effect is based on the difference of these predictions. It accounts for the influence of observed and unobserved variables, and systematic prediction errors. We show that under the assumption that prediction errors are independent of the treatment variable, the difference between two predictions yields an unbiased estimate of the impact on power-plant output caused by the carbon tax.

An important advantage of our approach is its ability to explicitly represent the channels through which the policy intervention affects the outcome variable. As our proxy for treatment—the relative fuel prices—is already observed before the CPS policy is introduced, we can use observations from both the pre- and post-treatment period to train the model. This improves the basis for learning about the key mechanisms between input prices and output through which the policy intervention affects power plant output decisions. In addition, the application

of ML techniques enables the development of nonparametric predictors and thus the nonparametric identification of treatment effects. Beyond the estimation of total treatment effects, we can perform simulations with the ML-trained model to assess the impacts of different (hypothetical) treatment intensities—a feature we exploit to analyze the empirical determinants of the environmental effectiveness and abatement costs of a carbon tax policy.

Our ex-post evaluation of the UK carbon tax policy yields the following main insights. First, our analysis provides strong evidence that a carbon tax is an effective regulatory instrument to reduce CO₂ emissions: the CPS induced a substitution away from “dirty” coal to cleaner natural gas-fired power plants—replacing about 15 percent or 46 TWh of coal-based generation and reducing emissions by 6.2 percent between 2013 and 2016. Second, the abatement of one ton of CO₂ has brought about on average additional costs of €18.2 in total for consumers and fossil-based electricity producers. Third, there is substantial heterogeneity in the carbon tax-induced abatement quantity and costs over time. Simulating the machine-learned model, we characterize the empirical conditions which influence the environmental effectiveness and costs of the tax policy. We find that the heterogeneity is mainly driven by the variation in the relative carbon tax-exclusive prices for coal and natural gas and only to a limited extent by the carbon tax rate itself. The important implication for climate policy is that a higher carbon tax does not necessarily deliver higher emissions reductions. At the same time, a higher carbon tax need not necessarily result in higher abatement costs.²

Our paper contributes to the literature in several important ways. First, we add to the recent and emerging literature on the use of ML techniques in economics and quantitative social science. Traditionally, ML methods have been used for pure prediction problems such as demand estimations (Bajari et al., 2015). More recently, ML methods have provided important new tools to improve the estimation of causal effects from observational data in high-dimensional settings as they enable to flexibly control for a large number of covariates (for overview articles see, for example, Varian, 2014; Athey, 2017; Athey and Imbens, 2017; Mullainathan and Spiess, 2017). Burlig et al. (2019) is a recent example using ML algorithms to estimate causal effects.³ Estimating the impact of the German nuclear phase-out on production levels of other plants in the market, Jarvis, Deschenes and Jha (2020)

²A by-product of our ex-post evaluation of the UK CPS is the derivation of empirical marginal abatement cost (MAC) curves for the UK electricity sector, i.e. relationships between tons of emissions abated and the CO₂ price. MACs have been widely used as reduced-form tools to inform policy-making and to illustrate simple economic concepts such as the benefits of emissions trading (Ellerman and Decaux, 1998; Klepper and Peterson, 2006; Morris, Paltsev and Reilly, 2012).

³Our approach differs in two important ways. First, Burlig et al. (2019) use the pre-treatment period to train a model predicting the post-treatment outcome without the intervention. In contrast, in our setting treatment is continuous and the data generating process is invariant to treatment. We thus train the model on the full sample, but at the same time have to rely on the continuity of treatment or, alternatively, have to identify a (continuous) variable with the same causal impact as the treatment variable. Second, ML based predictions have to deal with prediction errors. Burlig et al. (2019) assume that prediction errors have similar trends across treatment and control groups. They employ a DiD estimator to eliminate biases caused by prediction errors. In contrast, we eliminate this bias comparing predicted values of observed and counterfactual values, i.e. we assume that prediction errors are independent of treatment levels. Importantly, this allows us to estimate the impact of treatment in the absence of a control group.

use ML methods to predict generation for years without observational data as well as without the phase-out. [Varian \(2016\)](#) mentions the possibility of estimating treatment effects by constructing the unobserved counterfactual when no control group is available. To the best of our knowledge this paper provides the first empirical implementation of this idea in economics.

Second, there exists only a handful of studies using econometric and program evaluation methods to quantify the environmental impacts of carbon pricing, be it through a tax- or quantity-based approach to regulation. An overview of the work focusing on the impacts of the EU ETS is provided by [Martin, Muûls and Wagner \(2016\)](#). The paper by [McGuinness and Ellerman \(2008\)](#) estimate the impact of permit prices on the output of power plants in the UK. Using a panel regression, they quantify the emissions offset in the British power sector for the pilot trading period of the EU ETS. [Martin, De Preux and Wagner \(2014\)](#) analyze the impacts of the *Climate Change Levy* on manufacturing plants in the UK. Using panel data on manufacturing plants in the UK, their identification strategy builds on the comparison of outcomes between plants subject to the full tax and plants paying only 20 percent of the tax. [Bretschger and Grieg \(2020\)](#) estimate the impact of a fuel tax in the UK on CO₂ emissions from traffic.⁴ Recently, a few studies have been looking at the impacts of the UK CPS itself: [Gissey et al. \(2019\)](#) analyze the impact of the UK CPS on wholesale electricity trading, [Chyong, Guo and Newbery \(2020\)](#) analyze the merit order effect of the CPS and its impact on carbon abatement induced by wind generation, [Leroutier \(2019\)](#) and [Gugler, Haxhimusa and Liebensteiner \(2020\)](#) estimate the impact of the UK carbon price floor on CO₂ emissions. The latter two are closely related to our research question, yet they differ in methods. [Leroutier \(2019\)](#) uses a synthetic control group method which relies on constructing a “no-policy” counterfactual UK power sector from a combination of other European countries, and [Gugler, Haxhimusa and Liebensteiner \(2020\)](#) use regression discontinuities in time. Both studies find a significantly larger impact of the CPS on emissions (26 % in the case of [Gugler, Haxhimusa and Liebensteiner \(2020\)](#) and 49 % in the case of [Leroutier \(2019\)](#)). One main explanation for this difference is that we concentrate on estimating the short-run impact of the CPS on fuel switching. That is, we assume that all observed changes in capacities have not been induced by the CPS but potentially by other policies such as the “Large Combustion Plant Directive”. Given the possibility that the CPS had an impact on plant closure, our results tend to be a lower bound estimate. With this paper, we contribute to the scarce empirical evidence on the economic impacts of carbon taxes by applying an estimation strategy which can be used in a setting without a control group.

Third, a recent and growing literature, following the U.S. shale gas boom after 2005, uses the variation in natural gas prices to empirically estimate the impact of fuel prices on CO₂ and other pollutants stemming from electricity generation

⁴[Fowle, Holland and Mansur \(2012\)](#) evaluate the NO_x emissions reduction delivered by the Southern California’s emission trading program. To construct the counterfactual, they exploit program-specific participation requirements allowing them to match regulated facilities with similar facilities in non-attainment areas.

(see, for example, Knittel, Metaxoglou and Trindade, 2015; Linn, Muehlenbachs and Wang, 2014; Holladay, Soloway et al., 2016; Holladay and LaRiviere, 2017). Cullen and Mansur (2017) and Lu, Salovaara and McElroy (2012) exploit the fact that the introduction of a carbon price impacts emissions through the same economic mechanism as a change in relative fuel prices. Similar to our approach, these studies use the variation in natural gas prices to estimate the impact of a *hypothetical* carbon pricing policy on emissions. We contribute with an ex-post assessment of a real-world carbon tax policy.

Fourth, studies investigating the environmental impact of carbon pricing in the electricity sector are abundant but the vast majority of the work relies on numerical simulation methods based on strong theory-driven behavioral assumptions and, sometimes, insufficiently validated empirical hypotheses (see, for example, Delarue, Ellerman and D’Haeseleer, 2010b; Delarue, Voorspools and D’Haeseleer, 2008; Rausch and Mowers, 2014; Goulder, Hafstead and Williams III, 2016; Abrell and Rausch, 2016). Some of the economic mechanisms at work, which we empirically identify in our analysis, have already been analyzed using ex-ante policy analysis based on analytical and simulation models. For example, Kirat and Ahamada (2011) show that the high permit prices induced a switch in the merit order from coal to gas. Delarue, Ellerman and D’Haeseleer (2010a) show that abatement does not only depend on the level of carbon prices but also on demand and the ratio between coal and gas prices. Some studies model the fuel switching potential for hypothetical carbon pricing policies as in Pettersson, Söderholm and Lundmark (2012) for the EU ETS and Chevallier et al. (2012) for the UK.

The remainder of this paper is organized as follows. Section I presents our empirical strategy to estimate the treatment effect of a policy intervention in the absence of a control group. Section II details how we apply the framework to assess the CO₂ abatement quantity and costs of the UK carbon tax, including a description of data sources. Section III scrutinizes the validity of our approach for estimating the causal effects of the policy intervention. Section IV presents our main findings. Section V analyzes the determinants of environmental effectiveness and costs of the UK carbon tax. Section VI concludes.

I. Estimation Strategy: Conceptual Framework

The primary challenge in assessing the emissions impact of the UK carbon tax on the electricity market is to estimate the treatment effect of a policy that affects all entities in the market, i.e. there exists no control group (or all units are assigned to treatment with probability one). Our proposed approach leverages economic theory on price and production decisions in electricity markets and ML techniques to estimate the treatment effect of a policy intervention. The main idea is to train a model, based on an economic equilibrium model and ML, which predicts outcomes under observed and counterfactual treatment. The estimator of the treatment effect is based on the difference of these predictions and accounts for the influence of observed and unobserved variables and prediction errors.




FIGURE 1. Prediction-based estimator of treatment effect

Notes: y_{it} denotes outcome of unit i at time t (i.e. electricity output of a given power plant) which has been subject to treatment (i.e., a carbon tax). z_t and \bar{z}_t refer to observed and counterfactual levels of treatment which apply to all units equally, respectively. Hats refer to predicted values. δ_{it} and $\hat{\delta}_{it}$ denote the true and estimated treatment effect, respectively.

A. Basic Idea

Consider a population model according to which the outcome y_{it} of unit i at time t —here, the output of power plant i —is generated according to

$$(1) \quad y_{it} = f_i(x_{it}, h_{it}, z_t) + \epsilon_{it}.$$

z_t is the time-varying treatment which is received uniformly by all units—here, a carbon tax levied uniformly on all power plants in the market at time t . x_{it} and h_{it} are vectors of observed and unobserved control variables, respectively. $\epsilon_{it} \sim (0, \sigma_\epsilon^2)$ is a random noise with mean zero $\mathbb{E}[\epsilon_{it}] = 0$ and variance σ_ϵ^2 . ϵ_{it} is independent of controls and treatment: $\epsilon_{it} \perp (x_{it}, h_{it}, z_t), \forall i, t$. We assume that the function f_i is invariant or autonomous to changes in the treatment and control variables (Haavelmo, 1944; Aldrich, 1989; Peters, Bühlmann and Meinshausen, 2016).

The goal is to identify the effect on outcome y_{it} caused by a change in the treatment from its observed values z_t to counterfactual values \bar{z}_t (i.e., power plant output which would have been chosen without a carbon tax). The treatment effect is defined as:

$$(2) \quad \delta_{it} := y_{it}^{z_t} - y_{it}^{\bar{z}_t} \quad \forall i, t.$$

Figure 1 graphically illustrates the treatment effect where solid and dashed lines denote observed and unobserved outcomes, respectively. Of course, the funda-

mental problem—also known as the missing data problem (Rubin, 1974; Holland, 1986)—is that we do not observe $y_{it}^{z_t}$ and $y_{it}^{\bar{z}_t}$ at the same time.

To obtain an estimate of the treatment effect δ , the basic idea of our approach is to derive a predictor \hat{f}_i for the function f_i that produces reliable predictions. To this end, we use economic theory to provide a basis for specifying the predictor function. Specifically, we use a dispatch and peak-load pricing model of the wholesale electricity market (Boiteux, 1960) to impose structure on the underlying data generating process and to construct a population model which “pre-selects” the factors influencing the outcome. In a second step, we harness the power of ML methods which—in contrast to traditional econometric methods focused on consistently estimating parameters of f —are optimized to predict the value of the outcome variable (Mullainathan and Spiess, 2017).⁵ In applying ML methods to estimate the predictor function, economic theory is of ultimate importance to inform about the choice of control variables. An agnostic “let the data speak” approach, which would leave model selection entirely to the ML algorithm, is highly problematic.

While making use of economic theory in a first step adds structure to the prediction problem, obtaining an unbiased estimate for δ requires controlling for the impact of unobserved variables and systematic prediction biases. We thus estimate the treatment effect as (see also Figure 1 where dotted lines represent predicted values):

$$(3) \quad \hat{\delta}_{it} := \hat{y}_{it}^{z_t} - \hat{y}_{it}^{\bar{z}_t} \quad \forall i, t.$$

\hat{y}_{it} denotes predicted values of the outcome variable under a situation with or without treatment which, naturally, involve prediction errors ξ_{it} :

$$(4) \quad \hat{y}_{it} = \hat{f}_i(x_{it}, z_t) = y_{it} + \xi_{it}(x_{it}, h_{it}, z_t) + \epsilon_{it}.$$

Our main identifying assumption to obtain $\hat{\delta}_{it}$ is that the prediction error only depends on observed and unobserved variables, but does not change between the prediction of outcomes under observed and counterfactual treatments keeping control variables constant. That is:

ASSUMPTION 1: *Prediction errors $\xi_{it}(x_{it}, h_{it}, z_t)$ are independent of the treatment: $\xi_{it}(x_{it}, h_{it}, z_t) = \xi_{it}(x_{it}, h_{it}, \bar{z}_t) = \xi_{it}(x_{it}, h_{it})$, $\forall z_t$.*

Using *predicted* outcomes under observed and counterfactual treatment differences out prediction errors then yields an unbiased estimate of δ_{it} :

$$(5) \quad \begin{aligned} \hat{\delta}_{it} &= y_{it}^{z_t} + \xi_{it}(x_{it}, h_{it}) + \epsilon_{it}^{z_t} - \left[y_{it}^{\bar{z}_t} + \xi_{it}(x_{it}, h_{it}) + \epsilon_{it}^{\bar{z}_t} \right] \\ &= y_{it}^{z_t} - y_{it}^{\bar{z}_t} + \phi_{it}, \end{aligned}$$

⁵Using cross-validation, i.e. repeated re-sampling methods, ML methods construct an estimate of expected prediction errors that are minimized by regularizing the underlying function \hat{f}_i (see, for example, Hastie, Tibshirani and Friedman, 2008; Gareth et al., 2013, as well as Appendix B for a more detailed explanation of the ML approach).

where $\phi_{it} := \epsilon_{it}^{z_t} - \bar{\epsilon}_{it}^{z_t}$ is random noise with mean zero. The estimator is thus unbiased: $\mathbb{E}(\hat{\delta}_{it}) = \delta_{it}$.

It is important to understand that the estimator in (5) involves using *predicted* outcomes under observed and counterfactual treatment. To see this, consider a naïve approach which would alternatively estimate the treatment effect as the difference between the *observed* outcome under treatment ($y_{it}^{z_t}$) and the *predicted* outcome under no treatment ($\hat{y}_{it}^{\bar{z}_t}$). Such an approach would yield a biased estimate of δ due to the prediction error.

