

Abrell, Jan; Kosch, Miriam; Rausch, Sebastian

Working Paper

How effective was the UK carbon tax? A machine learning approach to policy evaluation

Economics Working Paper Series, No. 19/317

Provided in Cooperation with:

ETH Zurich, CER-ETH – Center of Economic Research

Suggested Citation: Abrell, Jan; Kosch, Miriam; Rausch, Sebastian (2019) : How effective was the UK carbon tax? A machine learning approach to policy evaluation, Economics Working Paper Series, No. 19/317, ETH Zurich, CER-ETH - Center of Economic Research, Zurich, <https://doi.org/10.3929/ethz-b-000338148>

This Version is available at:

<https://hdl.handle.net/10419/207139>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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



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 and S. Rausch

Working Paper 19/317
April 2019

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*

APRIL 2019

Carbon taxes are commonly seen as a rational policy response to climate change, but little is known about their performance from an ex-post perspective. 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 novel approach for policy evaluation which leverages economic theory and machine learning techniques 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 a higher carbon tax does not necessarily lead to higher emissions reductions or higher costs. (JEL 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 work¹, surprisingly little is known about the *ex-post* effects of carbon pricing. This, however, is pivotal for designing effective and efficient climate policies in the future.

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*

* Abrell: Center for Energy and Environment of ZHAW Winterthur and Centre for Energy Policy and Economics ETH Zurich, Switzerland (email: jan.abrell@zhaw.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).

Support (CPS). The CPS was introduced to enhance economic incentives for CO₂ 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 main 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 difference-in-differences (DiD), regression discontinuity design, and synthetic control 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 establish causal inference 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 conduct causal inference of carbon pricing.

The proposed approach to estimate treatment effects in the absence of a control group comprises three main steps. We first derive a structural causal model of the observed outcomes (electricity output by power plant) based on an economic model of wholesale market activity. We then use ML techniques optimized for out-of-sample prediction (Mullainathan and Spiess, 2017) to train the causal model, i.e. to estimate a predictor function for electricity output of each plant in the market given fuel prices, available capacities, and demand. We train the model based on both pre- and post-intervention data combining hourly panel data of electricity output at the plant level with market information on hourly demand, production capacities by power plant, fuel and carbon prices, temperature. As the UK CPS is adjusted annually, thereby only providing four tax rate levels over our sample period, we exploit the variation in the relative market prices for coal and natural gas in order to overcome the problem of insufficient variation in the treatment variable. As the market substitution from carbon-intensive coal-fired power plants to less polluting natural gas plants hinges on relative fuel prices, we identify the impact of the CPS on plant output by using the carbon price-inclusive fuel price ratio as the pseudo-treatment variable. The core idea of our approach, following Pearl (2009) and his “do()” operator concept, rests on performing a counterfactual intervention based on a causal model: we predict the outcome for a counterfactual value of the treatment using a causal and ML-trained model in which one can control the treatment variable. Given the objective to evaluate the electricity market impacts of the UK CPS, we estimate the treatment effect for each power plant as the difference between predicted outcomes with and without the CPS policy. In doing so, we account for the potential impacts stemming from unobserved variables as well as systematic prediction errors related the ML algorithm.

An important advantage of our approach that is based on a causal, ML-trained model is its ability to explicitly represent the channels through which the policy intervention affects the outcome variable. As our pseudo-treatment variable—the carbon price-inclusive fuel price ratio, which trivially contains the case of no CPS

policy—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 impacts the economic behavior of firms’ output. 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 causal, 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, we find that there is substantial heterogeneity in the carbon tax-induced abatement quantity and costs impacts 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. (2017) and Cicala (2017) are two recent examples using ML algorithms to estimate causal effects. Our approach differs in two important ways. First, they deal with discrete treatment leading to a change in the data generating process (DGP) between the pre- and the post-treatment period. They, therefore, use the pre-treatment period to train a model predicting

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).

the post-treatment outcome without the intervention. In contrast, we deal with a setting with an invariant DGP and continuous treatments. Therefore, we are able to 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. \(2017\)](#) and [Cicala \(2017\)](#) assume that prediction errors have similar trends across treatment and control groups. Therefore, they use 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. This allows us to estimate the impact of treatment without a control group. [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.³ An overview of the work focusing on 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. [Leroutier \(2019\)](#) estimates the impact of the UK carbon price floor on CO₂ emissions using a synthetic control group method which relies on constructing a “no-policy” counterfactual UK power sector from a combination of other European countries.⁴ 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 (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

³The likely reason for the scarce literature is that it is difficult to find a good “no-policy” counterfactual. This is particularly true for the power sector where carbon pricing policies typically cover almost all carbon-emitting installations ([Leroutier, 2019](#)).

⁴[Fowlie, 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 nonattainment areas.

mechanism as a change in gas 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 methodological framework to estimate the treatment effects 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. Conceptual Framework

A. Overview

We begin by providing a conceptual description of our proposed framework to estimate the causal effects of a policy intervention when a suitable control group does not exist and when treatment intensity varies over time but not across treated units. The framework comprises three major steps:

- Step 1: Deriving a structural model of the observed outcomes based on economic theory which is invariant to the policy intervention;
- Step 2: Using machine-learning (ML) techniques to train the causal model, i.e. to estimate a predictor of outcomes based on the causal model;
- Step 3: Estimating the treatment effect as the difference between predicted outcomes under observed and counterfactual values of the policy intervention (while holding other controls constant).

Before turning to a detailed description of each of the three steps, two general features of our proposed method are important to emphasize. First, it does not rely on the existence of multiple units. The estimation of the predictor function requires sufficient data for a single unit. As a consequence, we derive a time-unit specific treatment effect which can also be computed for a single unit. Second, by relying on ML techniques we allow for non-parametric predictors and therefore for the non-parametric identification of treatment effects.

B. The Causal Model

Consider a population model according to which the outcome y_{it} of unit i in period t is generated according to

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

where z_t is the treatment received by all units at time t . x_{it} and h_{it} are vectors of observed and unobserved control variables, respectively. ϵ_{it} is a random noise which is distributed with zero mean, $\mathbb{E}[\epsilon_{it}] = 0$ and variance σ_ϵ^2 , $\epsilon_{it} \sim (0, \sigma_\epsilon^2)$. ϵ_{it} is independent of controls and treatment:

$$(2) \quad \epsilon_{it} \perp\!\!\!\perp (x_{it}, h_{it}, z_t) \quad \forall i, t.$$

For each unit i , we observe a sample of outcomes $Y_i := (y_{i1}, y_{i2}, \dots, y_{iT})^\top$ and control variables $X_i := (x_{i1}, x_{i2}, \dots, x_{iT})^\top$ of size T , where T is the number of time periods. While outcomes and controls are observed at the unit level, observed treatment levels are uniform across the population, i.e. we only observe the sample of treatment levels $Z := (z_1, z_2, \dots, z_T)^\top$ —as is, for example, the case for an environmental tax which is levied equally on all units in the market.

