Policy Papers

Carbon Pricing and Power Sector Decarbonisation : Evidence from the UK

Marion Leroutier

PP 2019-03

Suggested citation:

M. Leroutier *(2019)*. Carbon Pricing and Power Sector Decarbonisation : Evidence from the UK. *FAERE Policy Paper, 2019-03.*

Carbon Pricing and Power Sector Decarbonisation: Evidence from the UK

Marion Leroutier*

June 13, 2019

Abstract

The electricity and heat generation sector represents about 25% of global greenhouse gas (GHG) emissions. Policy-makers have implemented a variety of instruments to decarbonise their power sector. This paper examines the UK Carbon Price Floor (CPF), a novel carbon pricing instrument implemented in the United Kingdom in 2013. After describing the potential mechanisms behind the recent UK power sector decarbonisation, I apply the synthetic control method on country-level data to estimate the impact of the CPF on per capita emissions. I discuss the importance of potential confounders and the amount of net electricity imports imputable to the policy. Depending on the specification, the abatement associated with the introduction of the CPF range from 106 to 185 millions tons of equivalent CO_2 over the 2013-2017 period. This implies a reduction of between 41% and 49% of total power sector emissions by 2017. Several placebo tests suggest that these estimates capture a causal impact. This paper shows that a carbon levy on high-emitting inputs used for electricity generation can lead to successful decarbonisation.

Keywords: carbon tax, electricity generation, synthetic control method

JEL Codes: D22, H23, Q41, Q48

PRELIMINARY VERSION - PLEASE ASK FOR PERMISSION BEFORE QUOTING

*Paris School of Economics, Universite Paris I, CIRED. Contact details: Email: marion.leroutier@psemail.eu - Postal: 48 Boulevard Jourdan, 75014, Paris, France. I am grateful to Philippe Quirion, Katheline Schubert, an anonymous referee from the FAERE, as well as the participants to the 2018 OECD environmental microdata workshop for their comments and suggestions. Thanks to Sandbag for kindly giving me access to their dataset on EU ETS participants.

1 Introduction

All governments in the world are facing the need to reduce their greenhouse gas emissions in order to tackle climate change. In the past two decades, they have implemented a variety of abatement policies to address this challenge, including economic instruments in the form of carbon taxes and markets (World Bank and Ecofys, 2018). These carbon pricing instruments tend to cover the electricity and heat production sector, which represents 25% of worldwide emissions in 2010 (IPCC, 2014). Notably, power installations represent 66%[1] of the emissions covered by the European Emissions Trading System (hereafter EU ETS), the carbon market covering the largest share of global annual GHG emissions to date. Given the relatively low prices that prevailed for a long period on the EU ETS, several EU countries considered implementing unilateral measures to strengthen the price signal (Newbery et al., 2018). The UK responded by implementing a carbon price floor (hereafter CPF) in its power sector in 2013. Under the CPF, power generators have to pay a carbon tax on fuels used for electricity generation on top of their ETS carbon allowances. The rate of the tax has varied between around £5 and £18 over the 2013-2017 period.

At the same period, the UK power sector has undergone a remarkable transition: between 2012 and 2017, the coal share in electricity generation fell by 33 percentage points, gross consumption decreased by 6% (Source: Eurostat), and power sector greenhouse gas emissions expressed in tons of equivalent CO_2 (hereafter CO_2e) decreased by around 57%. The rapid transformation of the UK power sector has received significant coverage in the media and policy reports (Evans, 2019; Brown, 2017). However, none of these accounts tries to measure the contribution of specific policies, and in particular of the Carbon Price Floor, to such de-carbonisation.

In this paper, I adopt a policy evaluation approach and estimate the impact of the Carbon Price Floor (hereafter CPF) on per capita emissions from the UK power sector. Using the synthetic control method exposed in Abadie and Gardeazabal (2003) and Abadie

[1] Source: European Union Transaction Log

et al. (2010, 2015), I build a counterfactual UK power sector with a weighted combination of other European countries' power sector and compare the emission paths of the actual and synthetic UK power sector. I find that the introduction of the CPF is associated with a decrease of between 106 and 185 millions ton of equivalent CO_2 (hereafter $MtCO_2e$) in the UK power sector over the 2013-2017 period. Compared to the synthetic UK emissions, the actual UK emissions were lower by between 41% and 49% in 2017. I then run a set of placebo tests suggesting that the impact is causal. The lower bound of my estimate corresponds to a setting where the closure or conversion of some power plants is assumed to be independent from the CPF and entirely caused by other policies. The upper bound assumes that these decisions are either controlled for in the empirical strategy or caused by the Carbon Price Floor.

This paper contributes to several strands in the literature: first, it contributes to the growing empirical literature evaluating the impact of regional and national carbon pricing instruments, notably in Europe. Although the power sector represents a large share of total emissions covered by carbon taxes and markets, there is still relatively few evidence on the effectiveness of carbon pricing on power sector abatement. Furthermore, the estimation results can differ depending on the technique used and period considered, as suggested in Martin et al. (2016) for the EU ETS case: using detailed data from power generators, Ellerman and McGuinness (2008) estimate that the EU ETS caused an abatement of 13 to 21 millions tons of CO_2 in the UK power sector in 2005 and 2006. In contrast, the papers based on aggregate data - mostly using ex-ante simulations - point to a rather low impact. One explanation for the relative scarcity of such evaluations might be the difficulty to find a good counterfactual to the plants covered by the scheme. While this difficulty applies to other sectors covered by the EU ETS, it is particularly acute for the power sector, where almost all GHG-emitting installations are covered[2]. To my knowledge, only one other paper,

[2]The only installations exempted from the scheme are those with a rated capacity of less than 20 Megawatt thermal input (MWth) are exempted, which typically represent a very small share of a country's total capacity and production. For example, in the UK in 2015, only 13 of the 95 active fossil-fuel fired power plants have a capacity below 20 MWth. They represent only 0.2% of the installed capacity (Source:

unpublished to date, looks specifically at the impact of the CPF on CO_2 emission in the UK power sector: Abrell et al. (2019) estimate counterfactual plant-level generation using machine learning techniques, and find that the UK CPF resulted in a cumulated abatement of 26MtCO_2e over the 2013-2016 period. For the same period considered, I find an abatement between 61 MtCO_2e (lower bound) and 116 MtCO_2e (upper bound), which is 2.3 to 4.5 times higher. The difference in the method used and channels considered likely explains the difference in our respective results: Abrell et al. (2019) focus on one mechanism only via which the CPF affects power sector emissions, namely the fuel switching from coal to gas-fired generation[3]. In contrast, this paper uses less granular data but adopts a method allowing to take into account all the mechanisms via which a carbon tax on high-emitting input fuels may induce a decrease in emissions: fuel switching at the intensive margin, but also production re-allocation to renewable sources of energy, plant closure, and changes in demand or imports - although I show that these two last channels probably play an unsignificant role compared to changes in the emission intensity of domestic production.

Second, this paper relates to the scarce existing literature examining the rapid decarbonisation of the UK power sector in the recent years. The only published paper to date, by Staffell (2017), describes the evolution of key power sector variables (such as wholesale prices, supply, demand, generation mix) using granular - up to half-hourly - data. In contrast, this paper aims at identifying a causal link between the Carbon Price Floor policy and carbon emissions, and to do so it contrasts the evolution of UK emissions with that of other countries. More broadly, this paper draws from a broader literature examining the drivers of carbon emissions in the power sector. In particular, it uses insights from Van den Bergh and Delarue (2015), who look at the role and interactions of different abatement channels and their contributions to the shape of marginal abatement cost curves in the power sector. It also draws from Ellerman and McGuinness (2008) and Kirat and Ahamada (2011), who

Digest of United Kingdom Energy Statistics)

[3]Fuel switching from gas to coal arises because carbon pricing changes the relative marginal cost of coal- and gas- fired plants. This change affects the merit order (the ranking of power plants by ascending marginal costs), which determines each individual plant's output at every period (hour or half-hour)

emphasize the role of EU ETS carbon prices and the coal-to-gas price ratio in encouraging fuel switching. Third, this paper contributes to a recent literature - mainly unpublished to date - applying the synthetic control method to the analysis of environmental policies: for example, Lee and Melstrom (2018) estimate the impact of a regional carbon pricing initiative on electricity imports, Andersson (2017) evaluate the impact of the Swedish carbon tax, Gloriant (2018) the impact of the French carbon tax, and Isaksen (2018) evaluates the effect of international pollution protocols. A distinctive feature of this work is to rely on plant-level data aggregated at the country-level, allowing to account for specific shocks experienced by individual plants (e.g. plant closure induced by EU air pollution regulation). Beyond the academic contribution, this paper is also relevant from a policy perspective, since it tackles the topical question of how to achieve emission targets by decarbonising electricity. To be in line with the 2015 UN Paris Climate Agreement, OECD countries including most European countries need to be coal-power-free by 2030 (Rocha et al., 2016). While Germany - the country in Europe with the highest power sector emissions - recently announced that coal phase-out would only be achieved in 2038, other countries such as France, the Netherlands, Italy, Ireland and Sweden announced earlier dates. However, the means necessary to achieve such transition are still under discussion, and the announced date for phase-out (2022) already seems unrealistic for France (Le Hir, 2019). Lessons can probably be drawn from the UK experience exposed in this paper.