Besides removing prediction errors, the estimator in (5) also controls for the indirect impacts of observed and unobserved variables on outcome if the following assumptions hold true:

ASSUMPTION 2: *Observed controls are independent of the changes in the treatment variable: $x_{it} \perp\!\!\!\perp z_t$.*

ASSUMPTION 3: *Unobserved controls are conditionally independent to changes in the treatment variable given the observed controls: $h_{it} \perp\!\!\!\perp z_t | x_{it}$.*

Assumption 2 rules out effects of the treatment variable on observed controls. This assumption is necessary as the observed controls are held constant in the counterfactual simulation. Otherwise, if z influences x , there would be an indirect effect on the outcome, which would bias $\hat{\delta}$. Likewise, Assumption 3 rules out effects of the treatment variable on unobserved variables once observed variables are controlled for. Assumptions 1, 2 and 3 are essential for identification and can be viewed as analogous to the parallel trend assumption in a difference-in-differences setting.⁶

B. Additional Assumptions

Assumptions 1-3 are imposed on the population model for identification. For the feasibility of our estimation strategy, we need additional properties in the data that must be met.

To be able to identify the impact of treatment on outcome, there has to be sufficient variation in the treatment variable:

ASSUMPTION 4: *The variation in the level of treatment and controls over time is sufficiently large.*

Typically, the level of treatment—in our case the carbon tax—does not change frequently (if at all). One can, however, find a proxy variable which affects the outcome through the same mechanism as the treatment variable. A carbon tax is an input tax that directly affects input costs (and thus the output decisions of power plants) in the same way as input prices. It is thus possible to use the

⁶In fact, Assumption 1 is the key difference in comparison to the approach used by Burlig et al. (2019). They "...require treated and untreated schools to be trending similarly in prediction errors...". (Burlig et al., 2019, p. 15). In contrast, and given that we do not observe a control group, we need to assume that the prediction error is independent of the treatment in order to difference out the impact of unobservables and systematic prediction errors.

variation in carbon-based fuel prices, such as the prices of coal and natural gas, as a proxy for the insufficient variation in the carbon tax rate.

A potential concern for the validity of Assumption 1 is the quality of predictions based on *unobserved* counterfactual control and treatment values. Although ML algorithms are designed to produce reliable predictions, they only locally approximate the true model within the range of observed treatments and controls. It is thus unclear how the estimated functions behave for covariate and treatment combinations which lie outside of the range of observed combinations. The following *positivity* or *covariate overlap* assumption (Samii, Paler and Daly, 2016) rules out such cases:

ASSUMPTION 5: *Each combination of the counterfactual treatment z_t and covariates x_{it} has been observed. (i.e., $\text{Prob}[z|x_{it}] > 0$).*

While it is highly unlikely that all combinations of z and x were observed, Assumption 5 only requires that these combinations should lie within the range of the observed data.

II. Applying the Framework: The Case of the UK Carbon Tax

We apply the empirical strategy presented in Section I to assess the market impacts of the UK carbon tax. Sections II.A and II.B provide contextual detail about the carbon tax policy and draw on economic theory to derive the population model. Section II.C presents the data. Section II.D describes how we implement the conceptual idea behind our empirical strategy in the given context. Section III scrutinizes the validity of our identifying assumptions.

A. Policy Context and Confounding Factors

The main policy instrument of the UK government to decarbonize the heavily fossil-based UK electricity sector is the *Carbon Price Support* (CPS), an annual constant tax on fossil fuel use in the wholesale electricity market (Department of Energy & Climate Change, 2016). The CPS intends to close the gap between an envisaged minimum carbon price, the so-called *Carbon Price Floor* (CPF) and the price of European Emission Allowances (EUA) traded under the European Emissions Trading System (ETS).⁷ Table 1 shows the evolution of the EUA, CPS, and the total carbon price over time. Since the introduction of the CPS in 2013, the CPF always exceeded the EUA price, thus resulting in a positive CPS. In 2013, the modest level of the CPS led to a more than two-fold increase of the total carbon price for the UK electricity industry. In 2016, the CPS was set at the level of €21.60, six times higher than the annual EUA price in this year.

⁷Prior to the introduction of the CPS, the CPS level was conceptualized to be determined two years in advance as the difference between the EUA future price and the CPF. In 2013, the CPF was announced to increase up to 34.5 (69) €/tCO₂ in 2020 (2030). At the end of 2015, however, the UK government fixed the CPS rate to 21.6 €/tCO₂ until 2021 (Hirst, 2017). In the 2017 budget, the UK government expressed its confidence that “the Total Carbon Price, currently created by the combination of the EU Emission Trading System and the Carbon Price Support, is set at the right level [...]” (HM Treasury, 2017, Article 3.46), thus indicating that the CPS is likely to stay at its current level in future years.

TABLE 1. Descriptive statistics of UK electricity market: carbon prices, generation and import capacity, fuel prices, output, and demand.

	Year							
	2009	2010	2011	2012	2013	2014	2015	2016
<i>Carbon prices [€ per ton of CO₂]^a</i>								
EUA	13.23	14.36	13.02	7.37	4.76	6.22	7.34	5.26
CPS	—	—	—	—	5.85	12.17	24.70	21.60
Total carbon price (=EUA+CPS)	13.23	14.36	13.02	7.37	10.61	18.39	32.04	26.86
<i>Installed capacities [GW]</i>								
Coal	25.3	25.3	25.3	24.5	19.9	18.8	19	13.8
Gas	27.3	29.5	30.2	30.3	29.3	27.4	26.6	26.1
Import	2.5	2.5	3.5	3.6	4.0	4.0	4.0	4.0
<i>Fuel prices [€ per MWh thermal energy]</i>								
Coal	7.60 (0.74)	10.46 (1.55)	13.20 (0.45)	10.90 (0.68)	9.28 (0.54)	8.55 (0.35)	7.70 (0.56)	8.12 (2.27)
Gas	11.82 (4.47)	16.84 (3.53)	22.17 (1.31)	25.07 (2.01)	27.34 (2.79)	21.16 (3.29)	20.03 (2.19)	14.38 (2.53)
Ratio ^b	0.89 (0.19)	0.79 (0.07)	0.71 (0.05)	0.51 (0.06)	0.43 (0.04)	0.59 (0.09)	0.69 (0.08)	0.88 (0.08)
<i>Hourly demand and generation [GWh]</i>								
Demand	27.10 (6.51)	28.33 (6.58)	25.81 (6.63)	24.99 (6.77)	23.77 (6.93)	22.16 (6.23)	20.01 (6.36)	19.54 (6.43)
Gas generation	17.14 (3.01)	18.29 (3.07)	14.56 (3.79)	9.50 (4.16)	9.17 (5.12)	9.81 (4.87)	9.47 (4.43)	14.23 (4.75)
Coal generation	9.81 (5.80)	9.97 (5.29)	10.70 (5.14)	14.35 (4.04)	13.11 (3.18)	10.13 (4.10)	8.17 (3.45)	3.27 (2.88)

Notes: Standard deviations in parentheses. CPS taken from [Hirst \(2017\)](#) and [HM Revenue & Customs \(2014\)](#) converted with exchange rate data from [ECB \(2017\)](#). Daily European Emission Allowances (EUA) spot prices taken from [EEX \(2017\)](#). Further detail about data sources and calculations is provided in Section II.D. ^aAs the CPS is adjusted in April of every year, the annual EUA and CPS carbon prices for the years 2013-2016 are calculated based on the period from April to March of the subsequent year. ^bCoal-to-gas fuel price ratio, inclusive of EUA and CPS carbon prices, calculated according to equation (12).

To develop some first intuition for the impacts of the CPS on electricity supply and emissions, Figure 2 plots the short-run supply curve (i.e., ordering marginal cost of fossil-based power plants from low to high) for two situations:⁸ a hypothetical situation without the CPS where marginal emissions are only priced at the costs of an EUA (Panel a) and the observed situation with the CPS (Panel b). We observe two main changes. First, the supply curve shifts upward—indicating the increase in the marginal cost of all fossil plants. Second, as natural gas-fired power plants are less carbon-intensive, they are less affected by the carbon price increase and, therefore, become relatively cheaper. Gas plants are thus dispatched into the market and replace emissions-intensive coal-fired plants, in turn reducing emissions.

Consistent with this basic mechanism, Panel (a) in Figure 3 shows that starting with the introduction of the CPS in 2013 the annual market share of coal-fired

⁸The illustrative calculation shown in Figure 2 is based on one particular hour and assuming average heat efficiencies for plants; it ignores the fact that heat efficiencies, and hence the impact of CPS on individual plants, varies over time depending on temperature and other factors.




FIGURE 2. Illustrative impact of the UK carbon tax on the short-run market supply curve for electricity

Notes: The graph shows the merit order curve of fossil-based power plants on December 19, 2016, at 5:00 p.m. based on the data described in Section II.C. Hydro, nuclear, and renewable power plants are omitted and their total generation is subtracted from demand as they are always dispatched first given that their marginal cost are smaller than those of fossil-based plants. Marginal costs are calculated according to equation (16).



generation sharply decreased while the share of gas-fired plants increased; over the same period, UK’s electricity-sector emissions sharply declined.

While Figures 2 and 3 provide some first evidence that the CPS may have led to a reduction in electricity-sector CO₂ emissions, there is arguably a host of other factors which are likely to have affected the observed market outcomes. First, the fraction of electricity demand to be covered with domestic fossil-based generation from coal and natural gas has declined between 2013-2016. This is due to, at least, three factors: (i) energy efficiency improvements; (ii) targeted support policies have likely pushed in zero (or low) marginal-cost generation from renewable energy whenever the underlying natural resource (wind or solar) was available; and (iii) UK’s electricity imports have slightly increased likely due to both an expansion of newly built inter-connector lines (see Table 1) and the fact that the CPS has increased the domestic cost of generation relative to import prices.

Second, the switch from coal to natural gas was likely also triggered by substantial changes in relative fuel price. Between 2013-2016, natural gas prices declined by nearly 50 percent while coal prices remained largely constant (see Figure 3 Panel (b) and Table 1). This suggests that even without the introduction of the CPS there may have been a marked shift towards gas-fired generation in the UK electricity market.

Third, the decisions to shut down coal-fired plants, reflected in the available production capacity for coal (see Table 1), are likely influenced by factors which are unrelated to the CPS. A main reason for these closures is the European “Large Combustion Plant Directive”, which sets specific limits on local pollutant emissions for power plants constructed after the year 1987. It left electricity firms essentially the choice to either comply with the emissions limits or to “opt out” in which case a maximum operation time of 20’000 hours was granted until the end of 2015 when eventually the plant had to be shut down (European Commission, 2001).

In summary, there is ample evidence that the decline in coal generation and CO₂ emissions in the UK power sector which has occurred since the introduction of the

(a) Annual electricity generation by technology and CO₂ emissions over time

(b) Monthly coal and gas prices with and without CPS

FIGURE 3. Generation, emissions and fuel prices

Notes: Own calculations. Electricity generation by fuel is based on [ELEXON \(2016\)](#). “*Base*” comprises electricity generated from hydro and nuclear power plants. “*Renewables*” comprises wind, solar, and other (mainly biomass) generation where generation from wind and solar is corrected for generation embedded in final demand ([Nationalgrid, 2016](#)). “*Emissions*” refer to reported values from the EU Transaction Log ([European Commission, 2016](#)). Fuel prices for coal and natural gas are taken from [EIKON \(2017\)](#). CPS rates are reported by [Hirst \(2017\)](#) and [HM Revenue & Customs \(2014\)](#), and the EUA price by [EEX \(2017\)](#). Carbon price inclusive fuel prices refer to MWh of *thermal energy*.

CPS in 2013 has likely been the result of a multitude of factors comprising market developments (international fuel prices and electricity demand) and a variety of different policy measures (renewable energy support policies, transmission infrastructure measures, and the CPS). We next present our empirical framework we use to disentangle the market impacts brought about by the carbon tax policy alone.

B. Determinants of Wholesale Electricity Market Activity

We apply microeconomic theory based on a dispatch and peak-load pricing model of the wholesale electricity market (Boiteux, 1960) to pre-select the potentially relevant variables determining wholesale market impacts in response to a carbon tax. The pre-selected variables subsequently enter the ML algorithm to estimate the empirical prediction model which is used to estimate the treatment effect of the UK CPS.

COMPETITION IN UK'S WHOLESALE ELECTRICITY MARKET.—The UK wholesale electricity market is a liberalized market based on exchange and over-the-counter trades. In power exchanges, market participants can trade forward and real-time contracts.⁹ In the day-ahead market, market participants trade electricity for each hour of the next day. Given the new information in the market, these trades can be revised using the intra-day market which closes one hour before delivery time. In 2014 the UK regulator asked for an investigation of anti-competitive behavior in the UK energy market. In its final report, the “*Competition and Markets Authority*” (CMA, 2016) did not find evidence for anti-competitive behavior in the wholesale electricity market.

A SHORT-RUN EQUILIBRIUM MODEL OF WHOLESALE MARKET ACTIVITY.—We conceptualize the UK wholesale electricity market as being composed of firms which are assumed to operate under perfect competition maximizing profits using production quantities as the decision variable. Generation units of a firm are represented at the plant level where total production of plant $i \in I$ in hour $t \in T$ is denoted by X_{it} . The set I comprises thermal carbon-based generation plants (i.e., hard coal, lignite coal, natural gas) and other conventional plants (i.e., nuclear, hydro, pump storage, biomass). Generation from wind and solar is modeled exogenously. Production at any point in time cannot exceed the given effective production capacity K_{it} :

$$(6) \quad K_{it} \geq X_{it} \quad \perp \quad \mu_{it} \geq 0 \quad \forall i, t$$

where the time-dependency of capacity mainly reflects maintenance and unscheduled plant outages. μ_{it} is the shadow price of capacity for technology i at time t . The value of capacity in a given hour is zero ($\mu_{it} = 0$) if production is below the capacity limit; it is positive ($\mu_{it} > 0$) if the capacity constraint is binding.¹⁰

⁹Real-time trading of UK electricity mainly takes place in the EPEX-Spot and Nordpool power exchanges. Forward contracts are traded via the InterContinental Exchange (ICE) and NASDAQ.

¹⁰We use the “ \perp ” operator to indicate complementarity between equilibrium conditions and variables. A characteristic of economic equilibrium models is that they can be cast as a complementarity problem, i.e. given a function $F: \mathbb{R}^n \rightarrow \mathbb{R}^n$, find $z \in \mathbb{R}^n$ such that $F(z) \geq 0$, $z \geq 0$, and $z^T F(z) = 0$, or, in short-hand notation, $F(z) \geq 0 \perp z \geq 0$ (Mathiesen, 1985; Rutherford, 1995).

Marginal cost $c_{it}(\boldsymbol{\vartheta}_{it})$ of a generation unit at time t depend on exogenous factors

$$\boldsymbol{\vartheta}_{it} = \{p_t^f, \theta^f, \eta_{it}, p_t^{EUA}, p_t^{CPS}\}$$

comprising the time-dependent price of the fuel f used for electricity generation (p_t^f), the carbon content (θ^f), the time-varying EUA and CPS prices on CO₂ emissions (p_t^{EUA} and p_t^{CPS}), and time-specific heat efficiency (η_{it}) reflecting the influence of ambient temperature ($temp_t$) and potential efficiency losses due to part-load operation.