We are interested in identifying the causal effect on outcome which is induced by a change in the treatment level from its observed value z_t to a specific value \bar{z}_t . To derive the effect of a change in the treatment variables z , we make use of an important assumption on the data generating process given by equation (1): the function f_i is invariant to changes in the treatment and control variables (Peters, Bühlmann and Meinshausen, 2016) or, put differently, f_i is assumed to be autonomous (Haavelmo, 1944; Aldrich, 1989). Given the invariance property, we are able to change the treatment variable and use the autonomous process to calculate the outcome under the changed treatment. The treatment effect is then defined as the difference between observed outcomes y_{it} , which realized under observed treatment levels z_t , and counterfactual outcomes $y_{it}^{\bar{z}}$ under hypothetical treatment levels \bar{z}_t :

$$(3) \quad \delta_{it}^{\bar{z}} := y_{it} - y_{it}^{\bar{z}} \quad \forall i, t$$

The *fundamental problem of causal inference* (Holland, 1986), often also referred to as the *missing data problem* (Rubin, 1974), is that we do not observe $y_{it}^{\bar{z}}$ and hence cannot directly calculate the treatment effect. If the treatment level varies

across units, for example, matching or difference-in-differences (DiD) methods have been put forward to solve this problem by exploiting the existence of treated and untreated, i.e. control, units. The fundamental challenge of our policy evaluation problem is, however, that the treatment (i.e., the carbon tax in the electricity sector) is uniform across the entire population (i.e., it is imposed equally on *all* power plants in the market). We are thus charged with the problem of finding a way to estimate the causal effect of the policy intervention without the possibility of relying on an untreated control group.

We propose to overcome the *missing data problem* by making use of counterfactual simulation which can create the unobserved outcomes $y_{it}^{\bar{z}}$. The main idea of our proposed approach is to predict the outcome for a counterfactual level of the treatment using a causal model for which we can change, i.e. control, the treatment variable. Pearl (2009) conceptualizes such a counterfactual intervention based on a causal model by his *do()* operator. Given the possibility to perform *do()*-interventions, we can re-write the treatment effect as:

$$(4) \quad \delta_{it}^{\bar{z}} := y_{it} - f_i(x_{it}, h_{it}, do(z_t = \bar{z}_t)) - \epsilon_{it} \quad \forall i, t.$$

In order to calculate counterfactual outcomes, the following two assumptions concerning the interaction between controls and the treatment variable have to be satisfied:

ASSUMPTION 1: *Observed controls are independent of the changes in the treatment variable: $x_{it} \perp\!\!\!\perp z_t$.*

ASSUMPTION 2: *Unobserved controls are conditionally independent to changes in the treatment variable given the observed controls: $h_{it} \perp\!\!\!\perp z_t | x_{it}$.*

Assumption 1 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 our estimate of the treatment effect.

Assumption 2 rules out effects of the treatment variable on unobserved variables after controlling for the observed variables. Again, if z would influence h , there would be an indirect effect on the outcome. It is important to note that Assumption 2 does not rule out an effect of unobserved controls. It only implies that once we include all observed controls into the model, the impact of unobserved variables is independent of the treatment level, and, thus, a change in the treatment does not affect the outcome indirectly by changing unobserved variables.

C. 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 *out-of-sample* predictions. We harness the power of ML methods which—in contrast to traditional econometric methods focused on consistently estimating *in-sample* parameters of f —are optimized to predict the value of the outcome variable (Mullainathan and Spiess, 2017).

Out-of-sample optimization is typically achieved by minimizing 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](#)):

$$(5) \quad \mathbb{E}[\text{MSE}_i] = \mathbb{E}[(y_i - \hat{f}_i)^2] = \sigma_\epsilon^2 + \underbrace{\left(\mathbb{E}[\hat{f}_i] - f_i\right)^2}_{= \text{Bias}^2(\hat{f}_i)} + \underbrace{\mathbb{E}\left[\left(\mathbb{E}[\hat{f}_i] - \hat{f}_i\right)^2\right]}_{= \text{Variance}(\hat{f}_i)}.$$

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:

$$(6) \quad \hat{f}_i^\alpha := \arg \min_{f_i \in \mathcal{F}} \sum_t \left[\left(y_{it} - \hat{f}_i^\alpha(x_{it}, z_t) \right) \right]^2$$

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 (6) 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 $\xi(x_{it}, h_{it}, z_t)$:

$$(7) \quad y_{it} = \hat{f}_i^{\alpha^*}(x_{it}, z_t) + \underbrace{f_i(x_{it}, h_{it}, z_t) - \hat{f}_i^{\alpha^*}(x_{it}, z_t)}_{=: \xi(x_{it}, h_{it}, z_t)} + \epsilon_{it}.$$

D. Estimation of Treatment Effects through Counterfactual Simulation

In the last step, we use $\hat{f}_i^{\alpha*}$ to predict the missing outcome under a counterfactual level of the treatment. A simple estimator of the treatment effect would then compare *observed outcomes* under treatment with *predicted outcomes* without treatment as suggested by equation (4). Doing so would, however, result in biased estimates due to the prediction error shown in equation (7).⁵ To estimate the treatment effect, we therefore need to eliminate the prediction error. This requires a further assumption:

ASSUMPTION 3: *The prediction error $\xi(x_{it}, h_{it}, z_t)$ is independent of the treatment:*

$$\xi(x_{it}, h_{it}, z_t^0) = \xi(x_{it}, h_{it}, z_t^1) = \xi(x_{it}, h_{it}) \quad \forall z_t^0, z_t^1.$$

Assumption 3 implies that the prediction error only depends on observed and unobserved variables, but does not change between the prediction of observed and counterfactual outcomes. Consequently, it allows to estimate the treatment effect $\hat{\delta}_{it}^{\bar{z}}$ as the difference between the *predicted values of observed outcomes* and the *predicted values of counterfactual outcomes*:

$$\begin{aligned} (8) \quad \hat{\delta}_{it}^{\bar{z}} &= \underbrace{\hat{f}_i^{\alpha*}(x_{it}, z_t)}_{\text{Prediction based on observed treatment}} - \underbrace{\hat{f}_i^{\alpha*}(x_{it}, do(z_t = \bar{z}_t))}_{\text{Prediction based on counterfactual treatment}} \\ &= y_{it} - \xi(x_{it}, h_{it}) - \epsilon_{it} - \left[y_{it}^{\bar{z}} - \xi(x_{it}, h_{it}) - \epsilon_{it}^{\bar{z}} \right] \\ &= y_{it} - y_{it}^{\bar{z}} + \phi_{it}, \end{aligned}$$

where $\phi_{it} := \epsilon_{it}^{\bar{z}} - \epsilon_{it}$ is random noise with mean zero. As we only change the treatment variable and as observed and unobserved variables are independent of the treatment (Assumptions 1 and 2), Assumption 3 allows us to eliminate the prediction bias and the impact of unobserved variables in the estimation, and hence the identification of the treatment effect. Assumption 3 is analogous to the parallel trend assumption in a DiD setting.⁶

A potential concern for Assumption 3 and therefore the estimation of $\hat{\delta}_{it}^{\bar{z}}$ is the quality of predictions based on *unobserved* counterfactual values. To ensure a valid prediction, two additional assumptions regarding the data need to be satisfied.

⁵Using the definition of the treatment effect (3) and equation (7), an estimator comparing the observed values y_{it} with a predicted counterfactual value yield: $\hat{\delta}_{it}^{\bar{z}} = y_{it} - \hat{f}_i^{\alpha*}(x_{it}, z_t = \bar{z}_t) = y_{it} - y_{it}^{\bar{z}} + \xi(x_{it}, h_{it}, z_t = \bar{z}_t) + \epsilon_{it}$. This estimator would thus be biased by the prediction error.

⁶Assumption 3 is one of the main differences in comparison with the approach used by Burlig et al. (2017) and Cicala (2017). They assume that "...treated and untreated schools [are trending] similar on prediction errors..." (Burlig et al., 2017, pp. 18) or, likewise, "Parallel trends in unobservables..." (Cicala, 2017, Assumption 2. p. 23) in the sense that the "contemporaneous error", i.e. the prediction error, is behaving similar across regions. Given these assumptions, they are able to differentiate out the prediction error and impact of unobservables by using control groups in a DiD setting. In contrast, and as we do not observe a control group, we need to assume that the prediction error is independent of the treatment to differentiate out the impact of unobservables and systematic prediction errors.