The paper is organized as follows: Section 2 presents the Carbon Price Floor policy and the theoretical mechanisms underlying decarbonisation in the power sector; Section 3 describes the empirical strategy; Section 4 presents descriptive evidence and the main results from the Synthetic Control Method; Section 5 discusses the potential threats to identification strategy; Section 6 presents some sensitivity analyses; Section 7 concludes.

2 The UK Carbon Price Floor: context and expected effects

2.1 The UK Carbon Price Floor

As explained above, the Carbon Price Floor was introduced in the GB power sector in April 2013. The idea of providing a price floor to the low-standing emission allowances on the EU ETS had already been discussed as in 2009 but the Labour Government then at power was opposing it. The policy was put back on the table by the Coalition Government in 2010 and was part of their government agreement (Ares and Delebarre, 2016a). At the time, through the five-year carbon budgets introduced by the 2008 Climate Change Act, the UK had to abate 236 MtCO2e across all economic sectors between the 2008-2012 and the 2013-2017 periods. Furthermore, low prices on the EU carbon market were limiting the potential for high emission reductions among industrial installations covered by the EU ETS. In March 2011, after some expert consultations, the Government announced its decision to introduce a Carbon Price Floor in the 2013/2014 budget year[4]. The Carbon Price Floor would apply only to the generators located in Great Britain and not those located in Northern Ireland, since the latter are part of a single electricity market together with generators from the Republic of Ireland , who would not impose the same carbon tax on its generators. The official announced goal of the CPF was to tackle price uncertainty on the EU ETS and encourage investment in low-carbon technologies in the generation sector; in official communication documents, the Carbon Price Floor was labelled "support and certainty for low-carbon investment"(Hirst, 2018).

In practice, the CPF translates into a tax called the Carbon Price Support (CPS hereafter), which yearly rate depends on projected EU ETS allowance prices. The CPF was initially expected to increase over time, with an overall carbon price target of £30 due by

[4]The budget year over which the annual tax rate is set runs form 1st April of year T to to 31 March of year T+1

Table 1: Level of CPS rate for each period in pound per ton of CO2e

Period	CPS rate in £/tCO_2e
April 2013/March 2014	4.96
April 2014/March 2015	9.55
April 2015/March 2016	18.08
April 2016/March 2017	18
April 2017/March 2018	18
April 2018/March 2019	18

Source: Ares and Delebarre (2016a)

2020. The CPS rate was expected to increase accordingly. CPS tax rates in pound per ton of equivalent CO_2 (£/tCO_2e) were announced in March 2011 for the 2013/2014 and 2014/2015 periods, with indicative rates announced for 2015/2016 and 2016/2017. In the 2013 budget, the 2015/2016 rate was finally set at a higher level than the indicative rate, and even higher indicative rates were given for 2016/2017 and 2017/2018. However, in 2014, the Government decided to freeze the CPS rate to £18/tCO_2e (about the 2015/2016 level) until 2019/2020, after business representatives expressed concern over the competitiveness of the UK energy-intensive industries due to generators passing on the tax costs (Ares and Delebarre, 2016a). Table 1 shows the confirmed level of the CPS rate for each period.

Because of the freeze and since rates have been determined based on expected future EU ETS allowance prices[5] rather than actual prices, the trajectory of the Carbon Price Floor looks somewhat different from what was first announced. Figure 1 overlays the actual CPS rates converted to euro with the observed European Emission Allowances (EUA) prices since 2009, which sum gives the actual observed carbon price floor. As visible on the figure, the CPS component increases substantially the carbon price paid by GB power generators compared to non-GB generators covered by the EU ETS. In 2016, the year where the relative difference is the largest, GB power generators paid a carbon price more than five times higher than non-GB generators.

[5]More precisely, rates have been determined based on the average annual ICE-ECX benchmark end of day settlement price for carbon for delivery in the target year (Ares and Delebarre, 2016a)

Figure 1: The Carbon Price Support and EUA price on the EU ETS

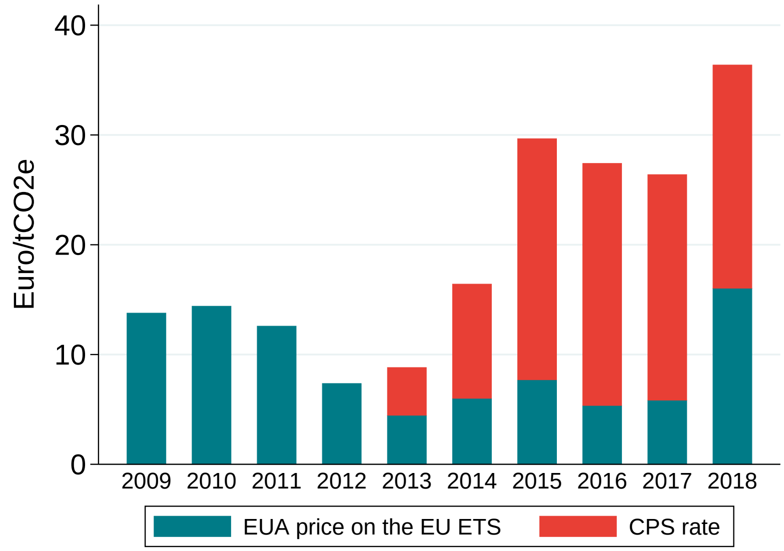

Note: EUA price data retrieved from Sandbag website. CPS prices retrieved from House of Commons library (2016), adjusted with appropriate weights to reflect the January to December period rather than April to March, and converted to Euro using yearly averages of monthly market exchange rates.

In terms of coverage, the tax applies to all generators other than stand-by ones[6], with a rated thermal input greater than 2 MWth, located in Great Britain, and producing electricity from fossil fuels. This includes conventional power plants as well as Combined Heat and Power (CHP) operators[7] and auto-generators (HM Revenue & Customs, 2017). The CPS tax is expressed as a specific component of pre-existing taxes[8] and its rate depends on the carbon content of the input fuel used for power generation. The only case where generators benefit from an abatement of the tax is if they are fitted with carbon capture and storage (CCS) technology - although no UK power plant has been equipped with such technology

[6]Stand-by generators are generators used to provide emergency electricity supplies in the event of a failure of a building's usual electricity supply and used only for that purpose. They are typically found in hospitals and other such facilities

[7]CHP plants are only liable to pay the CPS on the so-called "deemed supply of fuel", the share of fuel used to produce electricity for the grid. E.g. for a CHP which generation is 80% electricity and 20% heat, the 20% on the heat are exempted from the tax

[8]More precisely, for power plants using solid fossil fuels (such as coal), natural gas or LPG as an input fuel, the tax is expressed as a specific component of the Climate Change Levy (CCL). The CCL is an environmental tax levied on taxable commodities supplied to businesses and the public sector, of which power generators were exempted thus far. For power plants using oil as an input, which had been benefiting from a relief on tax duty since 2006, the CPS rate is expressed as a reduction to this excise tax relief(HM Revenue & Customs, 2017)

Table 2: Level of CPS rate by input fuel for each period, in pence per fuel-specific unit

Period	Natural Gas (p^1 per kWh)	Petroleum gas[2] (p per kg)	Coal[3] (p per GJ on GCV[4])	Fuel oil[5] (p per litre)	Gas oil[6] (p per litre)
2013/2014	0.091	1.146	44.264	1.568	1.365
2014/2015	0.175	2.822	81.906	3.011	2.642
2015/2016	0.334	5.307	156.86	5.730	4.990
2016/2017	0.331	5.28	154.79	5.711	4.916
2017/2018	0.331	5.28	154.79	5.711	4.916
2018/2019	0.331	5.28	154.79	5.711	4.916

Notes: [1]p stands for pence. [2]Or other gaseous hydrocarbon in a liquid state.[3]And other solid fossil fuels. [4]GCV stands for Gross Calorific Value and means that the amount of liquid water contained in the coal prior to combustion, that leaves as vapour, is taken into account in the energy calculation. [5]Or other heavy oil or rebated light oil. [6]Or kerosene or rebated bioblend. Source: HM Revenue and Customs 2014, 2016 and 2017 and Envantage website: https://www.envantage.co.uk/carbon-management/climate-change-levy-agreement/climate-change-levy-rates.html

thus far (Committee on Climate Change, 2018). Table 2 shows the CPS tax rates by input fuel, expressed in different units. Converting all tax rates to the same unit across fuels, one realizes that the CPS rate on coal is about 70% higher than the tax on natural gas, in line with the much higher emission factor of coal. The CPS rate thus substantially increases the relative cost of coal-fired power generation compared to gas-fired one.