In equilibrium, the following zero-profit condition, relating unit costs (comprising marginal costs and the opportunity costs for capacity) to unit revenues determines the output of generation unit i , y_{it} :

$$(7) \quad c_{it}(\boldsymbol{\vartheta}_{it}) + \mu_{it} \geq P_t \quad \perp \quad y_{it} \geq 0 \quad \forall i, t$$

where P_t measures unit profits or the wholesale electricity price at time t .¹¹ If unit cost exceed unit profit, positive generation would lead to losses and thus $y_{it} = 0$. Given perfect competition and no barriers for market entry or exit, zero profits in equilibrium (i.e., unit cost equal to unit profit) determine a positive level of electricity supply $y_{it} > 0$.

The market for electricity in a given hour balances if total supply is equal to hourly demand D_t which, given our short-run analysis, we assume to be given and price-inelastic:

$$(8) \quad \sum_i y_{it} = D_t \quad \perp \quad P_t \text{ "free"} \quad \forall t.$$

Equations (6)–(8) imply that given demand the equilibrium allocation of hourly electricity supplies is determined by the available capacity and the marginal cost ordering of technologies. The equilibrium outcome of each plant i , y_{it}^* , thus depends on demand, and its own as well as the marginal cost and available capacities of all other plants (indicated by $-i$):

$$(9) \quad y_{it}^* = \mathcal{F}_{it} (D_t, c_{it}(\boldsymbol{\vartheta}_{it}), K_{it}, c_{(-i)t}(\boldsymbol{\vartheta}_{(-i)t}), K_{(-i)t}) .$$

Equation (9) identifies the major determinants of the power plants' outputs, including the responses to a carbon tax policy, by modelling wholesale market activity based on first principles of producer behavior and equilibrium-based market interactions.

C. Data Sources and Construction

To obtain measurements for the empirical counterparts of all variables in (9), we combine data from different, publicly available, sources. We use panel data of hourly generation for each UK fossil-fuel power plant covering the 2009-2016 period.

¹¹Equation (7) determines the price as the marginal cost of the marginal generator.

In addition, we use data on available hourly generation and capacity, demand, daily fuel and carbon prices, and temperature.¹²

HOURLY OUTPUT BY PLANT (y_{it}).—We use “final physical notification” (FPN) data provided by the operator of the UK electricity balancing system (ELEXON, 2016) as the hourly generation of each fossil power plant unit for the whole sample period. FPN reports the final, five minutes before delivery time generation announcement of power plant owners to the grid operator. Although the grid operator might adjust this announcement due to the need for re-dispatching measures, these data can be viewed as a reasonable measures for generation (which is not directly observable for UK power plants). As the data on carbon emissions are only available at a plant level, we aggregate power plant units to power plants for our analysis.

FUEL PRICES (p_t^{fuel}).—Data on daily fuel prices for coal and natural gas are taken from EIKON (2017). For coal, we use the “ICE CIF ARA Near Month Future”. Natural gas prices are “NBP Hub 1st Day Futures”. All prices are converted to Euro values using daily exchange rates provided by the ECB (2017). Figure 4 plots the time series of monthly-averaged daily fuel price ratio with and without the CPS showing a substantial variation—ranging approximately between 0.4 and 1—over the sample period.

CARBON PRICES (p_t^{CPS} and p_t^{EUA}).—CPS rates are reported by Hirst (2017) and HM Revenue & Customs (2014) and the EUA price by EEX (2017). Note that the CPS rate is an annually constant tax in British Pound but reflects exchange rate variations due to conversion to Euro values ECB (2017).

AVAILABLE CAPACITY BY PLANT BY HOUR (K_{it}).—Installed capacities (shown in Table 2) are provided by Variable Pitch (2016) and Nationalgrid (2011). If observed generation exceeds installed capacity beyond the 95th percentile, we set the value of installed capacity equal to the 95th percentile of generation.

In addition, data on the maximal plant-level output in a given hour—accounting for permanent and temporary outages due to maintenance or other reasons—the so-called “maximum export limits” (MEL), are provided by ELEXON (2016). Using hourly MEL, we construct a measure of available generation units for each plant.^{13,14}

DEMAND (D_t).—We measure D_t as residual demand, defined as the total output generated by all coal- and natural gas-fired plants using data from ELEXON (2016) on hourly generation aggregated by fuel type.

TEMPERATURE.—We use data on daily temperature provided by ECA&D (2016) to account for time-specific effects on plant-level heat efficiency.

For the ex-post calculation of the abatement impact, we use the following data:

EMISSIONS, EMISSIONS FACTORS AND PLANT-SPECIFIC HEAT EFFICIENCIES (E_{iy} , θ^f and

¹²Table A1 in Appendix A provides descriptive statistics of demand, generation by technology, and imports on an hourly level.

¹³Specifically, we set the availability of a unit to zero if MEL is zero, and to one otherwise. Summing over all units of a power plant, we obtain a count variable indicating the number of units available per plant, which we use as a proxy for hourly available capacity.

¹⁴Not all plants in our data run over the entire sample period from 2009–2016 (see Table 2). For years in our sample period during which a plant has been shut down or not yet opened, we set the capacity to zero. In line with this, we also do not predict its counterfactual generation different from zero for these periods, i.e. the impact of the CPS will be zero by assumption.

TABLE 2. Power plant characteristics.

Plant	Installed capacity [MW]	Average heat efficiency η_i [-]	Emissions rate e_i [ton of CO ₂ /MWh]	Opening/closing date ^a
<i>Natural gas plants</i>				
Pembroke	2269	0.60	0.34	end 2012/-
Peterhead	2134	0.55	0.36	-/March 2014
Staythorpe	1792	0.58	0.34	2010/-
Didcot CCGT	1404	0.55	0.36	-/-
Connahs Quay	1380	0.48	0.42	-/-
West Burton CCGT	1332	0.51	0.40	-/-
Grain CHP	1305	0.56	0.36	-/-
South Humber	1239	0.50	0.40	-/-
Seabank	1169	0.55	0.36	-/-
Saltend South	1164	0.52	0.38	-/-
Teesside	1155	0.45	0.44	-/Feb. 2013
Immingham CHP	1123	0.44	0.46	-/-
Barking	945	0.46	0.44	-/Dec. 2012
Langage	905	0.55	0.37	-/-
Marchwood	898	0.58	0.34	-/-
Killingholme	854	0.48	0.42	-/March 2015
Severn	850	0.54	0.37	-/-
Spalding	830	0.54	0.37	-/-
Rocksavage	800	0.53	0.38	-/-
Sutton Bridge	796	0.52	0.39	-/-
Damhead Creek	783	0.53	0.38	-/-
Coryton	770	0.52	0.38	-/-
Little Barford	740	0.54	0.37	-/-
Rye House	715	0.43	0.46	-/-
Keadby	700	0.47	0.42	-/Feb. 2013
Medway	680	0.53	0.38	-/-
Baglan Bay	520	0.57	0.35	-/-
Deeside	498	0.47	0.42	Dec. 2011/-
Great Yarmouth	420	0.56	0.35	-/-
Shoreham	420	0.54	0.37	-/-
Enfield Energy	408	0.53	0.38	-/-
Corby	401	0.39	0.51	-/Oct. 2015
Cottam CCGT	395	0.55	0.36	-/-
Kings Lynn	325	0.52	0.39	-/March 2012
Peterborough	316	0.37	0.54	-/Dec. 2011
Average natural gas plant ^b		0.51	0.40	
<i>Coal plants</i>				
Longannet	2304	0.42	0.81	-/March 2016
Didcot COAL	2108	0.39	0.88	-/March 2013
Cottam	2000	0.39	0.86	-/-
Ratcliffe	2000	0.38	0.89	-/-
West Burton COAL	1972	0.38	0.90	-/-
Fiddlers Ferry	1961	0.37	0.92	-/March 2016
Ferrybridge	1960	0.38	0.89	-/March 2016
Drax COAL	1947	0.38	0.90	-/-
Kingsnorth	1940	0.36	0.94	-/Dec. 2012
Eggborough	1932	0.37	0.92	-/-
Aberthaw	1641	0.41	0.82	-/-
Cockenzie	1200	0.38	0.91	-/March 2013
Rugeley	996	0.39	0.88	-/June 2016
Ironbridge	964	0.35	0.98	-/March 2012
Uskmouth	363	0.33	1.04	-/-
Average coal plant ^b		0.38	0.89	

Notes: Installed capacities, fuel type, and plant opening and closure dates are provided by [Variable Pitch \(2016\)](#) and [Nationalgrid \(2011\)](#). For data sources and calculations of heat efficiencies and emission rates see text. “-/-” indicates that the plants’ opening or closure date lies outside of the sample period 2009–2016.

^b Calculated using installed capacities as weights.




FIGURE 4. Carbon-price inclusive (r_t) and exclusive (\bar{r}_t) ratio of coal to natural gas fuel prices over the sample period 2009–2016

Notes: Monthly average values based on daily fuel prices for coal and natural gas taken from on [EIKON \(2017\)](#). For coal, we use the “*ICE CIF ARA Near Month Future*”. Natural gas prices are “*NBP Hub 1st Day Futures*”. All prices are converted to Euro values using daily exchange rates provided by the [ECB \(2017\)](#).

η_i).—We take fuel-specific emissions factors from [IPCC \(2006\)](#): 0.34 and 0.20 tons of CO₂ per MWh of thermal energy for coal and natural gas, respectively. CO₂ emissions for each plant i and year y (E_{iy}) are taken from the official registry of the EU ETS ([European Commission, 2016](#)). Dividing total emissions by total generation per plant, we obtain plant-specific average emission rates: $e_i = \sum_y E_{iy} / (\sum_t y_{it})$. We then calculate average heat efficiencies for each plant as:

$$(10) \quad \eta_i = \theta^f / e_i .$$

Table 2 shows these technical characteristics for each plant in the sample. The average heat efficiency is around 51 percent for natural gas and 38 percent for coal plants. The emission rates, on the other hand, are significantly higher for coal (0.89 tCO₂/MWh) than for gas (0.40 tCO₂/MWh). As we only observe emissions on an annual level, we can only calculate average heat efficiencies. Therefore, hourly changes in heat efficiencies due to, e.g., start-up or ramping constraints, are not considered in our calculations of the emissions impact of the CPS. To allow for time-varying heat efficiencies in our prediction models, we do not use these average heat efficiencies in our estimations.

D. Empirical Implementation

This section describes how we implement the method for estimating the treatment effect in the absence of a control group put described in Section I to assess the market impacts of the UK carbon tax.

ESTIMATION EQUATION.—The empirical analogue of hourly output of power plant i from the equilibrium electricity model in (9) is given by:

$$(11) \quad y_{it} = f_i \left[r_t \left(p_t^{coal}, p_t^{gas}, \theta^f, p_t^{EUA}, p_t^{CPS} \right), temp_t, D_t, K_{it}, K_{(-i)t}, \Phi_t \right] + \epsilon_{it}.$$

Here, we include daily mean temperature ($temp_t$) as a proxy for heat efficiency and marginal costs which we do not observe directly. We further include time fixed effects for each hour of the day and each month of the year (Φ_t) to account for possible unobserved factors which may impact plant output. We also include the carbon price inclusive ratio of relative fuel prices:

$$(12) \quad r_t := \frac{p_t^{coal} + \theta^{coal} (p_t^{EUA} + p_t^{CPS})}{p_t^{gas} + \theta^{gas} (p_t^{EUA} + p_t^{CPS})}.$$

While we are interested in the impact of the CPS on plants' output decisions, there is not sufficient variation in the treatment variable (p_t^{CPS}) as the CPS changes only in annual steps. As the CPS directly impacts the fuel costs for coal and natural gas, we can, however, exploit the variation in carbon-inclusive fuel prices—instead of including fuel prices (p_t^{coal} and p_t^{gas}) and carbon prices (p_t^{CPS} and p_t^{EUA}) separately. The implicit assumption here is that a change in fuel prices has the same impact on plants' marginal cost and, hence, output as a change in the carbon price (taking into account the emissions factor of the respective fuel θ^{fuel}). Furthermore, the use of r_t in equation (11) is well in line with the view that it is not the absolute but the relative fuel prices that determine which plants leave or remain in the market.

MACHINE LEARNING ALGORITHM.—From the electricity model in (9), and its empirical counterpart in (11), we know which variables affect plants' output decisions; we do not know, however, the functional form of f_i . To obtain an estimator \hat{f}_i of the function f_i , we use ML techniques, allowing for flexible functional forms to predict plant-level output y_{it} .

We employ the LASSO algorithm (Tibshirani, 1996)—a penalized linear regression model—and use k-fold cross-validation dividing the sample into eight groups to train a prediction model $\hat{f}_i^{\alpha^*}$ for each plant individually. The LASSO algorithm requires a pre-defined set of input features. In addition to the variables which appear on the right hand side of (11), we include (i) interaction terms of all these variables with electricity demand, the coal-to-gas price ratio, and temperature, and (ii) second order polynomials of these three variables. Each prediction model consists of the set of coefficients $\hat{\beta}^{\alpha^*}$ and the optimal regularization parameter α^* , which lead to the best possible prediction.¹⁵

ESTIMATING THE IMPACT OF THE CPS.—To simulate plants' outputs that would have occurred in the absence of the UK carbon tax, we set the CPS treatment

¹⁵ Appendix C assesses the prediction performance of the ML algorithm as compared to standard regression analysis (OLS) for our data set. We find that the LASSO algorithm outperforms OLS, supporting the well-known result that ML techniques can be beneficially employed to use prediction to construct an unobserved counterfactual.

variable to zero while leaving all other data unchanged. The counterfactual “no-policy” level of the fuel price ratio is given by:

$$(13) \quad \bar{r}_t := \frac{p_t^{coal} + \theta^{coal} p_t^{EUA}}{p_t^{gas} + \theta^{gas} p_t^{EUA}} .$$

Based on the estimator in equation (5), the impact of the CPS on the output decision of each plant i in each hour t can then be calculated as:

$$(14) \quad \hat{\delta}_{it}^{CPS} = \hat{y}_{it}^{\text{with CPS}} - \hat{y}_{it}^{\text{without CPS}} ,$$

where

$$\begin{aligned} \hat{y}_{it}^{\text{with CPS}} &= \hat{f}_i^{\alpha*}(r_t, temp_t, D_t, K_{it}, K_{(-i)t}, \Phi_{it}) \\ \hat{y}_{it}^{\text{without CPS}} &= \hat{f}_i^{\alpha*}(\bar{r}, temp_t, D_t, K_{it}, K_{(-i)t}, \Phi_{it}) . \end{aligned}$$

As a closed-form solution of standard errors of the prediction is not available for the LASSO regression (see, for example, [Tibshirani, 1996](#)), we use bootstrapping to estimate the standard errors of $\hat{\delta}_{it}^{CPS}$ ([Venables and Ripley, 2002](#)). We generate a bootstrap sample with the same length as the original data by using random drawings with replacement. We individually bootstrap by year to get the same amount of values from each year, thus ensuring that all years are equally represented in each sample so as to not violate Assumption 5.