First, although ML algorithms are designed to produce reliable out-of-sample predictions, they only locally approximate the true model in the range of observed treatments and covariates. It is thus unclear how the estimated functions behave for covariate and treatment combinations which lie outside of the range of observed combinations. To rule these cases out, we need the *positivity* or *covariate overlap* assumption (Samii, Paler and Daly, 2016):

ASSUMPTION 4: *Each combination of the counterfactual treatment z and covariate level X has been observed (i.e., $Pr[z|X] > 0$).*

While it is highly unlikely that all combinations of z and X have been observed, Assumption 4 requires that these combinations should lie within the range of observed data.

The last assumption concerns the variation in the level of treatment and controls which is needed to estimate a valid predictor of the underlying structural process:

ASSUMPTION 5: *The variation in the level of treatment and controls over time is sufficiently large.*

Assumption 5 implies that the impact of a change in the treatment on the outcome can be predicted. For many policy interventions, however, treatments are discrete and do only change infrequently. A possible remedy is to find a control variable which affects outcome through the same causal mechanism as the treatment variable. In fact, changing such a control variable implies the same change in outcome as a change in the treatment variable itself. For example, an archetypical problem in economics is to estimate the impact of imposing an input tax (e.g., a carbon tax). Here, the tax change may be a one-time event or it may comprise only a few discrete tax changes. The impact of the tax on input costs follows, however, the same mechanism as a change in input prices. It is thus possible to use the variation in input prices to identify the causal mechanism of the input tax.

II. Applying the Framework to Evaluate Climate Policy: The Case of the UK Carbon Tax

We apply the proposed framework to assess the market impacts of a carbon tax policy using the case of the UK carbon tax. In Sections II.A and II.B we provide information about the policy background and draw on economic theory to derive the causal model (Step 1) Section II.C presents the data sources and construction. Section II.D describes our empirical framework to estimate the treatment effect (Steps 2 and 3). Section III scrutinizes the validity of our identifying assumptions within the context of our empirical application.

A. The Policy Intervention 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

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 (€ per ton)	–	–	–	–	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>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 (15).

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.

To develop some first intuition for the impacts of the CPS on electricity supply and emissions, Figure 1 plots the short-run supply curve (i.e., ordering marginal

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.

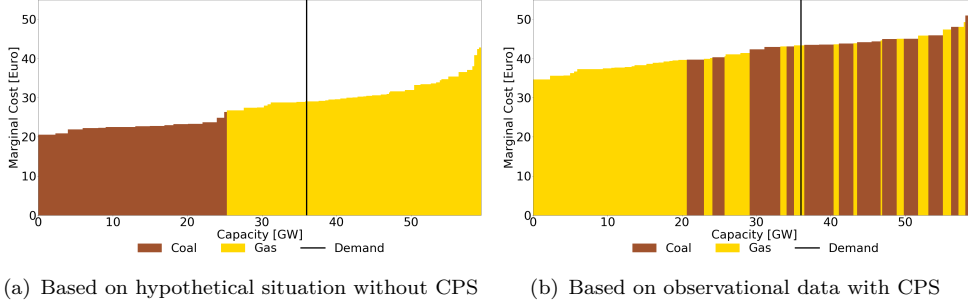


FIGURE 1. Illustrative impact of the UK carbon tax on the short-run market supply curve for electricity



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 (21).

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 2 shows that starting with the introduction of the CPS in 2013 the annual market share of coal-fired generation sharply decreased while the share of gas-fired plants increased; over the same period, UK’s electricity-sector emissions sharply declined.

While Figures 1 and 2 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) negative macro-economic shocks and 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

⁸The illustrative calculation shown in the figure 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.

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

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

FIGURE 2. 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 \(2007\)](#). 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*.

by nearly 50 percent while coal prices remained largely constant (see Figure 2 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” (LCPD), which sets specific limits on local pollutant emissions for power plants constructed after the year 1987. The LCPD 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 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

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.

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} :

$$(9) \quad K_{it} \geq X_{it} \quad \perp \quad \mu_{it} \geq 0 \quad \forall i, t$$

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

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 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} :

$$(10) \quad c_{it}(\boldsymbol{\vartheta}_{it}) + \mu_{it} \geq P_t \quad \perp \quad y_{it} \geq 0 \quad \forall i, t$$

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:

$$(11) \quad \sum_i y_{it} = D_t \quad \perp \quad P_t \text{ "free"} \quad \forall t.$$

Equations (9)–(11) 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

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).

¹¹Equation (10) determines the price as the marginal cost of the marginal generator, i.e. the generation that earns zero capacity rent in the given hour ($\mu_{it} = 0$).

on demand, and its own as well as the marginal cost and available capacities of all other plants (indicated by $-i$):

$$(12) \quad y_{it}^* = \mathcal{F}_{it} \left(D_t, c_{it}(\boldsymbol{\vartheta}_{it}), K_{it}, c_{(-i)t}(\boldsymbol{\vartheta}_{(-i)t}), K_{(-i)t} \right) .$$

Equation (12) 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.

GRAPHICAL REPRESENTATION OF THE STRUCTURAL MARKET MODEL.—We make use of directed acyclic graphs (DAGs) to graphically illustrate the causal relationships between variables of the structural electricity market model summarized in equation (12) and to briefly describe how one can obtain treatment effects.

In a DAG, model variables are represented by nodes. If an arrow runs from node A to node B, A is called a parent of B. A graph is called acyclic, if there is no series of arrows which starts and ends in the same variable. A causal DAG associates “parent” with “direct cause”, with an arrow indicating the directed causal relation between two variables. The important insight from Pearl (2009) is that in the absence of any indirect paths, one can use the do-operator to perform a hypothetical intervention on the treatment variable while holding all other controls constant. This implies that it is possible to use the structural model to obtain the unobserved counterfactual outcome of “no policy” by setting the treatment variable to zero. The treatment effect can then be derived as the difference between outcomes with and without policy intervention.

Figure 3 uses two DAGs to depict the structural electricity market model in (12). The left-hand DAG shows treatment (p^{CPS}) and control variables determining the marginal cost of plant i (c_i). Consistent with equation (10), the plant-specific heat efficiency together with fuel prices, carbon prices, and the fuel-specific carbon intensity determine c_i . The right-hand DAG then shows how marginal cost and capacities of all plants determine aggregate supply S and, together with demand D , the equilibrium market output y_i^* of power plant i . This representation implies that the central model (12) should not be viewed as portraying the output decision of a single plant; the dependence of the equilibrium quantities on own and other generators' marginal cost and demand can also be understood in terms of bid functions on the market level. Under perfect competition each generator bids the whole capacity at marginal cost into the market, leading to the supply curve S . The market operator then chooses the cheapest bids until demand is fulfilled. Consequently, the acceptance of a bid depends on the ordering of marginal cost in the entire market as well as available capacities and demand.

Figure 3 also visualizes Assumptions 1 and 2: the only path from the treatment variable p^{CPS} to outcome y_i^* is through marginal cost and supply. In particular, the observed controls (i.e., white nodes) are independent of the changes in the treatment variable; and unobserved controls are conditionally independent to changes in the treatment variable given the observed controls (i.e., there are no effects of the treatment variable on unobserved variables after controlling for the observed




FIGURE 3. Hypothesized directed causal relations in the structural model of the electricity market

Notes: The treatment variable (the UK carbon tax, p^{CPS}) and outcome variable (power plant output, y_i^*) are shaded in dark grey. White-shaded variables represent exogenous controls. Light grey-shaded variables denote derived variables. An arrow represents a direct cause. The left graph shows how treatment and control variables affect the marginal cost of a generator. The right graph shows how marginal cost and available capacity of all generators as well as aggregate market demand determine the output of a generator in a given hour. Time subscript t is dropped for the ease of notation.

variables).