2.2 Theoretical mechanisms underlying decarbonisation

The outcome considered throughout the analysis is greenhouse gas emissions from the power sector taken per capita. The theoretical channels leading to a decrease in per capita emissions can be illustrated with a basic equation decomposing yearly per capita emissions in the power sector. Although all the variables defined should be interpreted as values for a given country c in a given year t, the indices are not written to make equations easier to read.

$$\frac{Q_{CO_2e}}{POP} = \frac{Q_{elec}}{POP}\frac{Q_{CO_2e}}{Q_{elec}} \tag{1}$$

Where POP is total population, Q_{CO_2e} is the quantity of emissions from the domestic power production, Q_{elec} is domestic gross power production (in GWh), and $\frac{Q_{CO_2e}}{Q_{elec}}$ is the emission intensity in the domestic power sector (in tCO2e/GWh).

Q_{elec} can be rewritten as the difference between domestic gross electricity consumption C_g and net imports $(M-X)$. Gross electricity consumption is itself the sum of net consumption C_n (equivalent to demand), the amount of network losses, and the amount of electricity used by power generators. The two latter components are grouped in the variable L. This leads to the following equation:

$$\frac{Q_{CO_2e}}{POP} = (\frac{C_n}{POP} + \frac{L}{POP} - \frac{(M-X)}{POP})\frac{Q_{CO_2e}}{Q_{elec}} \tag{2}$$

The right-end side of the equation suggests that four different channels may lead to a decrease in per capita emissions: a decrease in consumption per capita $\frac{C_n}{POP}$ (the *demand channel*), a decrease in the amount of network losses and self-consumption of electricity by power generators $\frac{L}{POP}$ (the *network efficiency* channel), an increase in net imports per capita $\frac{(M-X)}{POP}$ (the *trade channel*), and a decrease in the average emission factor of the domestic power sector (the *emission intensity channel*). The latter can be further decomposed as:

$$\frac{Q_{CO_2e}}{Q_{elec}} = \sum_i e_i q_i \tag{3}$$

Where e_i is the average emission intensity of GHG-emitting process i used for electricity production and q_i is the share of gross electricity production covered with process i (also called the electric mix). Accounting conventions usually imply that renewables and nuclear energy sources are considered to generate zero emissions[9], so only fossil-fuel-fired power plants matter in this equation.

The *network efficiency* channel can be considered a negligible mechanism since it is unlikely to be affected by the CPF and very stable over time[10]. In the rest of the analysis I

[9]For example, it is the case in the way carbon emissions are calculated in the EU ETS. While it is true that power generation itself is emission-free with these sources, manufacturing solar panels, wind plants or nuclear plants is not. For biomass, power generation does release greenhouse gases, but the EUTL does not count them since carbon released when solid biomass is burned is expected to be re-absorbed during tree growth : see https://ec.europa.eu/clima/sites/clima/files/ets/docs/com_2018_842_final_en.pdf

[10]L can be estimated with Eurostat data as gross production Q_{CO_2e} plus net imports $(I-X)$ minus net

focus on the three channels of demand, trade, and emission intensity. It is worth noting that these channels are not independent from each other: for example the emission intensity of domestic production is probably correlated with the amount of net imports. However, given the descriptive nature of the exercise, I consider each channel independently from the other in the rest of the section.

2.3 Disentangling the role of the Carbon Price Floor and other policies

While the CPF might have had an impact on each of these channels, other factors might also have contributed to the evolution of each of them.

Demand per capita might have been dampened by the wholesale price effect of the Carbon Price Floor[11]. Yet other factors such as the continuous improvement of energy efficiency in buildings and electric appliances might have played a role too, which would be consistent with the continuous decreasing pattern of Figure 4a.

Regarding the trade channel, an increase in net imports would be an expected consequence of the CPF, which increases the relative cost of domestically produced electricity compared to imported electricity. However, the low interconnection between the UK and neighbouring countries is a physical constraint to the potential for increased imports: Being an island, the UK is relatively limited in its ability to trade electricity, which can only occur via undersea inter-connectors for GB, and via undersea connections or ground connections to the Republic of Ireland for Northern Ireland. Notably, there was a 50% increase in the UK interconnection capacity between 2010 and 2012: In 2010, the UK was only connected

consumption Y_n. It represents a constant share of around 18% of total gross production for all 21 countries

[11] Given past evidence of carbon cost pass-through for the European power sector as a whole in both phases of the EU ETS (Zachmann and Hirschhausen (2008) for the first phase and Hintermann (2016) for the second phase), the CPF may well have led to an increase in electricity prices too, as suggested by Newbery et al. (2018) and Ares and Delebarre (2016b). However, for all customers the effect on demand might have been mitigated by a low price-elasticity of demand ; furthermore, for large buiness customers the price effect was mitigated by a compensation scheme introduced for electro-intensive industries via a specific component of a larger Energy Intensive Industries support (meant to compensate the cost increase induced by climate change policies in general) Ares and Delebarre (2016b)

to France and Ireland, with a total capacity of 2,500 MW. In 2011, GB became interconnected with the Netherlands, and in 2012 a new undersea interconnector with the Republic of Ireland was completed (OFGEM, 2013). Given the timing of these new interconnections shortly before the introduction of the CPF, any observed increase in net imports could also from this new infrastructure.

Regarding the emission intensity channel, the Carbon Price Floor can directly impact both e_i and q_i. The CPS leads to a higher increase in marginal generation cost for less efficient plants (i.e plants using more coal or more gas by kWh of electricity output). Within a given fuel, it will thus lead to a production reallocation to more efficient plants. The magnitude of such impact depends on the potential for production re-allocation within a given input (there must be room for producing more for the most efficient plants), and of the heterogeneity of plants (there must be some plants more efficient than others). The CPF can also induce a change in the electric mix, with two distinct impacts: in the short-run, the CPF increases the marginal cost of coal generation compared to gas generation and induces fuel switching. In the long run, this change in the cost of running high-emitting power plants can also affect investments in low-carbon generation (Van den Bergh and Delarue, 2015).

However, two other policies implemented over the period of interest might have impacted the fuel mix. First, a 2001 EU regulation called the Large Combustion Plants Directive induced the closure of several high-emitting power plants. This air pollution regulation imposed emission limit values for three air pollutants, which large combustion plants had to respect by 2008. Targeted installations could however choose to opt out of the regulation if they committed not to operate the plant for more than 20,000 hours between 1 January 2008 and 31 December 2015 (European Commission, 2001); in practice, the concerned plants had to either retrofit or shut down by 2015. This policy applied to all EU countries. In the UK, nine plants chose this opt-out option and shut down between 2012 and 2015. Three other plants chose to opt-out for part of their sites only (Source: EEA website). In total, these 12 fully or partly opted-out plant represented 11% of UK power sector emissions in 2011. To the extent that these plants were not replaced by new coal-fired power plants respecting the EU air quality regulation, their closure would have contributed to the decrease of UK power sector emissions.

While the first policy is a European one, the second one is specific to the UK: the UK government started supporting coal-fired power plants converting to biomass in 2012. This support first took the form of dedicated Renewables Obligation Certificates (ROCs). These certificates were embedded in the broader Renewables Obligation scheme designed to support the deployment of large-scale renewable electricity generation; they create an obligation for electricity suppliers to source a proportion of their electricity from plants with ROCs. The scheme was replaced by the Contract for Difference scheme in 2014, which guarantees a government flat payment to generators converting to biomass[12]. Two power plants received government support for the conversion: Drax power station, representing 14% of UK power sector emissions in 2012, benefited from these two schemes and intensified its biomass conversion - started in 2009 at its own initiative - in 2012. Lynemouth power station received support under the CfD scheme and converted to biomass in 2016. For a given power production, such conversions to biomass automatically reduce the total CO_2e emissions from these power stations, since biomass is considered emissions-free. The timings of implementation of the CPF and the support policy for biomass conversion are close. It makes wonder to what extent one led to the other. The decision to grant this subsidy might have been partly motivated by industry lobbying and the government wish to weather the transition for coal-fired power plants facing the Carbon Price Support. On the other hand, even without the Carbon Price Floor it might have made sense to promote coal transition to biomass via subsidies.

Given the difficulty to assess how much these policies were independent from the CPF and led to additional emission decreases, they will be taken into consideration in the estimation

[12]The value of this flat payment is supposed to reflect the difference between the strike price derived from an auctioning process and reflecting generators' true cost (incl. investment in conversion in the case of coal-to-biomass conversion), and the average market price of electricity. See https://www.emrsettlement.co.uk/about-emr/contracts-for-difference/ for more details

strategy presented in the next section.