MEASURING EMISSIONS AND ABATEMENT COSTS.—To calculate electricity-sector emissions, which derive from the combustion of coal and natural gas in power generation, we aggregate CO₂ emissions from all plants:

$$\bar{E}_t := \sum_i \underbrace{e_i \hat{y}_{it}^{\text{without CPS}}}_{\text{Plant-level emissions}}$$

where the emissions of plant i are obtained by multiplying output by the plant-specific emissions rate e_i (see Table 2). Given the estimator for the CPS impact on plant-level output ($\hat{\delta}_{it}^{CPS}$), we can calculate the change in electricity-sector emissions impact due to the CPS as follows:

$$(15) \quad \Delta E_t := \sum_i \underbrace{e_i \hat{\delta}_{it}^{CPS}}_{\text{Policy-induced change in emissions of plant } i (=:\Delta E_{it})} .$$

For our ex-post calculations, we assume marginal cost to be linear in fuel and carbon prices. Specifically, based on average heat efficiencies (given by equation

(10) and shown in Table 2) marginal cost are calculated as

$$(16) \quad c_{it}(\boldsymbol{\theta}_{it}) = \frac{1}{\eta_{it}} \left(p_t^f + \theta^f (p_t^{EUA} + p_t^{CPS}) \right).$$

Aggregate production costs are obtained by summing over marginal generation costs of all plants in the market at time t :

$$\Psi_t = \sum_i \hat{y}_{it}^{\text{with CPS}} c_{it}(\boldsymbol{\theta}_{it}) - \hat{y}_{it}^{\text{without CPS}} c_{it}(\boldsymbol{\theta}_{it}) \Big|_{p_t^{CPS}=0}.$$

Using the definition of the treatment effect from equation (14) and plant-specific heat efficiency from equation (10), this can be rewritten as follows:

$$(17) \quad \Psi_t = \underbrace{\sum_i \hat{\delta}_{it}^{CPS} \frac{1}{\eta_{it}} \left(p_t^f + \theta^f p_t^{EUA} \right)}_{\substack{=:T_t \\ \text{Technical abatement} \\ \text{cost}}} + \underbrace{\sum_i p_t^{CPS} e_i \hat{y}_{it}^{\text{with CPS}}}_{\substack{=:R_t \\ \text{Tax payments} \\ \text{due to CPS}}}.$$

Ψ_t can thus be decomposed into two parts. T reflects the technical abatement costs for the supply side of the market as the CPS affects plant output by re-ordering the supply or merit order curve. In other words, the CPS leads to an increase in (expensive) natural gas, and a decrease in (cheap) coal generation. This results in higher total production cost for the same amount of electricity generation.

R takes into account the costs incurred due to the CPS tax paid on each unit of generated emissions. While Ψ reflects the costs borne by the supply side of the electricity market, this decomposition is useful as the tax payments by electricity firms are typically recycled in a way which does not destroy the value of R . If, for example, the tax revenues from the CPS are fully rebated to electricity consumers, the costs of the CPS aggregated over both sides of the markets amount to T only.

III. Evaluating the Assumptions Underlying the Treatment Effect Estimation

This section examines the validity of the assumptions underlying the estimation strategy presented in Section I for the specific context of our data and application. We first evaluate the assumptions regarding the properties of the data (Assumptions 4 and 5) and the assumptions for identification (Assumptions 1-3).

A. Assumptions on Data

SUFFICIENTLY LARGE VARIATION IN TREATMENT AND CONTROLS (ASSUMPTION 4).—As described in Section II.D, we exploit the fact that the relative price of natural gas to coal influences input costs and power plant output decisions in the same way as the carbon tax (which exhibits little variation as it is only adjusted on an annual basis). Figure 5 shows that there is substantial variation in our modified

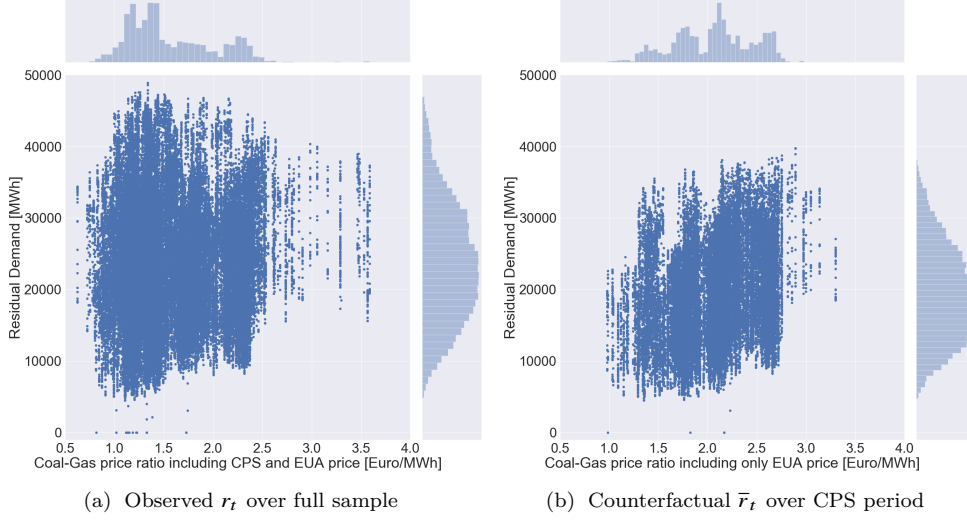


FIGURE 5. Joint distribution of electricity demand and the observed (r_t) and counterfactual (\bar{r}_t) coal-to-gas fuel price ratio

treatment variable (r_t and \bar{r}_t) and demand, as one key observed control variable, over the sample period.

POSITIVITY OR CO-VARIATE OVERLAP (ASSUMPTION 5).—ML methods construct predictors of the outcome variable based on a local approximation of the underlying unknown data generating process. The predictor is thus based on the observed sample but may not necessarily perform well in predicting out of the sample. To ensure a high prediction quality under counterfactual treatment, Assumption 5 requires that the counterfactual fuel price ratio \bar{r}_t lies within the range of observed fuel price ratios r_t conditional on observed control variables. Figure 5 shows the joint distribution of the observed fuel price ratio and residual demand over the full sample period 2009–2016 (Panel (a)) and the joint distribution of the counterfactual fuel price ratio and residual demand for the period after the CPS became effective, i.e. from April 2013 until the end of 2016 (Panel (b)). A comparison of Panel (a) and (b) shows that the imposed counterfactual fuel price ratios are well covered by the observed distribution.

B. Identifying Assumptions

INDEPENDENCE OF PREDICTION ERROR AND TREATMENT (ASSUMPTION 1).—We find strong evidence that this assumption is valid in the context of our application. First, we can assess the independence of the prediction error under *observed* treatment (i.e., ξ^{z_t} in Figure 1) with the treatment level. Computing the correlation between the fuel-price ratio r_t and $\xi_{it}^{z_t}$ for each power plant i at time t , we find that they are virtually independent.¹⁶ Second, we observe that counterfactual levels of

¹⁶The mean of correlation coefficients for all power plants i at different times t is 0.01 with a standard deviation of 0.07.

r lie within the range of the observed levels—as is evident from Figure 5 (and the discussion of Assumption 5). Given the low correlation between the *observed* prediction error and r_t , it is quite plausible that the prediction error is also independent from counterfactual treatment levels \bar{r} . Third, as we argue below related to the discussion of Assumption 3, we do not seem to miss important unobserved variables which could affect the errors differently for the observed and counterfactual predictions in the treatment period.

INDEPENDENCE OF OBSERVED CONTROLS FROM TREATMENT (ASSUMPTION 2).—This assumption excludes indirect effects of controls on outcome. In terms of our application context, this means the following. First, we require that the price of European Emission Allowances (EUA) traded under the European Emissions Trading System (ETS) and fuel prices are independent of the CPS. The EUA carbon price is determined by the EU ETS market, of which the UK electricity sector covers only a negligible part. The market share of UK electricity companies on international fuel markets is too small to influence fuel prices. Hence, both assumptions seem to be reasonable.

Assumption 2 excludes the possibility that the CPS will lead to a decrease in EUA carbon prices and coal prices due to lower demand for carbon and coal. If these prices were to react, we may slightly overestimate the treatment effect, as the price decline would favor carbon-intensive coal production.

Second, we require that available capacity is independent from treatment. As we measure the short-term market reactions to the CPS, this assumption is innocuous, as the installed capacity cannot easily be adjusted in shorter periods of time. It also rules out that plant closures are caused by the introduction of the CPS. The introduction of the CPS in 2013 coincides with the closure of several coal-fired power plants. Although the official reason for decommissioning is to be seen in a different policy—namely the European “Large Combustion Plant Directive” (see Section II.A)—we cannot entirely exclude the possibility that the closure of a few power plants may have been influenced by the announcement of the CPS policy.

Third, residual electricity demand is assumed to be inelastic, reflecting the short-term nature of our analysis. Hence, electricity demand is not responsive to hourly wholesale electricity prices, which could in turn be affected by a carbon tax. Moreover, the exogeneity of *residual* electricity demand also implies that the production of renewable energy sources, i.e. wind and solar, and base-load technologies, i.e. nuclear and hydro, are independent of the CPS. Output by these renewables and hydroelectric power is essentially governed by the availability of natural resources, and nuclear power plants typically do not change output levels much.

Essentially, assuming capacities and residual demand to be independent of the CPS reflects our focus on the short-term abatement induced by the CPS through a switch from coal to gas-based generation. Our estimates neglect longer-term effects such as plant closure, energy efficiency improvements, and investments into wind and solar generation. Thus, our results should be viewed as a lower-bound estimate of short-run abatement induced by the CPS.

CONDITIONAL INDEPENDENCE OF UNOBSERVED CONTROLS FROM TREATMENT (ASSUMPTION 3).—Unobserved controls comprise known and unknown unobserved vari-




FIGURE 6. Monthly average impacts of the UK carbon tax (CPS) on electricity output by technology

Notes: Shaded areas represent 95 % confidence intervals (based on bootstrapped standard errors). Values shown refer to estimated plant-level impacts $\hat{\delta}_{it}^{CPS}$, based on model specification M1 and equation (14), aggregated by technology category and month.

ables. For example, our population model of the UK electricity market ignores transmission and network restrictions which affect power generation. Whether these restrictions and other unknown factors are influenced by the level of CPS is naturally difficult to assess. We test the robustness of our model using a variety of different specifications for fixed effects. We find that our results are robust across different model specifications (see Table 3), suggesting that there do not seem to be unobserved variables, with significant systematic variation at the monthly and/or hourly level, that impact power plant output decisions. Moreover, we find that in all specifications the total net impact of the CPS on generation, i.e. the sum of the impacts on coal and natural gas, both on an annual and monthly basis does not statistically differ from zero (see Table 3 and Figure 6). A priori, this result was not to be expected as (1) the output decision of each power plant is estimated separately and (2) we do not impose an *explicit* market-clearing constraint in the empirical model—in contrast to the theoretical market model which requires that demand always equals supply as in equation (8). We interpret the statistical rejection of a violation of *implicit* market clearing as additional evidence that our empirical model is correctly specified and that we are not missing unobserved variables which affect power plant performance and depend on treatment.

IV. Results I: How Effective Was the UK Carbon Tax?

This section summarizes our main results on the plant-level and aggregate output, emissions, and cost impacts of the UK carbon tax in the period 2013-2016. Importantly, we also explore the drivers for environmental effectiveness and abate-

TABLE 3. Assessing unobserved heterogeneity: impact of the UK carbon tax (CPS) on aggregated power plant output by technology category for different model specifications.

	Model specification			
	M1	M2	M3	M4
Monthly fixed effects	<i>yes</i>	<i>no</i>	<i>yes</i>	<i>no</i>
Hourly fixed effects	<i>yes</i>	<i>no</i>	<i>no</i>	<i>yes</i>
<i>Coal</i>				
TWh	-46.29	-42.78	-43.17	-42.72
	(1.69)	(1.01)	(1.71)	(1.20)
% of total generation ^a	14.7	13.6	13.7	13.6
<i>Natural gas</i>				
TWh	45.55	45.00	46.01	45.23
	(1.06)	(0.92)	(1.07)	(0.75)
% of total generation ^a	15.0	14.9	15.2	14.9
<i>Total (TWh)</i>	-0.75	2.23	2.84	2.51
	(2.00)	(1.37)	(2.02)	(1.42)

Notes: Plant-level impacts δ_{it}^{CPS} based on equation (14). ^aRefers to situation without the CPS. Bootstrapped standard errors are shown in parentheses.

ment costs of the carbon tax.

A. Output and Emissions Impacts

PLANT-LEVEL AND MARKET OUTPUT.— Figure 7 shows the yearly impact of the CPS on generation of each plant in our sample. A clear main finding emerges: the CPS was quite effective in affecting the output decisions of UK power plants, leading to a pronounced decrease in coal- and an increase in gas-fired electricity generation (also see Table A3 in the Appendix A).

Table 4 shows the aggregate generation impacts of the CPS on coal and gas power plants for each year and the cumulative impact since its introduction in April 2013 until the end of 2016. We find that, in aggregate over all fossil-based power plants and until the end of 2016, the CPS caused a reduction in the output from coal-fired plants of 46.3 TWh and an increase from gas-fired plants of 45.6 TWh. These changes amount to a fuel switch from coal to natural gas in the order of 15 percent compared to a situation where the CPS policy would not have been introduced. Notably, the impact of the CPS on generation varies substantially over time. The fuel switch was initially low at a level of 4 TWh in the 2013 period and increased over the years with the highest value of 22 TWh in 2015. The impacts for both natural gas and coal are much larger in the 2015 than in the 2016 period. In relative terms, coal experienced the largest decrease in the 2016 period. In addition to the heterogeneity of the annual generation effects, Figure 6 shows that the CPS-induced effects on the monthly production of coal- and gas-fired power plants also vary considerably over time.

EMISSIONS AND ABATEMENT COSTS.—Table 5 summarizes the total and yearly impacts of the CPS on electricity-sector CO₂ emissions and abatement cost. We estimate that over the period 2013–2016, the CPS has reduced cumulative emissions




FIGURE 7. Generation impacts by plant

Notes: Gas-fired power plants are shown on the left side; coal-fired on the right.

by 26.1 million tons—corresponding to a 6.4 percent reduction of total emissions as compared to a situation without a CPS. Applying our measure of technical abatement costs T (see equation (17)), the CPS has reduced one ton of CO₂ emissions at an average cost of €18.2 over this period.¹⁷

B. Temporal Heterogeneity of Impacts

An important empirical finding is that there is considerable temporal heterogeneity in the abatement quantity and cost impacts of the UK carbon tax—both across and within CPS periods. First, aggregate CO₂ emissions reductions vary between 1.7 and 11.9 percent relative to a situation without a CPS and average technical abatement cost amount to €2.7 in 2016 to €47.5 in 2013 per ton abated CO₂ (see Table 5). Second, Figure 8 shows unequivocally that abatement quantity and costs impacts largely vary within a given period for a given level of the carbon tax level. Hence, the variation of the CPS level alone can thus not explain the observed variation in the impacts of the carbon tax.