To summarize, up to this point we have simply used economic theory to derive a structural model for equilibrium outputs of power plants. In order to perform counterfactual analysis and identify the causal impact of the UK carbon tax, we need to ensure that the relationships which are asserted in Figure 3 are valid. Before turning to this issue in Section III, we next present our data and empirical framework.

C. Data Sources and Construction

To obtain measurements for the empirical counterparts of all RHS variables in (12), we combine data from different (and publicly available) sources. We use panel data of hourly generation for each UK fossil-fuel power plant covering the 2009-2016 period. In addition, we use data on available hourly capacity, technical characteristics of each plant, non-fossil generation, demand, daily fuel and carbon prices, and weather data.

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 balancing power or 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.

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 \(2007\)](#). 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 \(2017\)](#).

FUEL PRICES (p_t^{fuel}).—Data on daily fuel prices for coal and natural gas are taken from [EIKON \(2007\)](#). 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 \(2017\)](#).

EMISSIONS FACTORS AND PLANT-SPECIFIC HEAT EFFICIENCIES (θ^f and 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 EUTL ([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:

$$(13) \quad \eta_i = \theta^f / e_i.$$

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.

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 output that a plant can provide 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: 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.

Not 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.

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.

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

D. The Empirical Framework.

We now turn to the implementation of our conceptual framework established in Section I within the context of the UK CPS. First, we use the information about electricity markets to derive a model of observed outcomes (i.e., generation of each plant), which is invariant to the policy intervention (i.e., the CPS). Second, we use available data and ML algorithms to estimate a predictor of outcomes (i.e., predicted generation). Third, we use the prediction model to estimate the treatment effect as the difference between predicted generation with treatment (i.e. observed CPS) and the unobserved counterfactual (i.e., generation without CPS). Fourth, we present our ex-post calculations to get from the impact on generation to the impact on emissions and abatement cost. Finally, in the subsequent section, we discuss the validity of our approach addressing our four main assumptions.

ESTIMATION EQUATION (STEP 1).—Based on the electricity market model in Section II.B, we know that the equilibrium outcome of each plant i depends on demand D_t , and its own as well as the marginal cost and available capacities K_{it} of all other plants. As we do not directly observe plants’ (or generators’) marginal costs and

heat efficiencies (η_{it}), we exploit the fact that they depend on ambient temperature and thus additionally include daily mean temperature ($temp_t$). The empirical analogue of (12) then becomes:

$$(14) \quad y_{it} = f_i \left[r_t \left(p_t^{coal}, p_t^{gas}, \theta^f, p_t^{EUA}, p_t^{CPS} \right), temp_t, D_t, K_{it}, K_{(-i)t}, \Phi_t \right] + \epsilon_{it},$$

where we 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; and the carbon price inclusive ratio of relative fuel prices:

$$(15) \quad r_t := \frac{p_t^{coal} + \theta^{coal} (p_t^{EUA} + p_t^{CPS})}{p_t^{gas} + \theta^{gas} (p_t^{EUA} + p_t^{CPS})}.$$

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}). Moreover, using r_t in equation (14) nicely concords with the view that it is not the absolute but the relative fuel prices determining which plants exit or stay in the market.

MACHINE LEARNING ALGORITHM (STEP 2).—While we know from the theoretical electricity model in (12) and its empirical counterpart in (14) which variables affect plants' output decisions, we do not know the functional form of f_i . To obtain an estimator \hat{f}_i of the function f_i , we therefore apply ML algorithms, which allow for flexible functional forms, to produce reliable out-of-sample predictions of each plants' output, y_{it} .

We employ the LASSO¹² algorithm (Tibshirani, 1996)—a penalized linear regression model—and use k-fold cross-validation dividing the sample into eight groups (often called folds) to train a prediction model $\hat{f}_i^{\alpha^*}$ for each plant individually.¹³ Each prediction model consists of the set of coefficients $\hat{\beta}^{\alpha^*}$ and the optimal regularization parameter α^* , which lead to the best possible out-of-sample prediction.¹⁴

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

¹²We also used other algorithms such as random forest. However, the LASSO, an algorithm which is linear in coefficients, lead to the most convincing simulation of the electricity market as a whole: While other algorithms failed at implicitly fulfilling the market clearing condition (see equation (11)), the LASSO algorithm was able to meet this condition—although it was not explicitly modelled (see also III).

¹³The LASSO algorithm requires a pre-defined set of input features. In addition to the variables which appear on the RHS of (14), 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.

¹⁴Appendix B assesses the out-of-sample 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 broader insight 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:

$$(16) \quad \bar{r}_t := \frac{p_t^{coal} + \theta^{coal} p_t^{EUA}}{p_t^{gas} + \theta^{gas} p_t^{EUA}} .$$

Based on the estimator in equation (8) detailed in Section I, the impact of the CPS on the output decision of each plant i in each hour t can then be calculated as:

$$(17) \quad \hat{\delta}_{it}^{CPS} = \hat{y}_{it}^{\text{with CPS}} - \hat{y}_{it}^{\text{without CPS}} ,$$

where

$$(18) \quad \hat{y}_{it}^{\text{with CPS}} = \hat{f}_i^{\alpha*} (r_t, temp_t, D_t, K_{it}, K_{(-i)t}, \Phi_{it})$$

$$(19) \quad \hat{y}_{it}^{\text{without CPS}} = \hat{f}_i^{\alpha*} (do(z_t = \bar{z}_t), temp_t, D_t, K_{it}, K_{(-i)t}, \Phi_{it}) .$$

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 (with sample size $N=1000$) 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 4.

MEASURING CO₂ EMISSIONS AND ABATEMENT COST.—To calculate electricity-sector emissions (from combustion of coal and natural gas in electricity generation) at time t , we aggregate CO₂ emissions from all plants operating in the market:

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

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:

$$(20) \quad \Delta E_t := \sum_i \underbrace{e_i \hat{\delta}_{it}^{CPS}}_{\text{Policy-induced change in emissions of plant } i (=:\Delta E_{it})} .$$

Next to its impact on generation and consequently emissions, the CPS also leads to a change in aggregate production costs. 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 (13) and shown in Table 2) marginal cost are

calculated as

$$(21) \quad c_{it}(\boldsymbol{\vartheta}_{it}) = \frac{1}{\eta_{it}} \left(p_t^f + \theta^f (p_t^{EUA} + p_t^{CPS}) \right).$$

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

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

Using the definition of the treatment effect from equation (17) and plant-specific heat efficiency from equation (13), this can be rewritten as follows:

$$(22) \quad \Psi_t = \underbrace{\sum_i \hat{\delta}_{it}^{CPS} \frac{1}{\eta_{it}} \left(p_t^f + \theta^f p_t^{EUA} \right)}_{\substack{=:T_t \\ \text{Technical abatement} \\ \text{cost}}} + \underbrace{\sum_i p_t^{CPS} e_i \hat{y}_{it}^{\text{with CPS}}}_{\substack{=:R_t \\ \text{Tax payments} \\ \text{due to CPS}}}.$$

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 Underlying Assumptions

As established in Section I.B, the validity of our approach to estimate the treatment effect of the policy intervention as the simple difference between predicted outcomes with and without the policy intervention relies on the existence of a causal model, f_i , and a set of assumptions. A general caveat applies here: given observational data, it is impossible to fully verify the validity and completeness of a causal model. There are, however, certain aspects and assumptions of the model which can be examined given the specific empirical context of the application.