3 Empirical strategy

3.1 The synthetic control method

To estimate the impact of the Carbon Price floor on the power sector in the UK, I use the synthetic control method (SCM) exposed in Abadie and Gardeazabal (2003) and Abadie et al. (2010, 2015). This method consists in building a counterfactual UK power sector by applying appropriate weights to the set of other European countries' power sectors. It is particularly appropriate in the context of the CPF since the "treatment" applies to one country only, and within the country it affects almost all power installations, without time variation in treatment. Within the UK, there is therefore no obvious group of installations that would perfectly reproduce the counterfactual of how the treated British power plants would have evolved absent the policy. Using the notation traditionally used in the policy evaluation literature, the challenge is to estimate $\beta_{\mathrm{UK}t}$ when t≥2013, defined as:

$$\beta_{\mathrm{UK}t} = Y^1_{\mathrm{UK}t} - Y^0_{\mathrm{UK}t} = Y_{\mathrm{UK}t} - Y^0_{\mathrm{UK}t} \tag{4}$$

$Y^1_{\mathrm{UK}t}$ designates, at each period, UK power sector emissions in the presence of the CPF policy. $Y^0_{\mathrm{UK}t}$ designates, at each period, UK power sector emissions in the absence of the policy. $\beta_{\mathrm{UK}t}$ designates the difference between the two. The challenge to estimate $\beta_{\mathrm{UK}t}$ comes from the fact that while $Y^1_{\mathrm{UK}t}$ is observed when t≥2013, $Y^0_{\mathrm{UK}t}$ is not.

Let us assume, after Abadie et al. (2010), that for each country c and period t the outcome in the absence of intervention Y^0_{ct} can be modelled as the following linear factor model:

$$Y^0_{ct} = \delta_t + Z_{ct}\alpha + f'_t\lambda_c + \epsilon_{ct} \tag{5}$$

δ_t is a time fixed effect, Z_{ct} is a vector of observed exogenous country characteristics, f_t

is a vector of unobserved time effects or factors (and f_t' denotes its transpose), λ_c is a vector of unobserved country-level effects or factor loadings, and ϵ_{ct} is the error term with mean 0 (typically transitory shocks at the country level).

One can easily see that such a model is more flexible than the typical difference-in-difference equation, since time effects and individual effects are allowed to interact. Abadie et al show that with such specification, it is possible to use as an estimator of $\beta_{\mathrm{UK}t}$ a function of outcomes observed post-treatment in other countries:

$$\hat{\beta}_{\mathrm{UK}t} = Y_{\mathrm{UK}t} - \sum_{j=1}^{J} w_j^* Y_{jt} \tag{6}$$

Where $\sum_{j=1}^{J} w_j^* Y_{jt}$ is a weighted combination of the outcome for J countries having not implemented the policy, and the vector $W^* = (w_1^* ... w_J^*)'$ should satisfy the three following conditions:

$$\begin{cases} \sum_{j=1}^{J} w_j^* & = 1 \\ \overline{Y}_{\mathrm{UK}}^{K} & = \sum_{j=2}^{J} w^* \overline{Y}_j^K \\ Z_{\mathrm{UK}} & = \sum_{j=2}^{J} w^* Z_j \end{cases}$$

With $\overline{Y}_{\mathrm{UK}}^{K}$ a linear combination of pre-intervention per capita power sector emissions in the UK, and $\overline{Y}_j^K$ a linear combination of pre-intervention per capita power sector emissions for country j (for example it can be the simple mean of pre-intervention outcomes $\overline{Y}_j^K = 1/T_0 \sum_{t=1}^{T_0} Y_j$). Abadie et al also show that the estimator gets closer to the true parameter $\beta_{\mathrm{UK}t}$ when the number of pre-treatment periods is high compared to the scale of transitory shocks affecting countries.

In practice, to find the appropriate W vector I rely on an algorithm created by Abadie et al. The algorithm minimizes the distance between a vector of pre-intervention characteristics in the treated region, X_{UK} (with dimensions $K \times 1$) and a weighted matrix of pre-intervention characteristics in the non-treated regions, $X_0 W$ (with dimensions $K \times K$). Pre-intervention characteristics are of two types: 1) the linear combinations of pre-intervention outcomes $\overline{Y}_j^K$, and 2) the country characteristics Z_j not affected by the intervention. To obtain the W vector, the programme starts with a positive and semi-definite matrix V that defines a dot product. The distance between X_{UK} and X_0W can then be written as

$$X_{\text{UK}} - X_0W = \sqrt{(X_{\text{UK}} - X_0W)'V(X_{\text{UK}} - X_0W)} \tag{7}$$

The goal is to find the vector $W^*(V)$ that minimizes this distance. Such minimization actually comes down to finding the right V matrix, which can be shown to be equivalent to a diagonal matrix assigning weights to linear combination of characteristics in X_{UK} and X_0W. Like Abadie and Gardeazabal (2003), I choose the V minimizing the mean squared prediction error of the outcome variable in the pre-treatment periods. Formally, let Y_{UK} be the (8×1) vector of pre-2013 power sector emissions from 2005 to 2012 for the UK and Y_j be the $(8 \times J)$ matrix of pre-2013 power sector emissions for the J other European countries. Then V^* is chosen such that:

$$V^* = argmin(Y_{\text{UK}} - Y_jW^*(V))'(Y_{UK} - Y_jW^*(V)) \tag{8}$$

where V is the set of all non-negative diagonal $(K \times K)$ matrices.

The ability to build a good synthetic control can be assessed by at least two criteria: first, the closeness of pre-intervention characteristics between the treated and synthetic units. This depends on how well these characteristics predict the outcome and can be assessed by comparing pre-intervention characteristics for the treated and synthetic country. The second criterion is the closeness of pre-intervention outcomes between treated and synthetic control, which can be seen graphically or by computing the Mean Squared Prediction Error (MSPE)[13].

Beyond identification, inference can be derived by running a set of placebo studies, which

[13]The MSPE gives the average of the squared difference between the treated unit's and the synthetic control's pre-intervention outcomes.

consist in applying the same synthetic control methodology to states that did not implement such a policy (Abadie et al., 2010).

3.2 The Data

The empirical strategy relies on a comparison between the UK and other European countries. Therefore, the data used cover all European countries and not only the UK. To obtain data on CO_2e emissions in the power sector, I use individual plant-level data from the European Union Transaction Log (hereafter EUTL), the official register of the EU ETS managed by the EU Commission. The register checks, records and authorises all transactions taking place between participants in the EU ETS. Participants have to monitor and report their CO_2e emissions each year and surrender enough emission allowances to cover their annual emissions. Participants' annual emissions are verified by an accredited verifier. The EUTL provides data on annual verified emissions for every installation covered by the EU ETS from 2005 to 2017.

I restrict the analysis and associated database construction to the EU-countries which have been part of the EU ETS since the beginning of the scheme in 2005, and which are big enough. Twenty-one countries are left in my main dataset[14].

The EUTL gives some sectoral information on individual installations, but not whether they are a power plant or not. To retrieve this crucial information - since the analyses focuses on the power sector -, several data sources were combined. The specific steps followed are explained in Appendix A.

I obtain an unbalanced plant-level panel of 14,065 plants subject to the EU ETS and located in 21 countries, split into 4,938 power plants (302 for the UK) and 9,127 non-power plants. It is worth noting that the power plants located in the UK and covered by my data

[14]The excluded countries are: Romania and Bulgaria (who joined in 2007); Croatia (who joined in 2013); Slovenia (which was part of the scheme from the beginning but only joined the EU in 2015); Norway, Liechtenstein and Iceland, all non-EU countries which joined the scheme in 2008; and the three countries having less than ten power plants subject to the EU ETS: Luxembourg (only nine power plants), Cyprus (only three) and Malta (only two)

do not fully overlap with the power plants subject to the Carbon Price Floor. Two types of power installations are included in my data but not subject to the CPF: first, power installations located in Northern Ireland, which represent a very small share of total UK power sector emissions (2.4% in 2012). Second, standby generators such as those part of hospitals, likely to be only used in case of power shortages. Given that their combined CO2e emissions are also very low, I deem them negligible compared to the rest of UK power plants[15]. In turn, some power plants subject to the CPF are not in my data: these are the fossil-fuel power plants with a rated thermal input between 2 and 20 MWth. These small plants may represent a substantial share of the total number of power plants, but logically they represent a very small share of total emissions. Overall, most emissions covered by the CPF are in my data. I am also confident that this EUTL emission data accurately reflect other countries' power sector emissions: even if some countries have a larger proportion of small power plants, the bulk of emissions presumably comes from the few large conventional power plants.

Plant-level emission data is aggregated at the country-level separately for power and non-power plants to obtain country-level power sector emissions. I add to this panel a set of country-level variables from Eurostat on electricity production by source, consumption and net imports used in section 4.1. I also add country-level variables used for the synthetic control method, described more specifically in the next section and in Appendix B.