These empirical findings bear out two important results—which run counter to

¹⁷ Although not the primary focus of our analysis, Table 5 also reports the tax revenues collected with the CPS instrument. Since its introduction and until the end of 2016, the British government received around €5.2 billion in tax revenue from the CPS policy. There is temporal heterogeneity in the magnitude of tax revenues collected: the highest tax revenues (around €2 billion) accrued in 2015 when both, emissions and the CPS level, were high; in the subsequent period, the CPS tax revenue dropped significantly as CO₂ emissions remaining in the market were considerably lower.

TABLE 4. Impacts of the UK carbon tax (CPS) on aggregated power plant output by fuel type.

	Period				Total impact
	2013	2014	2015	2016	2013-2016
<i>CPS</i> [€ per ton of CO ₂]	5.85	12.17	24.70	21.60	–
<i>Change in output from coal plants</i>					
TWh	-4.17	-9.26	-21.92	-10.94	-46.29
	(0.27)	(0.57)	(0.86)	(0.21)	(1.69)
% of total generation	-3.7	-9.8	-27.0	-43.6	-14.7
<i>Change in output from natural gas plants</i>					
TWh	4.27	9.37	21.19	10.72	45.55
	(0.10)	(0.23)	(0.57)	(0.40)	(1.06)
% of total generation	6.1	12.1	29.7	12.8	15.0
<i>Total</i> [TWh]	0.10	0.11	-0.73	-0.22	-0.75
	(0.29)	(0.62)	(1.03)	(0.45)	(2.00)

Notes: As the CPS is adjusted in April of every year, all reported variables refer to the period from April to March of the subsequent year. As data is available until December 2016, the 2016 period comprises only nine months. Values shown refer to estimated plant-level impacts $\hat{\delta}_{it}^{CPS}$, based on the model specification with time fixed effects (M1) and equation (14), aggregated by technology category. Bootstrapped standard errors are shown in parentheses.

the common intuition about the economic impacts of carbon taxes:

RESULT 1: *A higher carbon tax does not necessarily lead to a larger reduction in CO₂ emissions.*

RESULT 2: *A higher carbon tax does not necessarily imply greater average abatement costs.*

The upshot of these results is that the empirical relationships between the tax level, abatement quantity, and abatement costs are highly non-linear. This raises the fundamental question for the design of an effective price-based climate policy: what drives the environmental effectiveness and abatement costs of a carbon tax? We next turn to an investigation of this question.

V. Results II: Why Did the Environmental and Cost Effectiveness of the UK Carbon Tax Vary?

This section uses simulations with the ML-trained model to investigate what drives the heterogeneity in the abatement quantity and costs impacts caused by the UK carbon tax. Specifically, we aim to explore the nature of the non-linear empirical relationships behind Results 1 and 2. We investigate to what extent the insights gained from our empirical model are compatible with first principles of microeconomic theory for cost-optimizing firm behavior in electricity markets.

A. Conceptualizing the Determinants of Marginal Abatement Cost

Basic microeconomic theory suggests that cost-optimizing firms choose the level of abatement which equalizes marginal abatement costs (MAC) and marginal abatement benefits (MAB). MAB reflect the avoided tax payments per unit of emissions,

TABLE 5. Impacts of the UK carbon tax (CPS) on aggregate emissions and abatement costs

	Period				Total impact
	2013	2014	2015	2016	2013-2016
<i>CPS</i> [€/t]	5.85	12.17	24.70	21.60	—
<i>Emissions without CPS</i> (\bar{E}) [Mt]	125.8	112.0	98.0	71.3	407.1
<i>CO₂ abatement</i>					
ΔE_t [Mt]	2.1 (0.25)	4.7 (0.53)	11.6 (0.81)	7.6 (0.24)	26.1 (1.60)
% of total emissions	1.7	4.2	11.9	10.7	6.4
<i>Abatement cost</i> $\Psi_t = T_t + R_t$					
Technical cost T_t [mio. €]	101.1 (9.2)	129.1 (18.4)	195.1 (29.1)	20.5 (16.6)	445.0 (58.7)
Avg. tech. cost $T_t/\Delta E_t$ [€/t]	47.5 (12.5)	27.2 (8.7)	16.8 (4.0)	2.7 (2.3)	18.2 (4.0)
Tax payments R_t [mio. €]	725.7	1309.6	2129.4	1372.8	5194.3

Notes: Values shown refer to estimated plant-level impacts $\hat{\delta}_{it}^{CPS}$, based on model specification M1 and equation (14), aggregated by period. As the CPS is adjusted in April of every year, all reported variables refer to the period from April to March of the subsequent year. As data is available until December 2016, we can only estimate the impacts of the CPS for a nine month period. To ensure comparability with previous years, we scale model values for 2016 to a 12-month basis. Bootstrapped standard errors are shown in parentheses.

i.e. the level of the carbon tax (p^{CPS}). We next provide a theory-founded explanation for the empirically observed non-linear relationships between the carbon tax level, abatement quantity, and costs behind Results 1 and 2.

SIMPLE MICROECONOMICS.—Consider an electricity firm characterized by a plant portfolio including coal and gas power plants. The firm seeks to minimize its carbon tax-induced impact on production costs by choosing abatement:

$$\begin{aligned}
 \min_{a \geq 0} \Psi &= \underbrace{T(a; \bar{r})}_{\text{Technical abatement cost}} + \underbrace{(\bar{E} - a)p^{CPS}}_{\text{Tax payments due to CPS (=R)}} \\
 s.t. \quad a &\leq \underbrace{\Gamma(\bar{r})}_{\text{Maximum abatement potential}} \quad (\mu).
 \end{aligned}$$

The total impact on production costs $\Psi = T + R$ —in line with equation (17)—is given by the sum of technical abatement costs $T(a; \bar{r})$, which are a function of chosen abatement a and a *given* carbon tax-exclusive fuel-price ratio \bar{r} , and tax payments on unabated emissions $R = (\bar{E} - a)p^{CPS}$, where \bar{E} denotes “no-policy” emissions in the absence of a carbon tax.

The constraint simply expresses the fact that given a certain portfolio of fossil-based power plants in the market, there exists a maximum potential or capacity of abatement $\Gamma(\bar{r})$ that is attainable. The maximum potential depends on \bar{r} as the relative fuel prices of coal and natural gas affect the technology mix of gas-




FIGURE 8. Empirical relationship between average abatement costs and the quantity of CO₂ emissions abated for different periods (i.e., levels) of the UK carbon tax (CPS)

Notes: Average abatement costs shown refer to daily averages of hourly average abatement costs (HAC), where for a given hour HAC are calculated as $T_t/\Delta E_t$ using estimated treatment effects in equations (15) and (17).

vs. coal-fired power plants. For example, if the price of coal increases relative to the gas price the abatement potential decreases as natural gas generation starts to replace coal even in the absence of a carbon tax. $\mu \geq 0$ denotes the multiplier associated with the abatement potential constraint.

Deriving the Karush-Kuhn-Tucker (KKT) conditions for the optimal choice of abatement yields:¹⁸

$$(18) \quad \underbrace{\frac{\partial T(a; \bar{r})}{\partial a}}_{\text{Marginal technical abatement costs (MTAC)}} + \underbrace{\mu}_{\text{Marginal rent on abatement potential}} \geq \underbrace{p^{CPS}}_{\text{Marginal benefits of abatement (MAB)}} \perp a \geq 0.$$

Marginal abatement costs (MAC)

The firm seeks to equate the MAB, i.e., the carbon price, to the MAC which comprise two components: the marginal technical abatement costs (MTAC) and the marginal rent on the abatement potential (μ). The MTAC component reflects the fuel costs incurred to lower emissions by reducing electricity output from coal-

¹⁸Here, the “ \perp ” operator expresses complementarity between the difference of MAC and MAB, on the one hand, and optimal abatement a , on the other hand. It is short-hand notation for writing the KKT conditions: $\frac{\partial T}{\partial a} + \mu \geq p^{CPS}$, $a \geq 0$, $(\frac{\partial T}{\partial a} + \mu - p^{CPS})a = 0$. For example, in the absence of a carbon tax (i.e., $p^{CPS} = 0$), the KKT conditions imply that in the optimum $a = 0$; a positive amount of abatement requires that $\text{MAC} = \text{MAB}$ in the optimum.

fired while increasing output from gas-fired power plants. For given fuel costs \bar{r} , MTAC are typically referred to as the engineering-based estimate of marginal abatement costs. μ represents the shadow price on the maximum capacity or the potential for abatement. It measures how strongly the abatement constraint binds at the optimal solution. If the maximum abatement potential constraint is not binding, μ is zero and MAC are given by the MTAC only. Conversely, if only a limited abatement potential remains (e.g., because most of the coal power plants have already been driven out of the market), μ is large and the MAC exceed the MTAC.

CONJECTURES.—The KKT conditions in (18) enables deriving several conjectures about what drives the MAC and how this impacts the environmental effectiveness and costs of a carbon tax.

CONJECTURE 1: *For a given fuel price ratio \bar{r} , marginal abatement cost weakly increase with abatement ($\partial \text{MAC} / \partial a = \partial^2 T(a; \bar{r}) / \partial a^2 + \partial \mu / \partial a \geq 0$).*

CONJECTURE 2: *The marginal technical abatement cost weakly decreases in the fuel price ratio ($\partial^2 T(a; \bar{r}) / \partial a \partial \bar{r} \leq 0$).*

CONJECTURE 3: *The maximum abatement potential weakly decreases in the relative price of coal to natural gas ($\partial \Gamma(\bar{r}) / \partial \bar{r} \leq 0$), implying that the marginal rent on the abatement potential weakly increases in \bar{r} ($\partial \mu / \partial \bar{r} \geq 0$).*

Conjecture 1 describes the behavior of MAC regarding abatement. It simply states that MAC, and both of its constituent components, increase in abatement for a given fuel price ratio.

The next two conjectures state that the two components which make up total MAC depend on the level of the fuel price ratio \bar{r} . Conjecture 2 expresses the idea that fuel-switching between coal and natural gas becomes cheaper with an increasing fuel price ratio. This is directly implied by the definition of technical abatement cost as the cost of switching from coal to gas: if the fuel price of coal is already relatively high compared to the fuel price of natural gas (high \bar{r}), a given abatement level can be achieved at smaller MTAC. Conjecture 3 can be understood as follows. As \bar{r} increases, the cost of gas-fired power plants relative to coal plants decrease, driving some coal plants out of the market even without a carbon tax. To the extent that fewer coal plants are available for fuel-switching in response to a carbon tax, the maximum abatement potential declines with \bar{r} . If the abatement potential constraint is binding, a smaller Γ due to an increased \bar{r} implies that the shadow price on abatement capacity μ , which positively contributes to MAC, must be higher. If Conjectures 2 and 3 hold true, a change in the fuel price ratio has an ambiguous effect on the MAC: MTAC decrease in \bar{r} while μ increase in \bar{r} .

IMPLICATIONS FOR THE QUANTITY AND COSTS OF CO₂ ABATEMENT.—What do the conjectures imply about the drivers of environmental effectiveness and costs of a carbon tax? According to Conjecture 1, abatement increases with an increasing carbon tax and cost increase for a *given* fuel price ratio. For a *given* carbon tax the effect of an increase in the fuel price ratio \bar{r} on abatement is ambiguous. If MAC decrease, the effect of decreasing MTAC outweighs the increase in the shadow price

of abatement potential μ . Consequently, environmental effectiveness increases with increasing \bar{r} .

The impact on total abatement cost, however, is ambiguous as abatement increases but MAC decrease. In contrast, if MAC increase, the increase in μ exceeds the effect of decreasing MTAC. Consequently, the environmental effectiveness is decreasing as μ only becomes positive if the abatement potential is fully used and as the potential is decreasing in \bar{r} . Total abatement cost then also decrease, due to a decrease in MTAC and abatement.

B. Empirical Marginal Abatement Costs

This section presents results from performing simulations with our ML-trained model to investigate Conjectures 1–3 in the given empirical context of the UK carbon tax. This enables us to empirically analyze the determinants of the environmental effectiveness and costs of the UK carbon tax with a handshake on microeconomic theory.

COMPUTATIONAL DERIVATIONS.—To obtain empirical counterparts of the MAC (left hand side of (18)), we perform simulations with the ML-trained model deriving abatement quantities for different levels of the carbon tax (increasing the CPS level in increments of 1 from 0-50 €/tCO₂). To analyze the dependence of MAC on the fuel price ratio, we derive MAC curves for three different ranges for \bar{r} representing “Low” ($\bar{r} < 0.55$), “Intermediate” ($0.55 \leq \bar{r} \leq 0.88$), and “High” ($\bar{r} > 0.88$) values.¹⁹

MAIN RESULTS.—Figure 9 shows the empirical MAC curves for the different ranges of \bar{r}_t . Several insights emerge. First, for a given level of the fuel price ratio, the empirical MAC curve is monotonically increasing in abatement—which is consistent with expectations from economic theory that $\partial \text{MAC} / \partial a > 0$ and therefore Conjecture 1. We also find that empirical MAC curves are convex (i.e. $\partial^2 \text{MAC} / \partial a^2 > 0$). Second, we find a non-monotonous impact of the fuel price ratio on MAC: moving from “Low” to “Intermediate” values of \bar{r}_t slightly decreases MAC, while for “High” values of \bar{r}_t the MAC increase substantially again. This non-monotonicity is due to the opposing effects of the different MAC components in (18) with respect to a change in \bar{r} hypothesized in Conjectures 2 and 3. As \bar{r} increases, MTAC decrease as it becomes cheaper to substitute coal by gas-fired plants. At the same time, however, as gas plants become more favorable, coal plants are driven out of the market (even without carbon price), in turn lowering the remaining abatement potential which escalates MAC by increasing the shadow costs of available abatement capacity μ .

DISENTANGLING MTAC AND ABATEMENT POTENTIAL EFFECTS.—Figure 10 provides a more detailed analysis of the two MAC components by visualizing the change in the empirically-measured counterparts of the MTAC and the abatement potential

¹⁹The choice of cutoff points for \bar{r}_t is motivated by the following considerations. The “Low” value corresponds to the carbon tax-exclusive fuel price ratio for which the most efficient gas plant (*Pembroke* plant) substitutes for the most inefficient coal plant (*Uskmouth* plant)—given the observed, plant-specific heat efficiencies in Table 2. The lower end of “Intermediate” range thus contains values of \bar{r}_t for which gas-fired plants begin to move, in the absence of a carbon tax, to the left of the merit order dispatch curve. The “High” value corresponds to the fuel price ratio for which the least efficient gas plant (*Rye House* plant) breaks even, in terms of fuel costs, with the least efficient coal plant (*Uskmouth* plant).




FIGURE 9. Empirical MAC curves for different carbon tax-exclusive fuel price ratios \bar{r}_t




FIGURE 10. Empirical relationships between hourly MTAC, hourly abatement potential, and the fuel price ratio \bar{r}_t

Notes: Each dot corresponds to an hourly value which is computed based on plant-specific heat efficiencies in Table 2 and fuel costs. Empirical MTAC $[\partial^2 T(a; \bar{r}) / \partial a \partial \bar{r}]$ are measured by the minimum carbon price necessary to induce a switch from coal to natural gas triggering a “small” amount of abatement. The maximum abatement potential $\Gamma(\bar{r}_t)$ is measured as the quantity of CO₂ emissions abated if *all* coal plants were replaced by gas plants.

as the fuel price ratio \bar{r}_t varies. We measure MTAC as the minimum carbon price necessary to induce abatement, or equivalently, a switch from coal to natural gas (where we use the data on heat efficiencies from Table 2 and hourly electricity

demand). The maximum abatement potential $\Gamma(\bar{r}_t)$ is calculated as the quantity of CO₂ emissions abated if *all* coal power plants were replaced by gas power plants.