EXISTENCE OF f_i .— We assume that the function f_i (see equation 1) did not change over time and is independent of the treatment, i.e. the level of the carbon price. This assumption is supported by the fact that during our sample period there were no major institutional changes concerning the market operation of the UK power market.

STABLE UNIT VALUE TREATMENT.—By modeling the output of one plant as a function of the characteristics of all other plants in the market, we do not model the output decision of single plant but rather the decision of the market maker which plants to use. As a result, the output function is independent of the output function of the other plants, implying that the stable unit value treatment assumption is fulfilled.¹⁵

INDEPENDENCE OF CONTROL VARIABLES FROM TREATMENT (ASSUMPTION 1).—We argue that it is safe to assume that control variables are independent from the level of the CPS based on the following six observations. First, the CPS level is determined exogenously at a fixed rate two years in advance. Second, the EUA carbon price is determined by the EU ETS market of which the UK electricity sector only covers a negligibly small part. Third, the market share of UK’s electricity firms on international fuel markets is not large enough to affect fuel prices. Fourth, the short-run nature of our analysis means that electricity demand does not react to hourly wholesale electricity prices which may be impacted by a carbon tax. Fifth, likewise installed capacities cannot be adjusted in the short run and are thus not impacted by a carbon price. Sixth, temperature—which we use as a proxy for unobserved time-dependent heat efficiencies—is determined by exogenous weather conditions which are entirely independent of carbon tax policy.

CONDITIONAL INDEPENDENCE OF UNOBSERVED CONTROLS FROM TREATMENT (ASSUMPTION 2).—Assumption 2 cannot be tested directly because one does not know which unobserved variables may influence plants’ output decisions and whether or not they are affected by the level of the CPS. One (imperfect) remedy is to test for the robustness of our model using a variety of different fixed-effects specifications.

Table 3 reports the impact of the CPS on coal and gas power plant generation from four different model specifications. M1 includes monthly and hourly fixed effects while M2-M4 exclude either monthly or hourly dummies or both.¹⁶ Our finding that the results are robust across model specifications M1-M4 suggests that there do not seem to be significant unobserved variables, with systematic variation at the monthly and/or hourly level, that impact plants’ output decisions. We interpret this as evidence that Assumption 2 is plausibly satisfied in the context of our application.

Table 3 also bears out another important insight: in all specifications the total net impact of the CPS on generation, i.e. the sum of the impacts on coal and natural gas, does not statistically differ from zero (at a 5 % significance level). This also holds true on a monthly basis (see Figure 5). A priori, this result is not to be expected as we (1) separately estimate each plant’s output decision and (2) do not impose an *explicit* market clearing constraint in our empirical model—unlike the theoretical model which hypothesizes that demand always equals supply according to equation (11). We interpret the statistical rejection of a violation of *implicit*

¹⁵In contrast, if we modeled plant output solely depending on own marginal cost, i.e. by a power plant’s bid function, output of a given plant would depend on the treatment applied to other plants as it changes marginal cost and hence the ordering of the plants in the supply function.

¹⁶Note that estimating a predictor for power generation at the plant level does not allow the inclusion of plant-specific dummies.

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 $\hat{\delta}_{it}^{CPS}$ based on equation (17). ^aRefers to situation without the CPS. Bootstrapped standard errors are shown in parentheses.

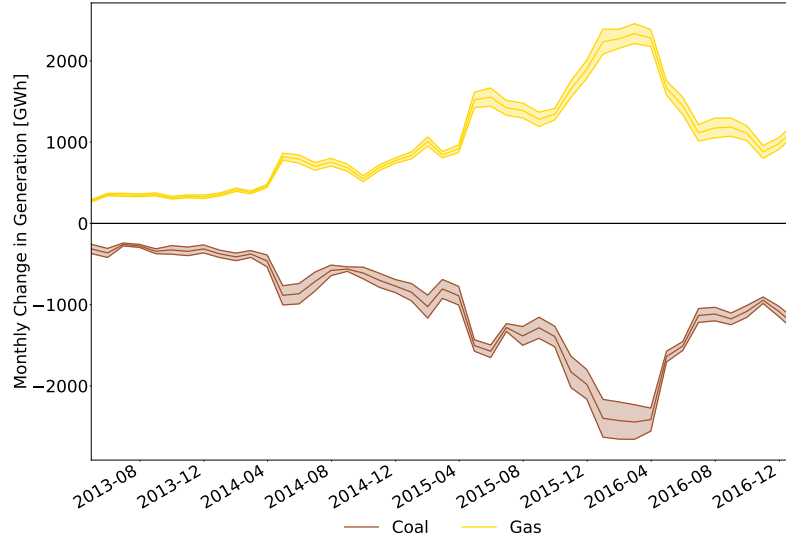


FIGURE 5. Monthly average impacts of the UK carbon tax (CPS) on electricity output by technology

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 (17), aggregated by technology category and month.

market clearing as additional evidence that our model is correctly specified. In particular, it suggests that we are not missing any unobserved variables that affect power plant performance and depend on treatment. For the subsequent analysis in this paper, we take M1 as our preferred model specification.

POSITIVITY OR CO-VARIATE OVERLAP (ASSUMPTION 4).—To ensure a high prediction

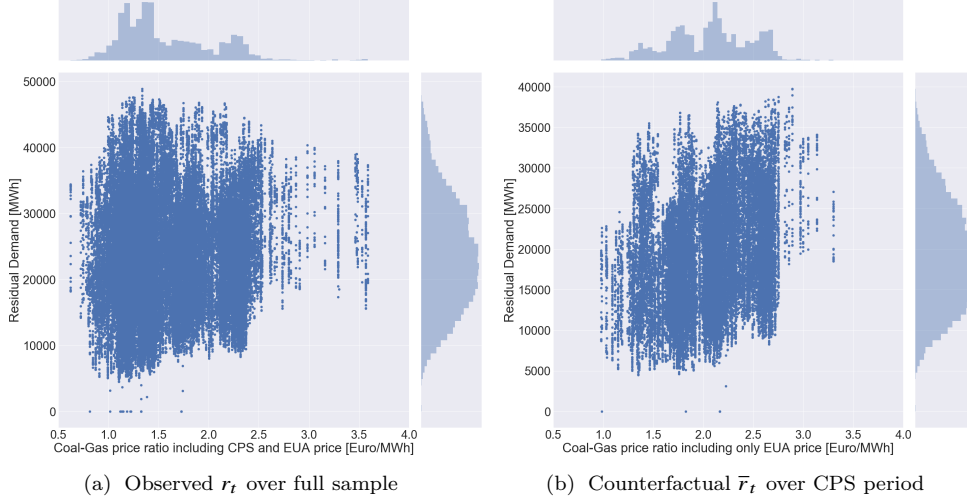


FIGURE 6. Joint distribution of electricity demand and the observed (r_t) and counterfactual (\bar{r}_t) coal-to-gas fuel price ratio

quality of the counterfactual simulation, the positivity assumption requires that the counterfactual fuel price ratio \bar{r}_{it} lies within the range of observed fuel price ratios r_t conditional on observed control variables. Apart from marginal cost, which are expressed through the fuel price ratio, residual demand is the main control variable determining which plant enters the market.

Figure 6 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 (only a small fraction of values with $r_t > 2.5$ fall outside the observed sample distribution).