3.3 Choice of predictors

The set of pre-intervention characteristics X_0 used to build the synthetic UK should be variables predicting aggregate per capita CO_2e emissions in the power sector. To select them, I rely on the literature looking at drivers of emissions (Ellerman and McGuinness, 2008; Van den Bergh and Delarue, 2015; Lee and Melstrom, 2018). For each predictor, the sources and data processing steps are described in Appendix B.

[15]in 2012, the six installations belonging to hospitals represent only 0.05% of the UK power sector's emissions

In countries relying both on coal- and gas-fired power plants for electricity generation, like the UK, fuel switching has been identified as an important determinant of emissions variation. Ellerman and McGuinness (2008) build a linear model of fuel switching and find that the coal-to-gas price ratio is an important predictor. The Carbon price floor directly impacts this ratio since it translates into a differentiated tax on input fuels depending on their carbon content. As a predictor, I therefore use the pre-2013 country-specific coal-to-gas price ratio. Given the non-linear effects of this variable, I also include a quadratic term[16].

While the coal-to-gas price ratio is an important predictor of the amount of high-emitting versus low-emitting fuels used for generation, it does not include polluting solid fuels other than conventional coal, in particular lignite. Lignite is a low-quality type of coal with a very poor calorific value and high emission intensity, that is mostly consumed domestically by the power sector (Berghmans and Alberola, 2013). To account for the large differences in lignite resources and use across European power sectors, I include in my predictors set a dummy variable identifying the countries with large lignite resources: Germany, Poland, Hungary, Greece, Czech Republic, and Bulgaria.

To account for demand from CO_2-emitting power plants, I also build a per capita residual load variable. Residual load measures the amount of electricity demand that needs to be covered by fossil fuels and biomass once generation from so-called "must-run" power generators (nuclear power plants) and those that generate with almost no marginal cost (solar, wind and hydro) is withdrawn. It is defined as the difference between electricity consumption and the generation from renewables and nuclear power plants, divided by total population. In a sensitivity analysis (see section 6.2), I use other proxies for demand from CO_2-emitting power plants. The results are unchanged.

Finally, to account for the impact of the opting-out regime of the LCP directive mentioned in section 2.3, I add as a predictor the amount of country-level per capita emissions coming from opted-out installations. The decision to opt out from the LCP Directive was

[16] In a sensitivity analysis, I also run a specification without this quadratic term. The coal-to-gas ratio weighs less, but the results are virtually unchanged

independent from the CPF since it had to be made by 2007 at the latest. Furthermore, the LCP directive was imposed to all EU countries, and the UK was not the only country with a significant number of plants deciding to opt out. Therefore I can estimate for each country the amount of annual emissions from opted-out plants and use it as a predictor. To do so I match my emission data available at the installation level with data from the European Environmental Agency listing the plants opted out from the directive. I take the value of this variable in 2009, shortly before the introduction of the CPF. This way, the synthetic UK will be built as a weighted combination of countries having a similar amount of emissions from plants subject to close or retrofit shortly before the announcement of the CPF.

In addition to these 5 predictors, I add lagged per capita emissions for the first and last year of the pre-treatment period (2005 and 2012), which is standard in the SCM literature. For the optimization described above, the predictors are averaged over the 2005-2012 period, except for the coal-to-gas price ratio and its square, averaged over the 2007-2012 period (see Appendix B for more details).

3.4 Choice of the donor pool

The "donor pool" designates the set of countries not affected by the CPF policy that will potentially enter the composition of the synthetic UK. To replicate the evolution of the UK power sector in the absence of the carbon price floor, it is important to discard from the donor pool the countries that are likely to be poor counterfactuals (Abadie et al., 2010). This can include three types of countries: First, countries that suffered idiosyncratic shocks to the outcome of interest, either by directly introducing a policy targeting the power sector or via a more generic exogenous shock likely to affect the electricity sector. Second, countries that are likely to be directly affected by the CPF. Third, countries with very different characteristics compared to the UK, which may cause severe interpolation biases.

Regarding the first type of countries, by 2017 no other European countries had adopted such a significant pricing policy as the UK and its carbon price floor. Although France and the Netherlands have recently discussed introducing a carbon price floor as well (Newbery et al., 2018), only the latter have passed a concrete law in August 2018, and the Dutch CPF will only start in 2020. The biggest change in other European countries' power sectors is the case of Germany, which unexpectedly decided to phase out of nuclear energy following the 2011 Fukushima nuclear accident. I therefore exclude Germany from the donor pool. Since the European debt crisis affected the Greek economic environment very heavily over the period, I also exclude Greece. However, including them in the donor pool does not change the results, as shown in section 6.3. Regarding the second type of countries, I do not exclude any country based on these grounds but I extensively discuss the risk of spillover in section 5. Finally, to avoid having countries too different from the UK, I eliminate the three Baltic countries - Estonia, Latvia and Lithuania - which unlike the UK do not use coal for power generation (as can be seen on Figure 6). Since coal-to-gas switching is expected to be one of the important drivers for de-carbonisation, it is most relevant to restrict the analysis to countries who can experience it. In the end, the donor pool includes 15 EU countries.

To ensure that it is possible to build a convex combination of countries that closely reproduce the UK's values for predictors and emissions, I check that there is common support between the distribution of these predictors in the donor pool and in the UK. This is the case for all variables. The corresponding histograms are shown in Appendix C

4 Results

4.1 Descriptive evidence on the UK power sector de-carbonisation

This section documents the recent changes in UK power sector CO_2e emissions compared to the rest of Europe and links it to the mechanisms exposed in section 2.2.

As shown on the left-hand side of figure 2, the UK power sector had the second largest total CO_2e emissions in 2005 (after Germany). From 2013 on - year of the introduction of the Carbon Price Floor, marked by a vertical line on the figure - the UK experienced a large decrease in emissions. The right-hand side of the figure shows that this is also true when emissions are taken per capita: While the UK was among the top emitters in per-capita terms in 2005, by 2017 it had joined the bulk of lower-emitting countries. The potential role of the CPF in explaining such a decrease becomes more plausible when one compares post-2012 emissions to the average for the 2005-2012 period. The left-hand side of figure 3 plots the difference between annual per capita emissions and their 2005-2012 mean. The UK trajectory stands out, especially from 2014 onwards: the majority of countries exhibit a mildly decreasing pattern while the UK undergoes a clear fall. A decrease is visible for three other countries, two of which exhibiting large deviations from the mean over the whole period (Finland and Denmark[17]), and the third one showing a decrease in emissions from 2012-2013 onwards (Greece[18]). The right-hand side of the figure shows the same graph for the non-power installations subject to the EU ETS: there, UK per capita emissions are extremely stable over time, like most other EU countries. This contrast suggests that there is something specific to the UK power sector around 2013. I now describe the potential underlying channels and the role the CPF may have played in such de-carbonisation.

[17] Their high variablity in emissions is likely explained by the inter-annual variation in the share of available renewable sources for electricity generation (due to varying hydro and wind resources for Denmark, and varying hydro resources for Finland)

[18] for which the decrease in emissions can not be traced back to a specific reform in my knowledge. However, the country underwent large reforms amidst the debt crisis around that period, which may have effected the power sector

Figure 2: Aggregate CO_2e emissions of power installations covered by the EU ETS

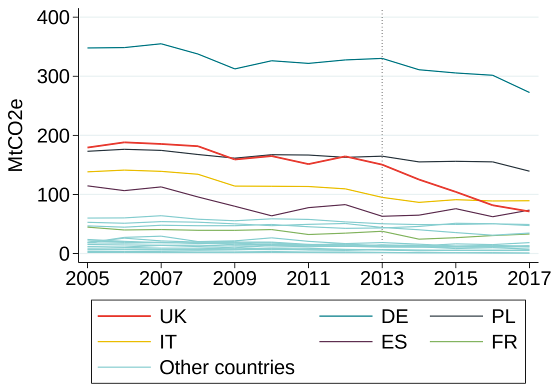

(a) Total absolute emissions

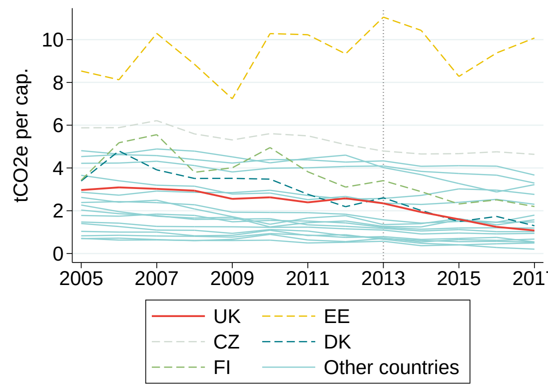

(b) Per capita emissions

Note: Emission values were obtained by aggregating, at the country level and every year, the verified emissions of power generators. The left-hand side value is a simple sum while the right-hand side is the sum of emissions divided by the average country population that year. The "Other countries" include all remaining EU countries except Romania, Bulgaria, Slovenia, Croatia, Malta, Cyprus, and Luxembourg.