Figure 10 documents empirical evidence in support of Conjectures 2 and 3. First, in the range of “*Low*” fuel price ratios, the MTAC rapidly diminish as \bar{r}_t increases; the abatement potential, however, largely remains on a high level. Second, at the lower bound of the “*Intermediate*” range (i.e., $\bar{r}_t = 0.55$), gas plants begin replacing coal plants even in the absence of a carbon tax, implying that Γ starts to decrease. For this range of fuel price ratios, the carbon tax-exclusive fuel costs of gas plants are roughly equal to those of coal plants, implying that the MTAC are close to zero, i.e. a very small carbon tax would be sufficient to create a cost advantage for gas plants. Third, at the transition from “*Intermediate*” to “*High*” values of \bar{r}_t , all gas plants are cheaper than the least efficient coal plant even without a carbon tax. Thus, MTAC are very low but the abatement potential is virtually exhausted.

The opposing effects of the constituent components of total MAC visualized in Figure 10 explain the change in MAC curves as the fuel price ratio varies (compare with Figure 9). MAC decrease when going from “*Low*” to “*Intermediate*” values of \bar{r}_t (i.e., green to red curve) due to the fact that MTAC fall while μ is small as the abatement potential constraint is slack. Further increases in \bar{r}_t drive up μ as the abatement potential diminishes, and in turn drive up total MAC even though MTAC are close to zero (i.e., red to blue curve). Moreover, the increasing shadow costs of abatement capacity imply that the degree of convexity of the MAC curves—for the range of abatement quantities shown in Figure 10—increases with \bar{r}_t . For “*High*” values of the fuel price ratio, MAC increase super-proportionally with the abatement quantity. MAC curves for “*Low*” to “*Intermediate*” values of \bar{r}_t , on the other side, are closer to linearity.




C. Does a Higher Carbon Tax Always Imply Larger Costs and Abatement?

Equipped with the theoretical and empirical insights on the counteracting effects of the fuel price ratio on MAC, environmental effectiveness and abatement cost, we now return to investigating Results 1 and 2. In particular, we examine the non-linear relationships between the tax rate, abatement quantity, and average abatement costs in order to investigate whether or not a higher carbon tax necessarily leads to more abatement and higher average costs.

ENVIRONMENTAL EFFECTIVENESS AND AVERAGE ABATEMENT COSTS.—To assess how abatement quantity and average cost impacts depend on the fuel price ratio, we compare CO₂ emissions reductions for different hypothetical levels of the carbon tax while using observational variation in the data for \bar{r}_t .²⁰ Figure 11 plots the empirical relationships between the fuel price ratio \bar{r}_t and CO₂ abatement (Panel a), total abatement costs (Panel b), and the average technical abatement costs (Panel c) for different levels of the CPS.

Panel (a) shows that the environmental effectiveness of a carbon tax largely depends on the prevailing relative (carbon tax-exclusive) fuel prices of coal and

²⁰This is equivalent to using the ML-trained model to computationally evaluate the KKT conditions in (18) to find cost-minimizing emissions abatement a for a given carbon tax rate p^{CPS} .

(a) CO₂ abatement (ΔE_t)(b) Total technical abatement costs (T_t)(c) Average abatement costs ($T_t/\Delta E_t$)FIGURE 11. Empirical relationships between CO₂ abatement, abatement costs, and the fuel price ratio \bar{r}_t for different carbon tax rates p^{CPS}

natural gas. This is consistent with the MTAC and abatement potential effects analyzed in Sections V.A and V.B. It is straightforward to see that a higher carbon tax does not necessarily imply a higher CO₂ abatement. Graphically speaking, there is a substantial overlap for the range of abatement induced by different levels of tax rates. For example, the carbon tax rate increases and abatement decreases from point A to point B to point C, i.e. a carbon tax rate of 20 €/ton CO₂ induces larger abatement than higher tax rates of 30 and 40 €/ton CO₂, respectively.

For a given carbon tax rate, we observe a humped-shaped pattern between abatement quantity and \bar{r}_t . A carbon tax is most effective at reducing CO₂ emissions for intermediate values of \bar{r}_t , i.e. when fuel costs of coal are neither “too” cheap nor “too” costly relative to the fuel costs of natural gas. In our empirical example, carbon abatement peaks at the point where the fuel costs of coal are about 60% of those of natural gas. The explanation is that for these fuel price ratios, gas plants are just as cheap as coal plants. MTAC are therefore near zero but the abatement potential is still large. Thus, a given carbon tax is effective at inducing a fuel switch at modest cost (reflected by low MTAC) while it can tap into a large abatement potential in the market (thus avoiding large shadow costs of abatement capacity μ).

Panel (b) shows that total abatement costs monotonically fall in \bar{r}_t , and they become zero once the abatement potential is exhausted (i.e. $\bar{r}_t \geq 0.7$). Comparing the different tax levels, the figure bears two main insights. First, a higher carbon tax does not necessarily imply larger total abatement costs as they crucially depend on the relative fuel price of coal to natural gas. For example, when \bar{r}_t is low, a carbon tax of 30 €/ton CO₂ (see point B’) implies much higher costs than what is borne out by a carbon tax rate of 40 €/ton CO₂ (see point C’) when \bar{r}_t is high. Second, comparing abatement B and cost B’ with abatement A and cost A’ we find that a lower abatement (B) can induce higher total abatement cost (B’) than a higher abatement (A).

Panel (c) combines the quantity and total costs impacts from Panels (a) and (b). Average abatement costs monotonically fall in \bar{r}_t . The important insight is that while average abatement costs do not vary much in the level of the carbon tax rate, they crucially depend on the coal-to-gas fuel price ratio.

D. Applying the Conceptual Insights: Heterogeneous Impacts of the UK Carbon Tax

While Figure 11 used hypothetical variations in the carbon tax rate to illustrate the relationships between policy stringency, abatement, and abatement costs, we can finally analyze the heterogeneous quantity and cost impacts triggered by the UK carbon tax. Figure 12 visualizes the effects of the UK carbon tax along the four relevant dimensions in a single diagram: average abatement costs (vertical axis), the coal-to-gas price ratio \bar{r}_t (horizontal axis), abatement quantity (color code), and CPS periods, corresponding to different levels of carbon tax (marker type). Several important insights emerge.

First, average abatement costs $T_t/\Delta E_t$ decrease as \bar{r}_t increases (similar to the pattern shown in Figure 11, Panel c): the more expensive coal becomes relative to gas, the smaller are the MTAC associated with a tax-induced fuel switch (T_t

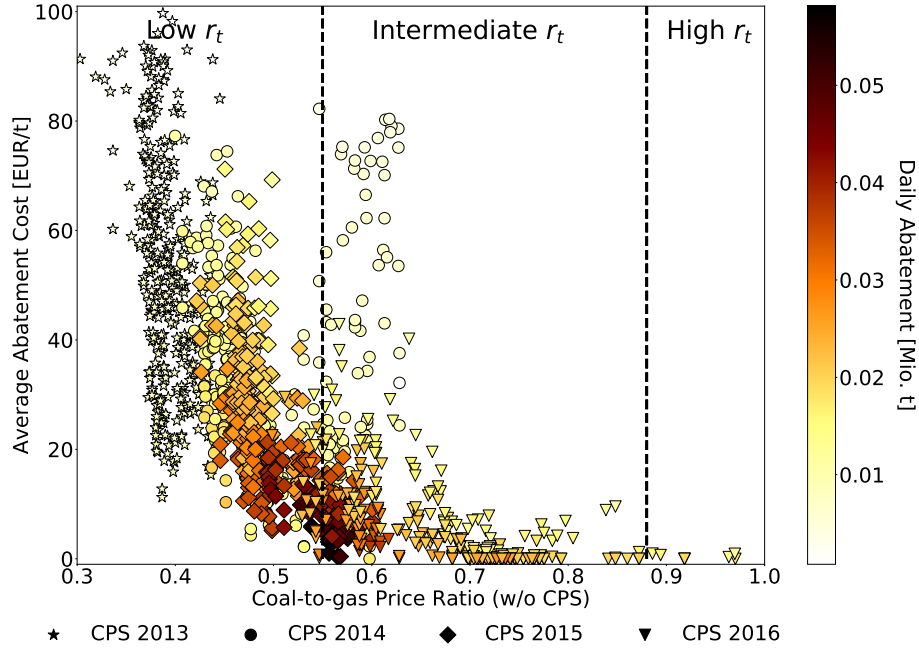


FIGURE 12. Relation between daily CO₂ abatement, and daily average technical abatement cost, and daily (average) fuel price ratio r_t .

Notes: All values refer to daily averages of hourly values. Average abatement costs shown refer to daily averages of hourly average abatement costs (HAC), where for a given hour HAC are calculated as $T_t / \Delta E_t$ using estimated treatment effects in equations (15) and (17).

declines). Second, a larger fuel price ratio increases the tax-induced quantity of CO₂ emission reductions (ΔE_t increases), up to the point where the abatement potential starts to decrease.

Third, it is evident that the level of the CPS does not solely determine the environmental effectiveness and abatement costs. In the 2013 period, the UK CPS was low and it coincided with fuel market conditions which implied a low fuel price ratio \bar{r}_t . Average abatement costs were thus high (see stars in the range “Low \bar{r} ”). In the 2014 period, the CPS was higher and the relative price of coal to gas increased as compared to 2013. Abatement was thus higher and average abatement cost decreased (see circles in the regions “Low \bar{r} ” and “Intermediate \bar{r} ”). In the 2014 and 2015 periods, the fuel price ratios were closely around the values generating peak abatement with low MTAC and a high abatement potential. While fuel price ratios were similar in 2014 and 2015, a higher CPS tax rate in 2015 implied higher abatement as compared to 2014. In the 2016 period, the fuel price ratio was high implying that the abatement potential was nearly exhausted. This implied that despite a still high CPS level, abatement in 2016 was lower relative to 2015 (see diamonds in the regions “Intermediate \bar{r} ” and “High \bar{r} ”).

In summary, our results indicate that—while the UK carbon tax has been effective in reducing CO₂ emissions in the electricity sector—there is considerable temporal heterogeneity in abatement quantities and costs, resulting from the varia-

tion of the relative fuel prices for coal and natural gas. The important implication for climate policy is that a higher carbon tax does not necessarily deliver high emissions reductions. At the same time, a higher carbon tax need not necessarily result in higher abatement costs.

VI. Conclusions

While economists see carbon pricing as one of the main policy instruments for mitigating climate change, knowledge about its performance from an ex-post perspective is limited. Assessing ex-post the effects of a broad-based carbon tax, i.e. one which affects virtually all CO₂-emitting units in a market, is fraught with difficulties because a suitable control group or “no-policy” counterfactual typically does not exist.

Against this background, this paper has made two contributions. First, we propose an approach for policy evaluation which combines economic theory and machine learning (ML) techniques in settings with high-frequency data when no control group exists. Specifically, we exploit economic theory of electricity market dispatch and peak-load pricing to select the variables of a prediction model which is then trained using ML to obtain an empirical prediction model for power plant output. We obtain the treatment effect of a carbon tax on plant-level electricity output as the difference between predicted outcomes with observed and counterfactual (i.e., no) carbon tax policy.

Second, this paper has applied this new approach to evaluate the environmental effectiveness and costs of the UK CPS—a carbon levy imposed on all fossil-based power plants in the electricity market. To our knowledge, this is the first paper in economics to incorporate ML methods to assess the ex-post effects of carbon pricing. Our analysis provides empirical evidence in support of the view that a carbon tax can be an effective regulatory instrument to reduce CO₂ emissions: the CPS induced a substitution away from “dirty” coal to cleaner natural gas-fired power plants—replacing about 15 percent or 46 TWh of coal-based generation and reducing electricity sector emissions by 6.2 percent between 2013 and 2016. Over that period, we find that the abatement of one ton of CO₂ incurred additional total costs of €18.2 for consumers and fossil-based electricity producers.

We find that there is substantial heterogeneity in the carbon tax-induced market impacts over time, which are mainly driven by the level of the tax rate and the ratio of carbon tax-exclusive prices for coal and natural gas. Our analysis thus contributes with an empirically-founded characterization of the conditions under which a tax-based climate policy can be more or less effective. An important policy implication is that in the short run higher carbon taxes does not necessarily bring about higher emissions reductions. At the same time, however, a higher carbon tax need not necessarily result in higher abatement costs.

Some limitations of our analysis should be kept in mind. First, we focus on analyzing the short-run market impacts. Thus, we abstract from potential effects of the CPS on energy demand, installed fossil capacities, and investments in low-carbon electricity production capacity. This implies that we also do not take into account the possible impacts of the CPS on plant closure. Although we assume

plant closures to be driven by existing regulation unrelated to the CPS, i.e. the European “Large Combustion Plant Directive”, we cannot rule out that the shut-down decision for some plants may also have been influenced by the announcement of the CPS as we observe that its introduction in 2013 coincides with the closure of several coal power plants. Second, by increasing domestic wholesale market prices relative to the costs of electricity imports, the CPS may have stimulated electricity imports. To the extent that such effects reduce (domestic) CO₂ emissions for a given tax level, our analysis should best be viewed as providing a lower-bound empirical estimate of the environmental effectiveness of the UK carbon tax.