VARIATION IN TREATMENT VARIABLE (ASSUMPTION 5).— To train the model and estimate a valid predictor, sufficient variation in treatment and control variables is necessary. While the CPS only varies on an annual level, the distribution of the carbon price-inclusive fuel price ratio depicted in Figure 6 shows that there is substantial variation in our modified treatment variable over the sample period.

PREDICTION ERROR IS INDEPENDENT OF TREATMENT (ASSUMPTION 3).— Since the “no-policy” counterfactual cannot be observed, one cannot assess the prediction error of the machine-learned “no-policy” counterfactual during the treatment period. We believe, however, that the following three arguments provide support for the conclusion that Assumption 3 holds. First, we can evaluate to what extent the *observed* prediction error depends on the level of treatment. We compute the correlation between r and the *observed* prediction error for each plant and find that the mean correlation coefficient (over all plants) is fairly low 0.01 (with a standard deviation of 0.07). Second, we observe that our counterfactual levels of

r lie within the range of the observed levels (see validity of Assumption 4). Given the low correlation between the *observed* prediction error and r , it seems plausibly to assume that prediction error is also independent from the treatment for counterfactual levels of r . Third, related to the discussion of Assumption 2, we argued that it seems plausible to assume that we do not miss important unobserved variables which could affect the errors differently for the observed and counterfactual predictions in the treatment period. Overall, we are thus confident that the errors of the observed and the “no-policy” counterfactual predictions can be assumed to be independent of the CPS level.

IV. Causal Impacts of the UK Carbon Tax

This section presents our treatment effect estimates of the UK CPS between 2013-2016 using the preferred model specification M1 (see Table 3). Our outcome of interest is the change in plant-level electricity output across different time periods based on which we can calculate tax-induced impacts in terms of aggregate market output by technology, CO₂ emissions, abatement costs, and tax revenues collected.

A. Plant-level and Aggregate Electricity Output

POWER PLANT IMPACTS.—Figure 7 plots monthly electricity output between 2012-2016 for the largest coal and gas-fired power plants in our sample for observational data as well as model-predicted values with the CPS and under the unobserved counterfactual without a CPS.

First, it provides a visualization of the main idea of our approach summarized by the model in (17) to estimate plant-level treatment effects over time: given the ML-trained causal model of plant-specific electricity output, we obtain the causal impact of the carbon tax policy by taking the vertical distance between model-predicted outcomes with the CPS ($\hat{y}_{it}^{\text{with CPS}}$, filled blue dots) and without the CPS ($\hat{y}_{it}^{\text{without CPS}}$, red dots).¹⁷ To the extent that Assumptions 1–5 can be taken as plausibly satisfied (see Section III), $\hat{\delta}_{it}^{\text{CPS}} = \hat{y}_{it}^{\text{with CPS}} - \hat{y}_{it}^{\text{without CPS}}$ identifies plant-level treatment effects.

Second, it becomes evident that the introduction of the UK CPS led to a decrease in coal- and an increase in gas-fired electricity generation. We observe the same pattern of output changes for the other power plants (see Table A3 in the Appendix A).

AGGREGATE IMPACTS.—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; relative to a situation without the CPS, these changes correspond to a fuel switch from coal to natural gas of around 15 percent.

¹⁷Comparing model-predicted values with a CPS policy (filled blue dots) to observational data (hollow blue dots), Figure 7 also shows that the model accurately predicts the observational data.

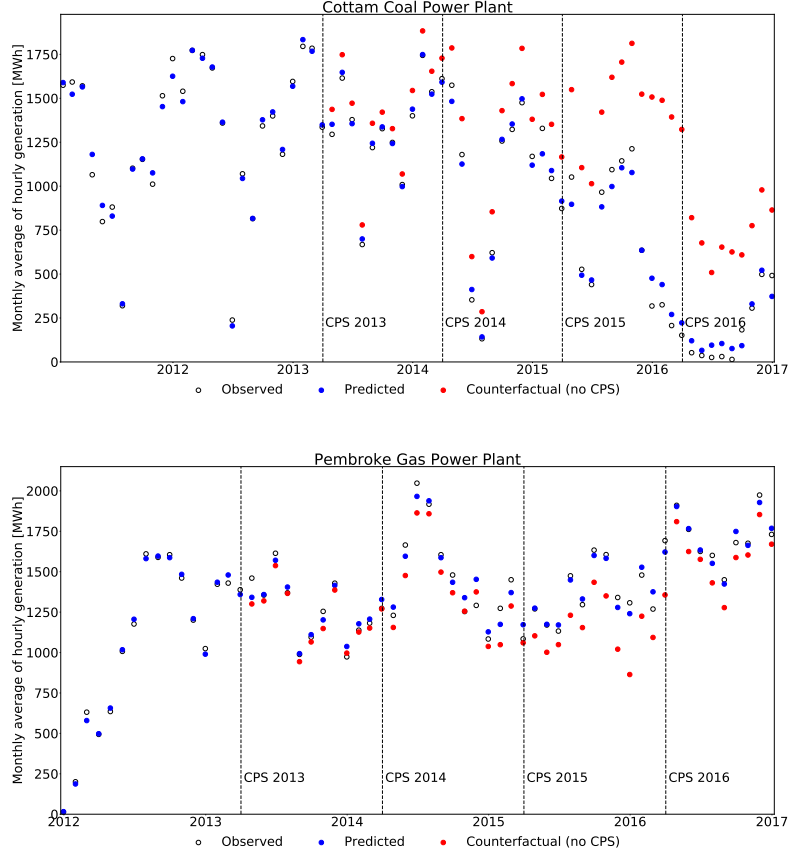


FIGURE 7. Electricity output (monthly averages, 2012-2017) for the largest coal plant (Cottam, upper row) and largest gas plant (Pembroke, lower row): observations versus model-predicted with CPS ($\hat{y}_{it}^{\text{with CPS}}$) and without CPS ($\hat{y}_{it}^{\text{without CPS}}$).

The impact of the CPS on generation varies substantially over time. The fuel switch was initially low at an absolute level of around 4 TWh in the 2013 period and then increased over the years with the highest value of around 22 TWh in 2015. The absolute 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. Adding to the heterogeneity in the annually aggregated impacts by CPS period, Figure 5 shows that there is also considerable variation over time in the CPS-induced impacts on monthly output of coal- and gas-fired power plants.

B. CO_2 Emissions and Abatement Cost

Table 5 summarizes the total and yearly impacts of the CPS on electricity-sector CO_2 emissions and abatement cost. Using equation (20), we estimate that over the period 2013–2016, the CPS has reduced cumulative emissions by 26.1 million

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 (0.27)	-9.26 (0.57)	-21.92 (0.86)	-10.94 (0.21)	-46.29 (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 (0.10)	9.37 (0.23)	21.19 (0.57)	10.72 (0.40)	45.55 (1.06)
% of total generation	6.1	12.1	29.7	12.8	15.0
<i>Total</i> [TWh]	0.10 (0.29)	0.11 (0.62)	-0.73 (1.03)	-0.22 (0.45)	-0.75 (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 model specification M1 and equation (17), aggregated by technology category. Bootstrapped standard errors are shown in parentheses.

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 $\Psi_t = T_t + R_t$</i>					
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 (17), 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.

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 from equation (22), the CPS has reduced one ton of CO₂ emissions at an average cost of €18.2 over this period.¹⁸

¹⁸Although not the focus of the paper, Table 5 also reports on the tax revenues collected with the CPS instrument. Since its introduction and until the end of 2016, the British government received around




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 (20) and (22).