Figure 3: Aggregate per capita CO_2e emissions of power and non-power installations covered by the EU ETS: Deviation from the 2005-2012 mean

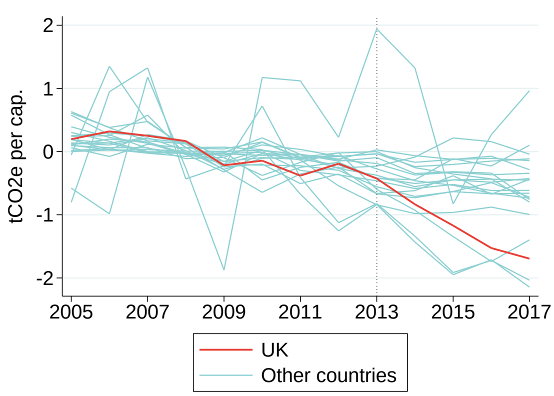

(a) Power installations

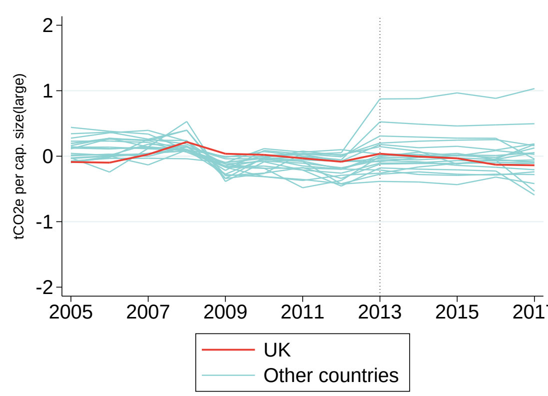

(b) Non-power installations

Note: Emission values appearing on these two graphs were obtained in 2 steps: 1)by aggregating, at the country level and every year, the verified emissions of installations and dividing them by the average country population 2)by taking the difference between the obtained yearly per capita emission and the 2005-2012 average. The "Other countries" include all remaining EU countries except Romania, Bulgaria, Slovenia, Croatia, Malta, Cyprus, and Luxembourg.

Figure 4 shows how the UK compares to the other European countries for each of the channels exposed in section 2.2. Demand per capita exhibits a decreasing trend over the whole period, making the UK reach among the lowest levels in 2017 compared to the other European countries. However there is no obvious break in trend in 2013. Given the relative isolation of the UK and its low interconnection capacity with other countries, UK net imports per capita appear extremely stable over time compared to other countries, although taken in isolation, the UK increased its net imports from 11,800 GWh to 17,500 GWh between 2012 and 2016. The emission intensity channel seems to be a)the largest source for the UK de-carbonisation and b)the channel contributing the most to the break in trend after 2013.

One can dive deeper into that channel following equation (3), by comparing the fuel mix q_i across European countries[19]. Figure 5 shows the fuel mix in the UK power sector in 2005, 2012 and 2017. While the shares are quite similar between 2005 and 2012 (apart from the increasing renewables share), between 2012 and 2017 the share of coal falls dramatically (by 33 percentage points). It is compensated by an increase in the share of gas and non-biomass renewables by 13 percentage points each, an increase in the share of biomass by 5 percentage points and in the share of nuclear by 2 percentage points. The case of the UK looks quite unique when compared with other EU countries like is done on Figure 6, showing the fuel mix for all EU countries in 2012 and 2017 and ranking countries in ascending order of their 2012 coal share. Between the two periods, the UK goes from the 6th highest coal share to one of the lowest, while most other countries keep the same ranking. This evolution is consistent with the theoretical impact of the CPF on the input mix.

[19] Data on e_i is not readily available for all European countries

Figure 4: Channels: evolution of the UK and other countries

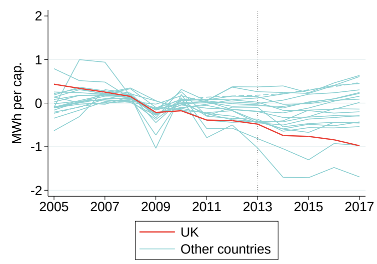

(a) Demand per capita

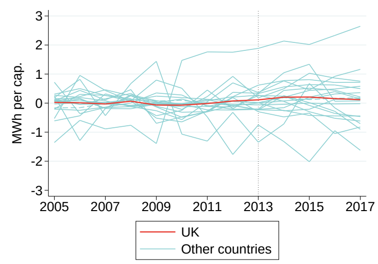

(b) Net imports per capita

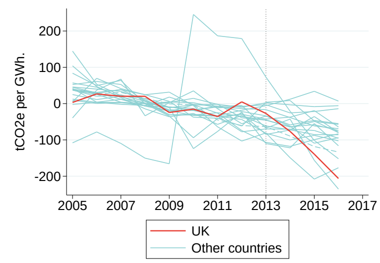

(c) Emissions per unit of electricity output

Note: The variables appearing on these two graphs were obtained by taking the difference between the original variable and the 2005-2012 average. The "Other countries" include all remaining EU countries except Romania, Bulgaria, Slovenia, Croatia, Malta, Cyprus, and Luxembourg.

Figure 5: UK power sector's input fuel mix in 2005, 2012 and 2017

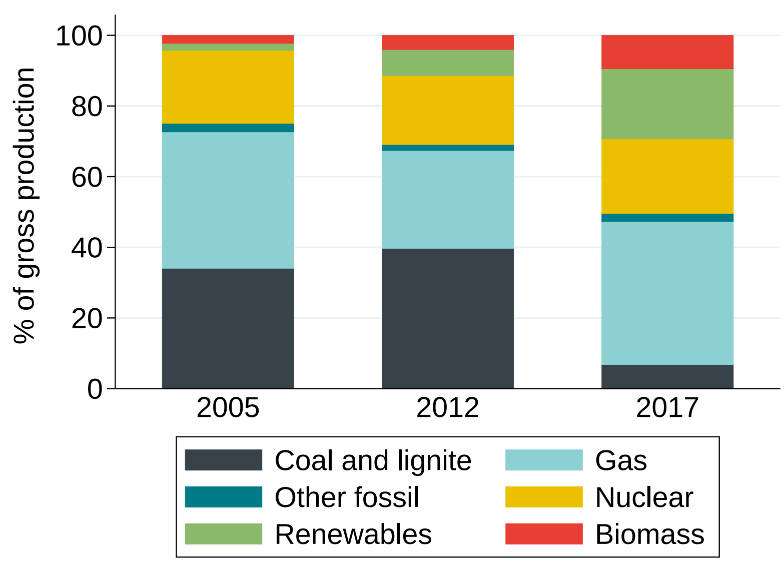

Note: Data come from Eurostat. Renewables include hydro, solar and wind

Figure 6: Power sector's input fuel mix in EU countries, 2012 and 2017

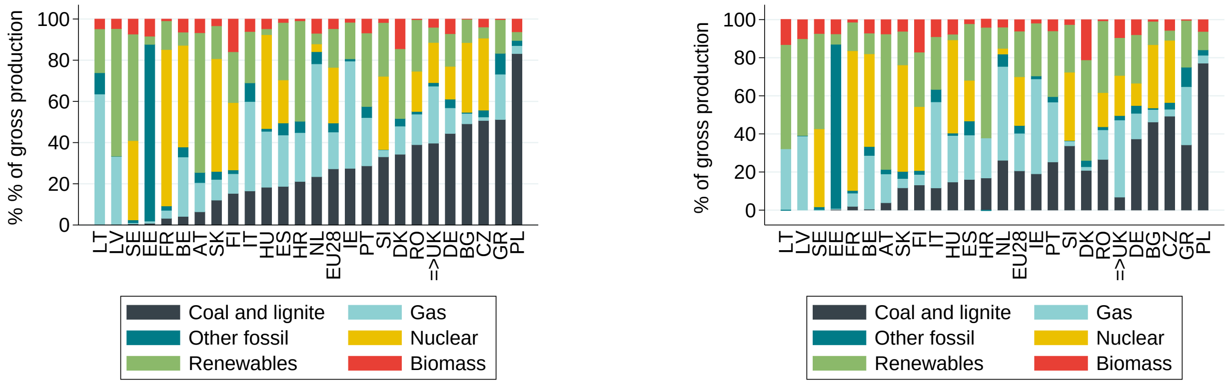

(a) 2012 (b) 2017

Note: EU countries are ranked by ascending order of the share of coal in electricity generation in 2012. The three smallest countries (Cyprus, Malta and Luxembourg) are not included, to improve readability. Data come from Eurostat. Renewables include hydro, solar and wind.