REFERENCES

- Abrell, Jan, and Sebastian Rausch.** 2016. “Cross-Country Electricity Trade, Renewable Energy and European Transmission Infrastructure Policy.” *Journal of Environmental Economics and Management*, 79: 87–133.
- Aldrich, J.** 1989. “Autonomy.” *Oxford Economic Papers*, 41: 15–34.
- Angrist, Joshua D., and Joern-Steffen Pischke.** 2008. *Most Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press.
- Athey, Susan.** 2017. “Beyond Prediction: Using Big data for Policy Problems.” *Science*, 335(6324): 483–485.
- Athey, Susan, and Guido Imbens.** 2017. “The State of Applied Econometrics-Causality and Policy Evaluation.” *Journal of Economic Perspectives*, 31(2): 3–32.
- Bajari, Patrick, Denis Nekipelov, Stephen P Ryan, and Miaoyu Yang.** 2015. “Machine learning methods for demand estimation.” *The American Economic Review: Papers & Proceedings*, 105(5): 481–485.
- Boiteux, M.** 1960. “Peak-Load Pricing.” *The Journal of Business*, 33(2): pp. 157–179.
- Bretschger, Lucas, Fillippo Lechthaler, Sebastian Rausch, and Lin Zhang.** 2017. “Knowledge Diffusion, Endogenous Growth, and the Costs of Global Climate Policy.” *European Economic Review*, 93: 47–72.
- Bretschger, Lukas, and Elise Grieg.** 2020. “Exiting the fossil world: The effects of fuel taxation in the UK.” Economics Working Paper Series - CER-ETH - Center of Economic Research at ETH Zurich.
- Burlig, Fiona, Christopher Knittel, David Rapson, Mar Reguant, and Catherine Wolfram.** 2019. “Machine learning from schools about energy efficiency.” E2e Working Paper 032.
- Chevallier, Julien, Erik Delarue, Emeric Lujan, and William D’Haeseleer.** 2012. “A counterfactual simulation exercise of CO₂ emissions abatement through fuel-switching in the UK (2008–2012).” *International Journal of Global Energy Issues*, 35(5): 311–331.
- Chyong, Chi Kong, Bowei Guo, and David Newbery.** 2020. “The impact of a Carbon Tax on the CO₂ emissions reduction of wind.” *The Energy Journal*, 41(1).
- CMA.** 2016. “Competition and Markets Authority: Energy market investigation.”
- Cullen, Joseph A, and Erin T Mansur.** 2017. “Inferring carbon abatement costs in electricity markets: A revealed preference approach using the shale revolution.” *American Economic Journal: Economic Policy*, 9(3): 106–33.
- Delarue, Erik D, A Denny Ellerman, and William D D’Haeseleer.** 2010a. “Robust MACCs? The topography of abatement by fuel switching in the European power sector.” *Energy*, 35(3): 1465–1475.
- Delarue, Erik D, A Denny Ellerman, and William D D’Haeseleer.** 2010b. “Short-term CO₂ abatement in the European power sector: 2005–2006.” *Climate Change Economics*, 1(02): 113–133.
- Delarue, Erik, Kris Voorspools, and William D’Haeseleer.** 2008. “Fuel switching in the electricity sector under the EU ETS: review and prospective.” *Journal of Energy Engineering*, 134(2): 40–46.
- Department of Energy & Climate Change.** 2016. “2014 UK Greenhouse Gas Emissions, Final Figures.” *Statistical release*.
- ECA&D.** 2016. “European Climate Assessment & Data. <http://eca.knmi.nl/dailydata/predefinedseries.php>.”
- ECB.** 2017. “European Central Bank: Daily exchange rates. www.ecb.europa.eu/stats.”
- EEX.** 2017. “European Energy Exchange: Daily Spot Prices EU Emission Allowances.”
- EIKON.** 2017. “Thomson reuters EIKON database.”
- ELEXON.** 2016. “ELEXON: www.elexon.co.uk.”

- Ellerman, Denny, and A. Decaux. 1998. "Analysis of Post-Kyoto CO₂ Emissions Trading Using Marginal Abatement Curves." *MIT Joint Program on the Science and Policy of GlobalChangeReport 40*, October. (http://web.mit.edu/globalchange/www/MITJPSPGC_Rpt40.pdf).
- European Commission. 2001. "Directive 2001/80/EC of the European Parliament and of the Council of 23 October 2001 on the limitation of emissions of certain pollutants into the air from large combustion plant." *Official Journal of the European Community*.
- European Commission. 2016. "DG Climate Action: Euroepan Union Transaction Log." <http://ec.europa.eu/environment/ets>.
- Fawcett, Allen A., Leon E. Clarke, Sebastian Rausch, and John P. Weyant. 2014. "The EMF Study on U.S. Technology and Climate Policy Strategies: Introduction to EMF24." *Energy Journal*, 35(1): 1–9.
- Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur. 2012. "What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NO_x Trading Program." *American Economic Review*, 102(2): 965–993.
- Gareth, James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An Introduction to Statistical Learning - with Application in R*. Springer Texts in Statistics.
- Gissey, G. Castagneto, B. Guo, D. Newbery, G. Lipman, L. Montoya, P. Dodds, M. Grubb, and P. Ekins. 2019. "The value of international electricity trading." A project commissioned by Ofgem.
- Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski. 2014. "Optimal Taxes on Fossil Fuel in General Equilibrium." *Econometrica*, 82(1): 41–88.
- Goulder, Lawrence H., and Ian W. H. Parry. 2008. "Instrument Choice in Environmental Policy." *Review of Environmental Economics and Policy*, 2(2): 152–174.
- Goulder, Lawrence H., Marc A. C. Hafstead, and Roberton C. Williams III. 2016. "General Equilibrium Impacts of a Federal Clean Energy Standard." *American Economic Journal: Economic Policy*, 8(2): 186–218.
- Gugler, Klaus, Adhurim Haxhimusa, and Mario Liebensteiner. 2020. "Carbon Pricing and Emissions: Causal Effects of Britains Carbon Tax." Working Paper.
- Haavelmo, T. 1944. "The probability approach in econometrics." *Econometrica*, 12: 1–115.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2008. *The Elements of Statistical Learning - Data Mining, Inference, and Prediction*. Springer Series in Statistics.
- Hirst, David. 2017. "Carbon Price Floor (CPF) and the price support mechanism." *House of Commons - Briefing paper (05927)*.
- HM Revenue & Customs. 2014. "Carbon price floor: Reform." www.gov.uk/government/publications/carbon-price-floor-reform.
- HM Treasury. 2017. "Autumn Budget 2017." www.gov.uk/government.
- Holladay, J Scott, and Jacob LaRiviere. 2017. "The impact of cheap natural gas on marginal emissions from electricity generation and implications for energy policy." *Journal of Environmental Economics and Management*, 85: 205–227.
- Holladay, J Scott, Steven Soloway, et al. 2016. "The Environmental Impacts of Fuel Switching Electricity Generators." *The Energy Journal*, 37(4).
- Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association*, 81(396): 945–960.
- IPCC. 2006. *Revised 1996 IPCC Guidelines for National Greenhouse Gas Inventories*. IPCC.
- IPCC. 2018. "An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty." <https://www.ipcc.ch/sr15/>.
- Jarvis, Stephen, Olivier Deschenes, and Akshaya Jha. 2020. "The Private and External Costs of Germany's Nuclear Phase-Out." Energy Institute WP 304.
- Kirat, Djamel, and Ibrahim Ahamada. 2011. "The impact of the European Union emission trading scheme on the electricity-generation sector." *Energy Economics*, 33(5): 995–1003.
- Klepper, Gernot, and Sonja Peterson. 2006. "Marginal abatement cost curves in general equilibrium: The influence of world energy prices." *Resource and Energy Economics*, 28(1): 1–23.
- Knittel, Christopher R., Konstantinos Metaxoglou, and Andre Trindade. 2015. "Natural Gas Prices and Coal Displacement: Evidence from Electricity Markets." National Bureau of Economic Research.
- Leroutier, Marion. 2019. "The Impact of the UK Carbon Price Floor on CO₂." *Unpublished working paper*.

- Linn, Joshua, Lucija Muehlenbachs, and Yushuang Wang. 2014. "How Do Natural Gas Prices Affect Electricity Consumers and the Environment?" *Resources for the Future Discussion Paper No. 14-19*.
- Liski, Matti, and Reyer Gerlagh. 2016. "Carbon Prices for the Next Hundred Years." *Economic Journal*, 128(609): 728–757.
- Lu, Xi, Jackson Salovaara, and Michael B. McElroy. 2012. "Implications of the Recent Reductions in Natural Gas Prices for Emissions of CO₂ from the US Power Sector." *Environmental Science & Technology*, 46(5): 3014–3021.
- Martin, Ralf, Laure B De Preux, and Ulrich J Wagner. 2014. "The impact of a carbon tax on manufacturing: Evidence from microdata." *Journal of Public Economics*, 117: 1–14.
- Martin, Ralf, Mirabelle Muûls, and Ulrich J Wagner. 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.
- Mathiesen, Lars. 1985. "Computation of Economic Equilibria by a Sequence of Linear Complementarity Problems." *Mathematical Programming Study*, 23: 144–162.
- McGlade, Christophe, and Paul Ekins. 2015. "The geographical distribution of fossil fuels unused when limiting global warming to 2 degree C." *Nature*, 517: 187–190.
- McGuinness, Meghan, and A Denny Ellerman. 2008. "CO₂ abatement in the UK power sector: evidence from the EU ETS trial period." *MIT CEEPR Working Paper*, 08-010.
- Metcalf, Gilbert E. 2009. "Market-Based Policy Options to Control U.S. Greenhouse Gas Emissions." *Journal of Economic Perspectives*, 23(2): 5–27.
- Morris, Jennifer, Sergey Paltsev, and John M. Reilly. 2012. "Marginal Abatement Costs and Marginal Welfare Costs for Greenhouse Gas Emissions Reductions: Results from the EPPA Model." *Environmental Modeling and Assessment*, 17(4): 325–336.
- Mullainathan, Sendhil, and Jann Spiess. 2017. "Machine Learning: An Applied Econometric Approach." *Journal of Economic Perspectives*, 31(2): 87–106.
- Nationalgrid. 2011. "Electricity Then Year Statement 2011 - Appendix F." www.nationalgrid.com/uk/publications.
- Nationalgrid. 2016. "Final electricity demand." www.nationalgrid.com/uk.
- Nordhaus, William D. 1994. *Managing the Global Commons: The Economics of Climate Change*. MIT Press, ISBN: 9780262140553, 223 pp.
- Peters, Jonas, Peter Bühlmann, and Nicolai Meinshausen. 2016. "Causal inference by using invariant prediction: Identification and confidence intervals." *Journal of the Royal Statistical Society: Series B*, 78(5): 947–1012.
- Pettersson, Fredrik, Patrik Söderholm, and Robert Lundmark. 2012. "Fuel switching and climate and energy policies in the European power generation sector: A generalized Leontief model." *Energy Economics*, 34(4): 1064–1073.
- Rausch, Sebastian, and Matthew Mowers. 2014. "Distributional and efficiency impacts of clean and renewable energy standards for electricity." *Resource and Energy Economics*, 36(2): 556 – 585.
- Rubin, Donald B. 1974. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology*, 66(5): 688–701.
- Rutherford, Thomas F. 1995. "Extension of {GAMS} for complementarity problems arising in applied economic analysis." *Journal of Economic Dynamics and Control*, 19(8): 1299 – 1324.
- Samii, Cyrus, Laura Paler, and Sarah Zukerman Daly. 2016. "Retrospective Causal Inference with Machine Learning Ensembles: An Application to Anti-recidivism Policies in Colombia." *Political Analysis*, 24: 434–456.
- Tavoni, Massimo, Elmar Kriegler, Keywan Riahi, Detlef P. van Vuuren, Tino Aboumahboub, Alex Bowen, Katherine Calvin, Emanuele Campiglio, Tom Kober, Jessica Jewell, Gunnar Luderer, Giacomo Marangoni, David McCollum, Mariosse van Sluisveld, Anne Zimmer, and Bob van der Zwaan. 2014. "Post-2020 climate agreements in the major economies assessed in the light of global models." *Nature Climate Change*, 5: 119–126.
- Tibshirani, Robert. 1996. "Regression Shrinkage and Selection via the Lasso." *Journal of the Royal Statistical Society. Series B (Methodological)*, 58: 267–288.
- Variable Pitch. 2016. "Variable Pitch." www.variablepitch.co.uk.
- Varian, Hal R. 2014. "Big data: New tricks for econometrics." *The Journal of Economic Perspectives*, 28(2): 3–27.
- Varian, Hal R. 2016. "Causal inference in economics and marketing." *Proceedings of the National Academy of Sciences*, 113(27): 7310–7315.

- Venables, W.N., and B.D. Ripley.** 2002. *Modern Applied Statistics with S - Fourth Edition*.
- Weyant, John, Brigitte Knopf, Enrica De Cian, Ilkka Keppo, and Detlef P. van Vuuren.** 2013. "Introduction to the EMF28 Study on Scenarios for Transforming the European Energy System." *Climate Change Economics*, 4(1): 1–3.
- World Bank.** 2018. "State and Trends of Carbon Pricing 2018." *World Bank, Washington, DC*. Doi: 10.1596/978-1-4648-1292-7.

APPENDIX A: ADDITIONAL FIGURES AND TABLES

TABLE A1. Descriptive statistics: annual means and standard deviations of observed hourly electricity demand, generation, and imports by technology category.

	2009	2010	2011	2012	2013	2014	2015	2016
Residual demand	27.10 (6.51)	28.33 (6.58)	25.81 (6.63)	24.99 (6.77)	23.77 (6.93)	22.16 (6.23)	20.01 (6.36)	19.54 (6.43)
Total demand	36.55 (7.76)	37.27 (8.15)	35.79 (7.68)	35.88 (7.52)	35.89 (7.74)	34.56 (7.40)	34.21 (7.47)	33.70 (7.74)
Gas	17.14 (3.01)	18.29 (3.07)	14.56 (3.79)	9.50 (4.16)	9.17 (5.12)	9.81 (4.87)	9.47 (4.43)	14.23 (4.75)
Coal	9.81 (5.80)	9.97 (5.29)	10.70 (5.14)	14.35 (4.04)	13.11 (3.18)	10.13 (4.10)	8.17 (3.45)	3.27 (2.88)
Nuclear	7.41 (1.03)	6.67 (1.12)	7.39 (1.13)	7.51 (0.83)	7.53 (0.97)	6.82 (1.04)	7.50 (0.61)	7.60 (0.66)
Hydro	0.41 (0.22)	0.24 (0.17)	0.42 (0.21)	0.37 (0.22)	0.33 (0.24)	0.45 (0.27)	0.47 (0.26)	0.38 (0.26)
PSP	-0.13 (1.14)	-0.11 (1.01)	-0.09 (0.95)	-0.11 (0.96)	-0.11 (0.92)	-0.11 (0.93)	-0.10 (0.90)	-0.12 (0.96)
Other	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.24 (0.25)	0.44 (0.34)	0.85 (0.26)	1.29 (0.53)	1.62 (0.46)
Wind	1.02 (0.66)	1.16 (0.82)	1.74 (1.15)	2.00 (1.43)	2.80 (1.79)	3.24 (2.17)	3.70 (2.26)	3.63 (3.08)
Solar	0.00 (0.00)	0.00 (0.00)	0.02 (0.03)	0.14 (0.21)	0.35 (0.56)	0.57 (0.85)	0.96 (1.48)	1.11 (1.64)
Imports	0.15 (1.28)	0.06 (1.44)	0.54 (1.17)	1.13 (1.13)	1.49 (0.86)	2.22 (0.51)	2.37 (0.65)	2.03 (1.20)

Notes: Standard deviations in parentheses. Data for generation by fuel type is based on [ELEXON \(2016\)](#). [Nationalgrid \(2016\)](#) provides data for final demand and embedded wind and solar generation.

TABLE A2. Descriptive statistics: installed annual generation capacities by technology category [GW].

	2009	2010	2011	2012	2013	2014	2015	2016
Gas	20.9	23.0	23.4	25.0	24.2	24.1	23.7	23.6
Coal	25.3	25.3	25.3	24.5	19.9	19.1	19.1	15.3
Hydro	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9
Nuclear	11.2	11.2	11.2	11.2	11.2	11.2	11.2	11.2
OCGT	1.4	1.4	1.4	1.4	1.3	1.3	1.3	1.3
Oil	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7
Other	2.9	2.9	2.9	2.9	2.9	2.9	2.9	2.9
PSP	2.7	2.7	2.7	2.7	2.7	2.7	2.7	2.7
Imports	2.5	2.5	3.5	3.6	4.0	4.0	4.0	4.0

Notes: Installed capacities are provided by [Variable Pitch \(2016\)](#) and [Nationalgrid \(2011\)](#). Plant characteristics of individual coal and gas plants, i.e. heat efficiencies, emission rates, installed capacities as opening and closure dates are shown in Table 2.