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, within period impacts are heterogeneous, i.e. abatement quantity and costs largely vary for a given level of the carbon tax (see Figure 8). Thus, the CPS level cannot solely explain the observed variation in the impacts of the carbon tax.

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

€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; already in the subsequent period, the CPS tax revenue dropped significantly due to the fact that the CO₂ emissions remaining in the market were considerably lower.

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. What Determined the Environmental Effectiveness and Costs of the UK Carbon Tax?

This section uses simulations with the ML-trained model to investigate what drives heterogeneity in the abatement quantity and costs induced by the carbon tax, shedding light on their ambiguous empirical relationships summarized in Results 1 and 2. We also examine the extent to which the insights gained from our empirical model are compatible with first principles of microeconomic theory for cost-optimizing firm behavior in electricity markets.

A. A Simple Model of Carbon Abatement

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, i.e. the level of the carbon tax (p^{CPS}). If the change in the tax level does not explain the heterogeneity in the tax-induced impacts (see Section IV.B), drivers for the different abatement impacts of the carbon tax must be related to changes in the MAC.

Consider a simple model of carbon abatement where a representative electricity firm seeks to minimize its carbon tax-induced impact on production costs (we drop the time index for simplification):

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

The total impact on production costs $\Psi = T + R$ —in line with equation (22)—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-

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 a yields:¹⁹

$$(23) \quad \underbrace{\underbrace{\frac{\partial T(a; \bar{r})}{\partial a}}_{\text{Marginal technical abatement costs (MTAC)}} + \underbrace{\mu}_{\text{Marginal rent on abatement potential}}}_{\text{Marginal abatement costs (MAC)}} \geq \underbrace{p^{CPS}}_{\text{Marginal benefits of abatement (MAB)}} \perp a \geq 0.$$

The MAC hence 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-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.

The KKT conditions in (23) allow us to derive several hypotheses 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 ($\partial MAC / \partial a = \partial^2 T(a; \bar{r}) / \partial a^2 + \partial \mu / \partial a \geq 0$).*

Conjecture 1 describes the behavior of MAC regarding abatement. It simply states that MAC increase in abatement for a given fuel price ratio. MAC are composed of two terms, MTAC and the shadow price of abatement potential, both increasing in abatement. The next two conjectures hypothesize that these two components defining total MAC in (23) depend on the level of the fuel price ratio \bar{r} :

CONJECTURE 2: *The marginal technical abatement cost weakly decreases in the relative price of coal to natural gas ($\partial^2 T(a; \bar{r}) / \partial a \partial \bar{r} \leq 0$).*

¹⁹Here, 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: $\partial T / \partial a + \mu \geq p^{CPS}$, $a \geq 0$, $(\partial T / \partial a + \mu - p^{CPS})a = 0$. 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 $MAC = MAB$ in the optimum.

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 2 simply 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.

Assuming that Conjectures 2 and 3 hold, an important insight is that a change in the fuel price ratio has an ambiguous effect on the MAC. On the one hand, MTAC are decreasing in \bar{r} . On the other hand, μ is increasing in \bar{r} .

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.

Conjectures 1–3 provide a theory-founded explanation for the empirically observed non-linear relationships between the carbon tax level, abatement quantity, and costs which were summarized in Result 1 and Result 2. We next perform simulations with our ML-trained model to investigate to what extent Conjectures 1–3 hold up in an empirical context. 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.

B. Empirical Marginal Abatement Costs

COMPUTATIONAL DERIVATIONS.—To obtain empirical counterparts of the MAC (LHS of (23)), 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

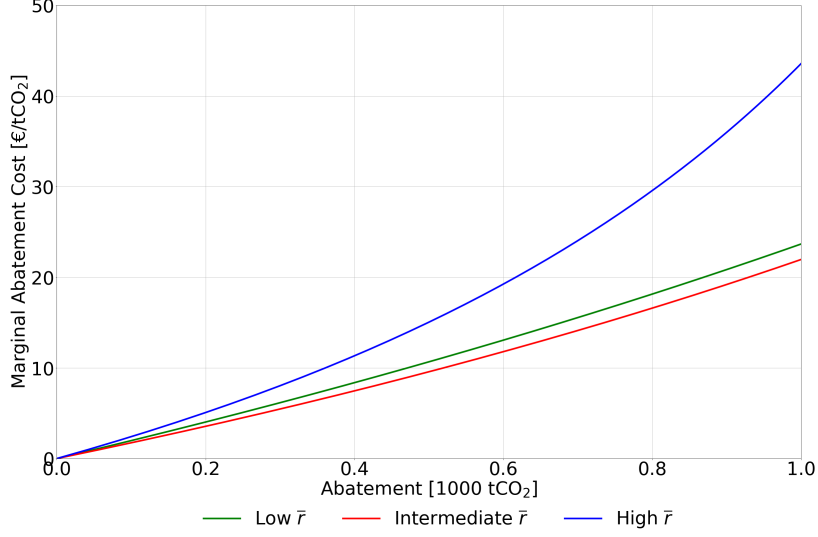


FIGURE 9. Empirical MAC curves for different carbon tax-exclusive fuel price ratios \bar{r}_t

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.

The 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).

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 (23) 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, 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

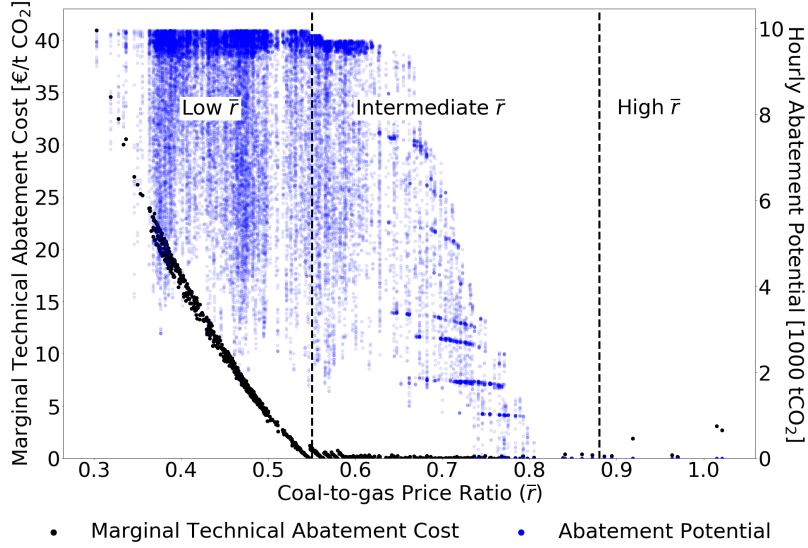


FIGURE 10. Empirical relationships between hourly MTAC, hourly abatement potential, and the fuel price ratio \bar{r}_t

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.

the empirically-measured counterparts of the MTAC and the abatement potential 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 provides empirical evidence in strong 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 High Carbon Tax Necessarily Lead to More CO₂ Reductions and Higher Costs?




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 and ask 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 carbon tax.

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 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

²⁰This is equivalent to using the ML-trained model to computationally evaluate the KKT conditions in (23) to find cost-minimizing emissions abatement a for a given carbon tax rate p^{CPS} .