4.2 SCM method: Main result

The main result is visible on figure 7a, displaying the trajectory of per capita emissions for the UK and synthetic UK over the 2005-2017 period. In dotted line, it also shows the simple average of per capita emissions for the 15 countries of the donor pool. The synthetic UK is made of five countries: Ireland (51.3%), Slovakia (25.7%), the Netherlands (11.1%), Finland(6.3%) and Czech Republic (5.5%). The remaining potential control countries have a weight of 0. The large weight of Ireland does not come as a surprise: the two countries have close institutions and energy markets, and like the UK, Ireland has a substantial portfolio of coal- and gas-fired power plants. The UK and synthetic UK have a similar trajectory before the treatment occurs in 2013, with a Mean Squared Prediction Error (MSPE) of 0.010. It suggests that the synthetic UK after 2013 accurately replicates the evolution of per capita emissions in the UK power sector if the CPF had not been introduced. Compared to the simple average for the donor pool, the synthetic UK has a relatively close trajectory but higher per capita emissions. The period where the fit is less good is 2012, year where the UK emissions peak. It is also the year where the share of coal in the UK fuel input mix is the highest and the coal-to-gas price ratio at its lowest point since 2007 (BEIS, 2018). If power could easily be stored, the 2012 peak could also be interpreted as an anticipation effect of the CPF, which was announced in 2011. Coal-fired power plants would then have an interest to use their coal before it becomes highly taxed in 2012, and sell part of the electricity produced with that coal the subsequent years. Given that electricity cannot be stored and production has to match demand at every point in time, such a strategic behaviour is conditioned on the willingness of some less efficient plants to generate more electricity in 2012 and sell it potentially at a loss, compared to more efficient plants. It is reasonable to think that the power plants scheduled to close because of the LCP directive had an interest in adopting such behaviour, especially if they had excess coal stocks from previous coal purchasing contracts that they wanted to get rid off before being taxed. Anecdotally, official data on annual coal consumption and stocks by electricity generators indicate that coal stocks as a share of stocks and consumption are lower in 2012 compared to previous periods (20% vs 27% on average over 2005-2012), although the difference is not large. (BEIS (Department for Business, Energy & Industrial Strategy), 2019). To capture this potential anticipation effect, I also test a specification where I assume that the treatment started in 2011 rather than 2013 (see section 6.1).

Table 3 shows the relative weights of each predictor in the V matrix, as well as the average value of each predictor for the UK, synthetic UK, and average of the donor pool. The predictors with the highest weight are the coal-to-gas price ratio and its square, followed by the residual load per capita. This is not surprising given the aforementioned potential for fuel switching in the UK and the countries forming the synthetic UK. Together, these three variables make up 95% of the total weighting. The pre-treatment predictors' values in the synthetic UK are very close to the values for the actual UK, which is one of the identification assumptions of the SCM. In contrast, compared to the average country in the donor pool, the UK has significantly higher average per capita emissions prior to the introduction of the CPF, a higher amount of emissions from opted out plans, a lower coal-to-gas price ratio, a

Table 3: Predictors' weights and values for UK, synthetic UK and average of the donor pool

Variable	**Weight**	**UK**	**Synth. UK**	**Avg. Donor pool**
Per capita residual load	36.6%	4.04	4.05	3.07
Coal-gas price ratio	15.5%	0.52	0.51	0.71
Coal-gas price ratio squared	43.1%	0.27	0.28	1.26
Per capita opted-out emissions in 2009	0.05%	0.28	0.24	0.22
Lignite dummy	1.2%	0	0.06	0.2
Per cap. emissions 2005	1.2%	3.0	3.1	2.6
Per cap. emissions 2012	2.2%	2.6	2.4	2.1

Note: The weights of the predictors correspond to the diagonal coefficients of the V matrix minimizing the distance between the UK and synthetic UK prior to the introduction of the CPF in 2013. The predictor values are their true value for time-invariant predictors, their average across 2005-2012 for the number of degree days and their average across 2007-2012 for the coal-to-gas price ratio and its squared. See Appendix B for more details.

higher residual load per capita, and no lignite. These differences further justify using the Synthetic Control Method to build an accurate counterfactual UK power sector.

Figure 7b displays the gap between the synthetic and treated UK, representing the annual decrease in per capita tCO_2e resulting from the introduction of the CPF. The gap between the UK and synthetic UK widens significantly between 2014 and 2016, with UK emissions slightly higher than synthetic UK emissions in 2013 and the gap stabilizing in 2017. Such a path makes sense given the timing of the introduction of the CPF (April rather than January 2013), the strong increase in the CPS rate between 2013 and 2015, and the CPS freeze in 2015/2016.

I estimate the corresponding annual abatement for each year $t \in [2013, 2017]$ by multiplying the annual gap in per capita emissions by the UK annual total population. I then add up all annual abatements and find a total cumulated abatement of 185 millions of tCO_2e ($MtCO_2e$) over the 2013-2017 period. By 2017, emissions were lower by 49% (69 $MtCO_2e$) in the UK than in the synthetic UK.

Figure 7: UK and Synthetic UK per capita emissions

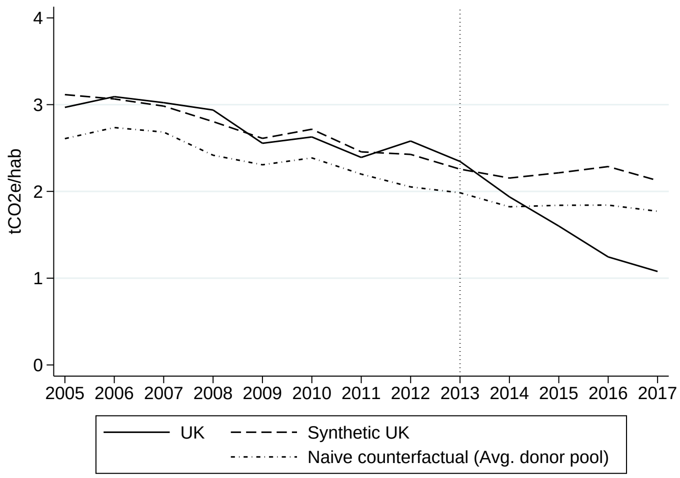

(a) Absolute per capita emissions for the UK, Synthetic UK and average of the donor pool

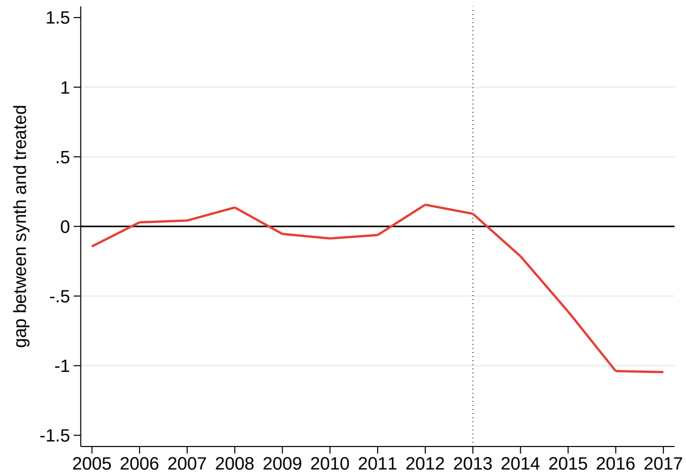

(b) Emission gap between treated and synthetic

Note: For each period, the variable of per capita emissions corresponds to the sum of CO_2e verified emissions from power installations subject to the EU ETS, divided by the average country population that year. The vertical line is set in 2013, date where the CPF is introduced. The synthetic UK is made of five countries: Ireland (51.3%), Slovakia (25.7%), the Netherlands (11.1%), Finland(6.3%) and Czech Republic (5.5%).

4.3 Lower bound removing emissions from potential confounders

Potential confounders An important assumption of my empirical strategy is that the difference in emissions observed between the UK and synthetic UK after 2013 can be imputed to the Carbon price floor only, not to other policies implemented in the UK and not elsewhere. As mentioned in section 2.3, two other policies may have contributed to the observed decrease in emissions in the UK: the European LCP Directive (LCPD), and the UK-specific support to biomass conversion. In the main specification, I control for the amount of emissions opted-out from the LCPD before the implementation of the CPF. I assume that once this is accounted for, any difference between the UK and Synthetic UK in the evolution of emissions from opted-out plants should be imputable to the CPF policy. For example, the CPF might affect the way in which the remaining operation hours of each opted-out plant has been spread over the 2005-2015 period; it may also influence the decision to shut down the plant rather than invest in low-carbon technologies that would bring these plant's emissions in line with the directive's requirements. Regarding the biomass conversion policy, I did not find evidence for similar policies in the rest of European countries. The CPF and support for biomass conversion are likely not to be independent from one another, as suggested in section 2.3. In the main specification, I make the somewhat strong but plausible assumption that the decision to subsidise biomass conversion was a consequence of the decision to implement the CPF. It means that UK plants would not have converted to biomass to a greater extent than non-UK plants in the absence of the CPF.