TABLE A3. Impacts of UK carbon tax (CPS) on power plant output [TWh].

	Period				Total impact 2013-2016
	2013	2014	2015	2016	
<i>Natural gas plants</i>					
Pembroke	0.38	0.84	2.01	0.70	3.94
Peterhead	0.00	0.00	0.00	0.00	0.00
Staythorpe	0.19	0.65	1.40	0.29	2.53
Didcot CCGT	0.52	0.85	2.26	1.05	4.68
Connahs Quay	0.28	0.58	1.04	0.68	2.58
West Burton CCGT	0.04	0.36	0.91	0.32	1.63
Grain CHP	0.21	0.66	1.39	0.37	2.63
South Humber	0.17	0.35	0.63	0.41	1.55
Seabank	0.36	0.76	1.36	0.88	3.36
Saltend South	0.07	0.17	0.67	0.49	1.41
Immingham CHP	0.18	0.37	0.66	0.43	1.64
Langage	0.23	0.29	1.00	0.83	2.35
Marchwood	0.04	0.08	0.14	0.09	0.35
Severn	0.12	0.25	0.44	0.28	1.09
Spalding	0.29	0.66	1.67	0.76	3.38
Rocksavage	0.05	0.11	0.46	0.29	0.92
Sutton Bridge	0.08	0.18	0.31	0.20	0.77
Damhead Creek	0.00	0.00	0.00	0.00	0.00
Coryton	0.11	0.24	0.43	0.28	1.07
Little Barford	0.00	0.00	0.00	0.00	0.00
Rye House	0.06	0.11	0.17	0.09	0.43
Medway	0.18	0.61	1.23	0.34	2.36
Baglan Bay	0.05	0.22	0.42	0.33	1.02
Deeside	0.07	0.15	0.26	0.17	0.65
Great Yarmouth	0.23	0.28	0.91	0.44	1.86
Shoreham	0.01	-0.05	0.17	0.12	0.25
Enfield Energy	0.10	0.21	0.37	0.26	0.94
Corby	0.08	0.14	0.13	0.00	0.35
Cottam CCGT	0.18	0.32	0.79	0.63	1.92
Fellside	0.00	0.00	0.00	0.00	0.00
Fawley Cogen	0.00	-0.01	-0.02	-0.01	-0.04
Grangemouth	-0.01	-0.01	-0.02	-0.02	-0.06
<i>Coal plants</i>					
Longannet	0.00	0.00	0.00	0.00	0.00
Cottam	-0.88	-2.15	-6.95	-3.47	-13.46
Ratcliffe	-0.39	-0.82	-1.46	-0.95	-3.61
West Burton COAL	-1.10	-2.47	-5.98	-3.33	-12.89
Fiddlers Ferry	0.00	0.00	0.00	0.00	0.00
Ferrybridge	0.00	0.00	0.00	0.00	0.00
Drax COAL	-0.69	-1.64	-3.71	-2.22	-8.25
Eggborough	-0.83	-1.77	-2.74	-0.59	-5.93
Aberthaw	0.00	0.00	0.00	0.00	0.00
Rugeley	-0.18	-0.40	-0.71	-0.14	-1.43
Uskmouth	-0.09	-0.01	-0.36	-0.26	-0.72

Notes: Values shown refer to estimated plant-level impacts $\hat{\delta}_{it}^{CPS}$, based on model specification M1 and equation (14). As the CPS is adjusted in April of every year, all reported variables refer to the period from April to March of the subsequent year. As data is available until December 2016, the 2016 period comprises only nine months. The plants are ordered from high to low according to their installed capacity (see Table 2).

APPENDIX B: USING MACHINE LEARNING FOR PREDICTION MODELS

To predict counterfactual outcomes $y_{it}^{\bar{z}}$, we need an estimator \hat{f}_i of the function f_i that produces reliable predictions. We harness the power of ML methods which—in contrast to traditional econometric methods focused on consistently estimating parameters of f —are optimized to predict the value of the outcome variable (Mullainathan and Spiess, 2017).

Machine learning methods typically minimize estimates of the expected prediction error. We use the mean squared error (MSE) as a measure of prediction quality whose expected value can be decomposed as follows (see, for example, Hastie, Tibshirani and Friedman, 2008; Gareth et al., 2013):

$$(B1) \quad \mathbb{E}[\text{MSE}_i] = \mathbb{E}[(y_i - \hat{f}_i)^2] = \underbrace{\sigma_\epsilon^2 + \left(\mathbb{E}[\hat{f}_i] - f_i\right)^2}_{= \text{Bias}^2(\hat{f}_i)} + \underbrace{\mathbb{E}[(\mathbb{E}[\hat{f}_i] - \hat{f}_i)^2]}_{= \text{Variance}(\hat{f}_i)}.$$

The expected prediction error thus consists of three parts: an irreducible population error, which corresponds to the variance of the random noise σ_ϵ^2 , and bias and variance terms which are both reducible. Standard econometric techniques such as OLS aim at minimizing the bias while allowing for high variance. While these methods are thus capable of representing very well the sample data, they are prone to over-fitting and they yield prediction outcomes that are highly dependent on the observed sample.

ML methods, in contrast, solve a bias-variance trade-off in order to find the best prediction model. They address this trade-off by introducing *hyper- or tuning parameters* in the estimation function. These parameters control for model complexity by decreasing the variance at the cost of a higher bias. The selection of hyper-parameters α is achieved through a process called cross-validation (CV), which makes optimal use of the available data. The CV process starts by splitting the observed sample into several subsets. One of the subsets, called the training set, is then used to estimate the predictor for a given set of hyper-parameters, \hat{f}_i^α , by minimizing the expected *in-sample* MSE:

$$(B2) \quad \hat{f}_i^\alpha := \arg \min_{f_i \in \mathcal{F}} \sum_t \left[(y_{it} - \hat{f}_i^\alpha(x_{it}, z_t)) \right]^2$$

where \mathcal{F} denotes the set of all possible functions f_i . The out-of-sample MSE is then computed on the remaining data—called the test or hold-out set—which has not been used for the estimation. Repeating this procedure for all subsets and averaging over all out-of-sample MSE yields an estimate of the expected prediction error for a given set of hyper-parameters α .

The optimal set of hyper-parameters α^* is the one that minimizes the expected prediction error which is obtained from using a grid search over different candidate sets. Given α^* , the final predictor $\hat{f}_i^{\alpha^*}$ is obtained by solving the problem in equation (B2) on the full sample of data. Finally, the true value of outcome in equation (1) can be written as the the sum of the predicted value and the prediction error $\xi(x_{it}, h_{it}, z_t)$:

$$(B3) \quad y_{it} = \hat{f}_i^{\alpha^*}(x_{it}, z_t) + \underbrace{f_i(x_{it}, h_{it}, z_t) - \hat{f}_i^{\alpha^*}(x_{it}, z_t)}_{=: \xi(x_{it}, h_{it}, z_t)} + \epsilon_{it}.$$

APPENDIX C: MACHINE LEARNING (LASSO) ALGORITHM VERSUS OLS

This section compares the prediction performance of the LASSO algorithm versus a standard linear OLS regression model. The comparison of both models is based on the same input variables (and data) as specified in equation (11).

To assess model performance, we proceed in three steps. First, we split out data into eight different pairs of train- and hold-out samples, i.e. each time we use all but one year to train the model and use the remaining year as a hold-out set. Consequently, each of the years 2009 to 2016 is used once as a hold-out set while the rest of the sample is used to train the model. Second, we use each train set to build the models which predict hourly generation y_{it} on a set of input features x_{it} and z_t for each $i \in I$, separately. In this step, we perform cross-validation to tune the regularization parameter α . The final step compares different types of models with respect to their in-sample and out-of-sample prediction

performance. We can assess for each plant the predictive performance by hold-out year and model type. We use the coefficient of determination—defined as $1 - \sum_i (y_i - \hat{y}_i)^2 / (\sum_i (y_i - \bar{y}_i)^2)$ —as the score function to evaluate model performance. A test score of 1.0 indicates that the model perfectly predicts the observed data. Note that, in contrast to the commonly reported R^2 , the test score can be negative because the model can be arbitrarily poor.

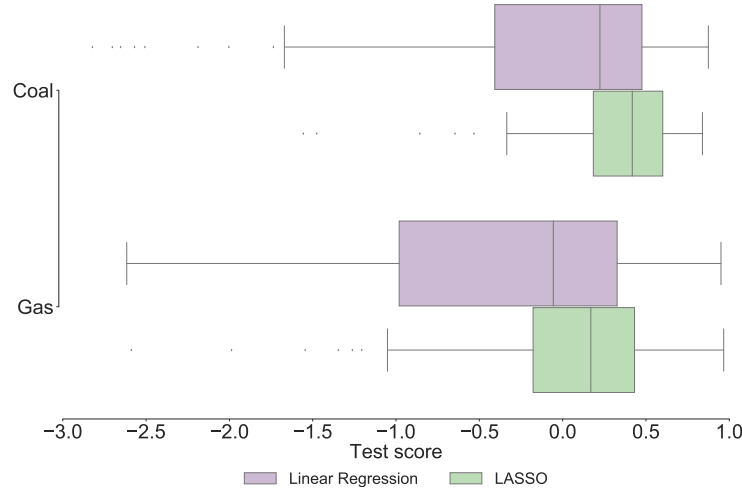


FIGURE C1. Comparison of the distribution of plant-specific performance scores by fuel type for LASSO vs. OLS models.

Figure C1 compares the test scores of the LASSO and OLS algorithms assessing the prediction of the hold-out set. It is evident that the LASSO outperforms the OLS model in terms of out-of-sample prediction: both average mean scores for coal- and gas-fired plants are higher for LASSO and the respective inter-quartiles ranges are significantly smaller under LASSO as compared to OLS. While from a conceptual perspective the qualitative ranking of LASSO and OLS models in terms of out-of-sample performance are not surprising, Figure C1 makes the important point that in the context of the suggested framework for policy evaluation (and given the specific empirical context), the use of a ML method is advantageous.

Working Papers of the Center of Economic Research at ETH Zurich

(PDF-files of the Working Papers can be downloaded at www.cer.ethz.ch/research/working-papers.html).

- 20/337 F. Böser, C. C. Senni
Emission-based Interest Rates and the Transition to a Low-carbon Economy
- 20/336 L. Bretschger, E. Grieg, P. J.J. Welfens, T. Xiong
Corona Fatality Development and the Environment: Empirical Evidence for OECD Countries
- 20/335 M. Arvaniti, W. Habla
The Political Economy of Negotiating International Carbon Markets
- 20/334 N. Boogen, C. Daminato, M. Filippini, A. Obrist
Can Information about Energy Costs Affect Consumers Choices? Evidence from a Field Experiment
- 20/333 M. Filippini, N. Kumar, S. Srinivasan
Nudging the Adoption of Fuel-Efficient Vehicles: Evidence from a Stated Choice Experiment in Nepal
- 20/332 L. Bretschger, E. Grieg
Exiting the fossil world: The effects of fuel taxation in the UK
- 20/331 H. Gersbach, E. Komarov
Research Bubbles
- 20/330 E. V. Dioikitopoulos, C. Karydas
Sustainability traps: patience and innovation
- 19/329 M. Arvaniti, C. K. Krishnamurthy, A. Crepin
Time-consistent resource management with regime shifts
- 19/328 L. Bretschger, K. Pittel
Twenty Key Questions in Environmental and Resource Economics
- 19/327 C. Karydas, A. Xepapadeas
Climate change financial risks: pricing and portfolio allocation
- 19/326 M. Filippini, S. Srinivasan
Investments in Worker Health and Labor Productivity: Evidence from Vietnam
- 19/325 H. Gersbach
Democratizing Tech Giants! A Roadmap

- 19/324 A. Brausmann, M. Flubacher and F. Lechthaler
Valuing meteorological services in resource-constrained settings: Application to smallholder farmers in the Peruvian Altiplano
- 19/323 C. Devaux and J. Nicolai
Designing an EU Ship Recycling Licence: A Roadmap
- 19/322 H. Gersbach
Flexible Majority Rules for Cryptocurrency Issuance
- 19/321 K. Gillingham, S. Houde and A. van Benthem
Consumer Myopia in Vehicle Purchases: Evidence from a Natural Experiment
- 19/320 L. Bretschger
Malthus in the Light of Climate Change
- 19/319 J. Ing and J. Nicolai
Dirty versus Clean Firms' Relocation under International Trade and Imperfect Competition
- 19/318 J. Ing and J. Nicolai
North-South diffusion of climate-mitigation technologies: The crowding-out effect on relocation
- 19/317 J. Abrell, M. Kosch and S. Rausch
How Effective Was the UK Carbon Tax? - A Machine Learning Approach to Policy Evaluation
- 19/316 D. Cerruti, C. Daminato and M. Filippini
The Impact of Policy Awareness: Evidence from Vehicle Choices Response to Fiscal Incentives
- 19/315 M. Filippini, N. Kumar and S. Srinivasan
Energy-related financial literacy and bounded rationality in appliance replacement attitudes: Evidence from Nepal
- 19/314 S. Houde and E. Myers
Heterogeneous (Mis-) Perceptions of Energy Costs: Implications for Measurement and Policy Design
- 19/313 S. Houde and E. Myers
Are Consumers Attentive to Local Energy Costs? Evidence from the Appliance Market
- 19/312 N. Kumar
A model-based clustering approach for analyzing energy-related financial literacy and its determinants

- 19/311 C. Karydas and A. Xepapadeas
Pricing climate change risks: CAPM with rare disasters and stochastic probabilities
- 19/310 J. Abrell, S. Rausch and C. Streitberger
Buffering Volatility: Storage Investments and Technology-Specific Renewable Energy Support
- 19/309 V. Britz
Negotiating with frictions
- 19/308 H. Gersbach and S. Papageorgiou
On Banking Regulation and Lobbying
- 18/307 V. Britz, A. Ebrahimi and H. Gersbach
Incentive Pay for Policy-makers?
- 18/306 C. Colesanti Senni and N. Reidt
Transport policies in a two-sided market
- 18/305 A. Schäfer and A. Stünzi
The impact of green preferences on the relevance of history versus expectations
- 18/304 M. Filippini and S. Srinivasan
Impact of religious participation, social interactions and globalisation on meat consumption: evidence from India
- 18/303 H. Gersbach and M.-C. Riekhof
Permit Markets, Carbon Prices and the Creation of Innovation Clusters
- 18/302 M. Hersche and E. Moor
Identification of Causal Intensive Margin Effects by Difference-in-Difference Methods
- 18/301 L. Kleemann and M.-C. Riekhof
Changing background risk and risk-taking - Evidence from the field
- 18/300 J. Blasch and C. Daminato
Behavioral anomalies and energy-related individual choices: the role of status-quo bias
- 18/299 S. Rausch and H. Schwerin
Does Higher Energy Efficiency Lower Economy-Wide Energy Use?
- 18/298 H. Gersbach, U. Schetter and M. Schneider
Economic Rationales for Investments in Science
- 18/297 K. Borissov and L. Bretschger
Optimal Carbon Policies in a Dynamic Heterogenous World