(a) CO_2 abatement (ΔE_t)(b) Total technical abatement costs (T_t)(c) Average abatement costs ($T_t/\Delta E_t$)FIGURE 11. Empirical relationships between CO_2 abatement, abatement costs, and the fuel price ratio \bar{r}_t for different carbon tax rates p^{CPS}

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. 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 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 is exhausted. Taken together, 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

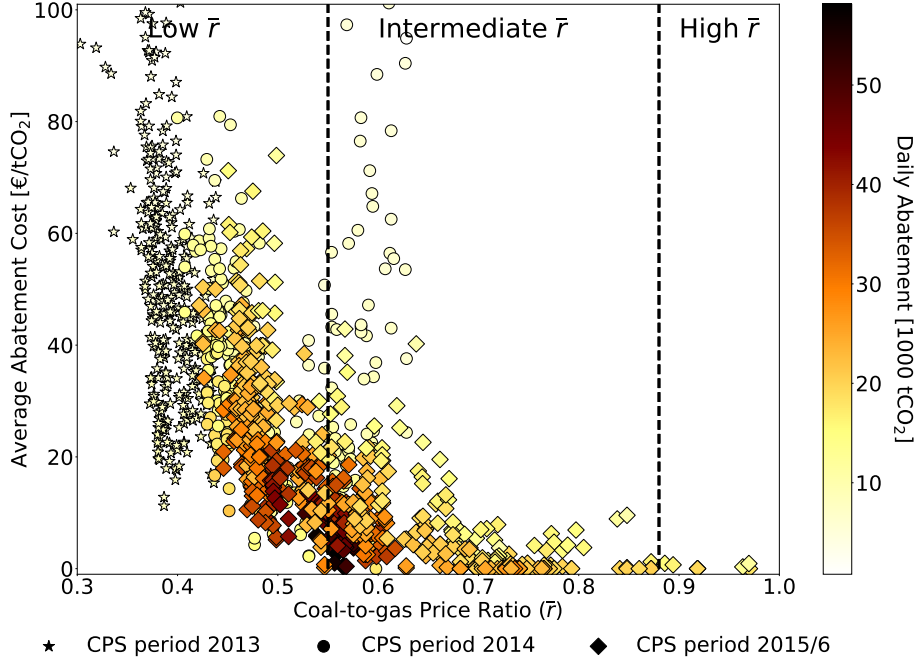


FIGURE 12. Relation between daily CO₂ abatement, and daily average technical abatement cost, and daily (average) fuel price ratio \bar{r}_t .

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 (20) and (22).

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 targeted sector—there is considerable temporal heterogeneity in abatement quantities and costs, resulting from the variation 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 arguably one of the main policy instruments for mitigating climate change, surprisingly little is known about its performance from an ex-post perspective. Causal inference of the impacts of a broad-based carbon tax, i.e. one which affects almost all CO₂-emitting units in a market, is difficult as typically no control group or counterfactual situation exists.

Against this background, this paper has made two contributions. First, we have developed and implemented a new approach which combines economic theory

and machine learning (ML) techniques to establish causal inference of a policy intervention 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 causal model which is then trained using ML to obtain an empirical model for out-of-sample prediction at the firm level. We obtain the treatment effect of a carbon tax on generation for each power plant as the difference between predicted outcomes with and without policy.

The developed framework is based on several conditions that have to hold—and which we think are plausibly satisfied in the empirical context of the UK carbon tax program evaluation. We deal with a situation in which the underlying structural causal model is constant over time. This allows us to use the full sample to train a predictor function ensuring a high prediction quality. Furthermore, the treatment variable has to be variable enough to allow to identify its causal impact in the predictor function. If treatment does not vary enough—as is the case for the CPS level in our application—one needs to exploit the variation of a control variable which exerts the same causal impact as the treatment variable. In our context of evaluating a carbon tax policy, this is the relative fuel price of coal to natural gas.

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 the best of our knowledge, this is the first paper in economics to incorporate ML methods to conduct causal inference 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 CPS 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 emerging from our analysis is that, in the short run, a higher carbon tax does not necessarily deliver higher emissions reductions; at the same time, however, a higher carbon tax need not necessarily result in higher abatement costs. When designing effective carbon tax regulation, policy makers should use as much market information as possible about current and future relative fuel prices.

Some limitations of our analysis should be kept in mind. First, focus on analyzing the short-run market impacts of the CPS. 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” (LCPD), 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 the introduction of the CPS 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.
- Burlig, Fiona, Christopher Knittel, David Rapson, Mar Reguant, and Catherine Wolfram.** 2017. "Machine learning from schools about energy efficiency." National Bureau of Economic Research.
- 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.
- Cicala, Steve.** 2017. "Imperfect Market versus Imperfect Regulation in the U.S. Electricity Generation Generation."
- 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.** 2007. "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 Global Change Report 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: European 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.
- Fowle, 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.
- 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.
- 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/>.
- 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.
- Pearl, Judea.** 2009. *Causality*. Cambridge University Press.
- 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, Mariette 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	2014	2015	2016	2013-2016
<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 (17). 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: MACHINE LEARNING (LASSO) ALGORITHM VERSUS OLS

This section compares the out-of-sample 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 (14).

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 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 B1. Comparison of the distribution of plant-specific performance scores by fuel type for LASSO vs. OLS models.

Figure B1 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 B1 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).

- 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
- 18/296 L. Bretschger and C. Karydas
Economics of Climate Change: Introducing the Basic Climate Economic (BCE) Model
- 18/295 A. Pattakou and A. Vlahakis
Effectiveness of renewable energy subsidies in a CO2 intensive electricity system
- 18/294 H. Gersbach, V. Hahn and Y. Liu
Macropprudential Policy in the New Keynesian World
- 18/293 H. Schwerin
Swap Bonds or Stocks, or Balance Trade! A Game of Implicit Environmental Policy
- 18/292 S. Houde
Bunching with the Stars: How Firms Respond to Environmental Certification
- 18/291 L. Bretschger and A. Vinogradova
Escaping Damocles' Sword: Endogenous Climate Shocks in a Growing Economy

- 18/290 S. Houde
The Incidence of Coarse Certification: Evidence from the ENERGY STAR Program
- 18/289 J. Blasch, N. Boogen, C. Daminato and M. Filippini
Empower the consumer! Energy-related financial literacy and its socioeconomic determinants
- 18/288 L. Bretschger and S. Soretz
Stranded Assets: How Policy Uncertainty affects Capital, Growth, and the Environment
- 18/287 S. Rausch and H. Yonezawa
The Intergenerational Incidence of Green Tax Reform
- 18/286 J. Abrell, S. Rausch, and C. Streitberger
The Economics of Renewable Energy Support
- 18/285 K. Borissov, L. Bretschger and A. Vinogradova
Carbon Pricing, Technology Transition, and Skill-Based Development
- 17/284 H. Gersbach, A. Mamageishvili and O. Tejada
Assessment Voting in Large Electorates
- 17/283 H. Gersbach, A. Mamageishvili and O. Tejada
Sophisticated Attacks on Decoy Ballots: A Devil's Menu and the Market for Lemons
- 17/282 S. Houde, J. E. Aldy
The Efficiency Consequences of Heterogeneous Behavioral Responses to Energy Fiscal Policies
- 17/281 Chiara Colesanti Senni
Energy Transition, Technological Spillovers and Elasticity of Substitution
- 17/280 Anna Alberini, Olha Khymych and Milan Scasny
Response to Extreme Energy Price Changes: Evidence from Ukraine
- 17/279 M. Filippini, G. Masiero and S. Steinbach
The Impact of Ambient Air Pollution on Hospital Admissions
- 17/278 M. Filippini and T. Wekhof
The Effect of Culture on Energy Efficient Vehicle Ownership
- 17/277 L. Bretschger, A. Pattakou
As Bad as it Gets: How Climate Damage Functions Affect Growth and the Social Cost of Carbon
- 17/276 J. Blasch, M. Filippini, N. Kumar, A. Martinez.Cruz
Narrowing the energy efficiency gap: The impact of educational programs, online support tools and energy-related investment literacy