I now estimate a lower bound of the impact of the CPF where I deal with these two potential confounders by removing the emissions from plants concerned by the policies.

Emission decomposition Having detailed plant-level emission data allows for a decomposition of UK and synthetic emissions by source and category of plants. In Figure 8, I decompose emissions distinguishing between the decrease in emissions which I deem fully imputable to the CPF, and the decrease in emissions potentially linked to the LCPD or biomass conversion policy. The amount of emissions coming from installations having opted out from the LCP directive in 2005 is shown in light blue. The amount of emissions coming from plants having converted to biomass is shown in grey (only for the UK).

Therefore heterogeneity in pre-CPF response to the LCP directive is partially controlled for by the corresponding predictor (the value of these emissions in 2009). Nevertheless, over the whole 2005-2012 period, the amount of "opted out" emissions somewhat differs between UK and synthetic UK. But the most striking difference is the strong decrease of these emissions after 2012 in the UK and their relative stability in the synthetic UK. This difference suggests that the CPF may have intensified UK plants' response to the LCP opting-out regime and accelerated their closure. Such an interpretation would confirm a Guardian journalist's statement that "[UK coal-fired] Plants have closed in recent years as EU pollution standards started to bite, but it was increases in the UK's carbon tax that sealed their fate" (Vaughan, 2018). The figure also confirms that the 2012 emission peak experienced in the UK is almost entirely caused by these LCP opted-out plants, suggesting that they were willing to use their polluting inputs before the introduction of the tax. On the other hand, the emissions from plants converted to biomass following the UK specific policy represent a substantial amount of UK emissions on figure 8b.

Once the emissions associated to plants concerned by these two policies have been identified, the remaining emissions can be divided in two, according to whether they correspond to changes at the intensive or extensive margin. At the intensive margin, production can be reallocated to either more efficient plants using the same fuel or plants using a lower-emitting fuel; at the extensive margin the CPF can cause plants to enter or exit the market. The red area shows the evolution of emissions at the intensive margin and represent emissions from plants present in the dataset the whole period throughout. They contribute significantly to the overall decrease in emissions after 2013 in the UK while they are stable in the synthetic UK in the same period. The dark blue area shows the evolution of emissions at the extensive margin and represent emissions from plants either appearing in the data after 2005 or disappearing before 2017.

Figure 8: per capita CO_2e emissions by source, UK and synthetic UK

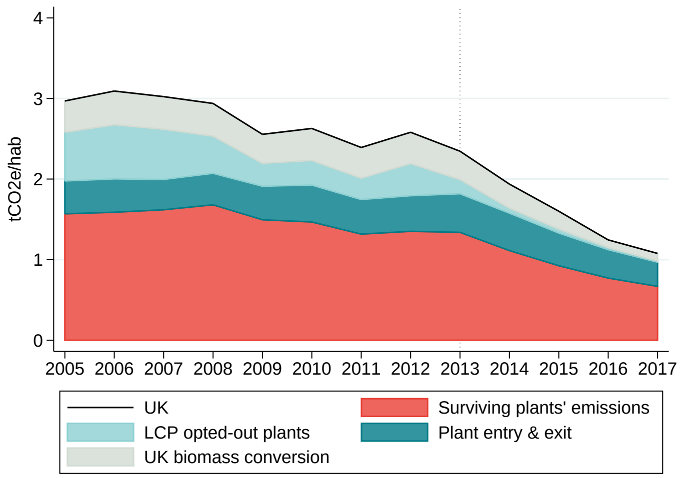

(a) UK

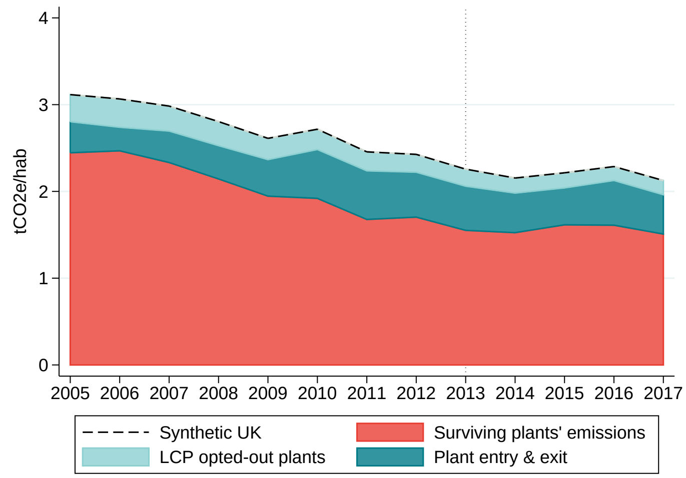

(b) Synthetic UK

Note: For each period, the variable of per capita emissions corresponds to the sum of CO_2e verified emissions from power installations subject to the EU ETS, divided by the average country population that year. The vertical line is set in 2013, date where the CPF is introduced. The synthetic UK is made of five countries: Ireland (51.3%), Slovakia (25.7%), the Netherlands (11.1%), Finland(6.3%) and Czech Republic (5.5%).

Lower bound Acknowledging that the present synthetic UK may not fully account for the role of the LCP opting-out regime and the UK biomass conversion policy, I estimate a rough lower bound of the impact by applying the synthetic control method on a modified per capita emissions value. In this modified outcome variable, CO2e emissions coming from opted out plants and plants having converted to biomass are removed. This comes down to keeping the dark blue and red areas of figure 8. The predictors used are slightly different: the lags of the outcome in 2005 and 2012 are taken at their modified value, a lag for 2010 is added (ensuring a good pre-treatment fit) and the amount of emissions from opted out plants is removed. I also calculate a second version of modified emissions, where I only remove emissions from UK plants having converted to biomass and keep the emissions from opted-out plants in both the UK and donor pool (as well as the corresponding predictor). The results are displayed in Figure 9. Since UK per capita emissions before 2013 are significantly lower in these two specifications, the synthetic UK is also made of different countries. For the specification without opted-out emissions and emissions from plants converted to biomass, the Synthetic

UK is made of the Netherlands (38.6%), Italy (35.4%), Slovakia (22.9%), Spain (1.7%) and Denmark (1.3%); For the specification without emissions from plants converted to biomass, it is made of Italy (65%), Ireland (15.8%), Slovakia (8.1%), the Netherlands (4.5%), Poland (4.2%) and Finland (2.4%).

The results go in the expected direction: the level of UK and synthetic UK per capita emissions are lower in these modified versions, and the gap between the two is also smaller than in the main specification. This confirms that part of the post-2013 difference in emissions is due to the change in emissions by power plants opted out from the LCP directive or converting to biomass (although again, these two policies are likely not independent from the CPF policy). Notably, removing emissions from both opted out plants and plants converting to biomass or only for plants converted to biomass does not change much the gap, suggesting that including the amount of emissions from opted out plants is a good enough way to control for the specific impact of opted-out plants on per capita emissions.

I compute again the equivalent annual abatement of tCO_2e for the specification without opted-out emissions and emissions from plants converted to biomass. Adding them up for the 2013-2017 period, I find a total cumulated abatement of 106 millions of tCO_2e, 79 $MtCO_2e$ less than in the main specification. Emissions are also lower by 41% in the UK than in the synthetic UK, 8 percentage points less than in the main specification. I consider this abatement of 106 $MtCO_2e$ as the lower bound of the impact of the CPF.

Figure 9: Synthetic control method excluding emissions from opted out plants and plants having converted to biomass

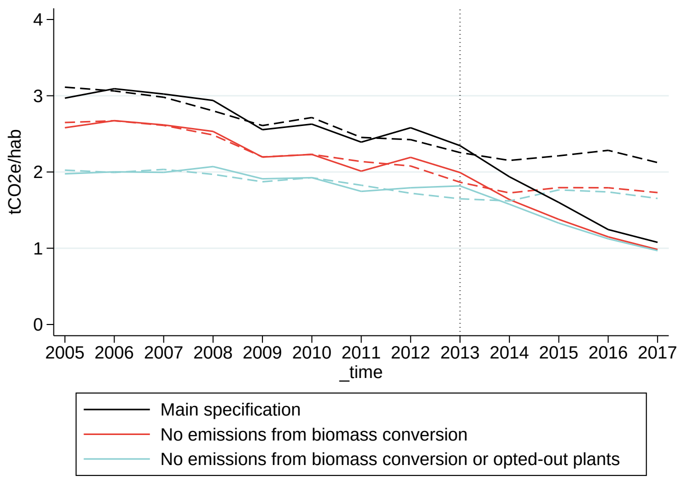

(a) Absolute per capita emissions

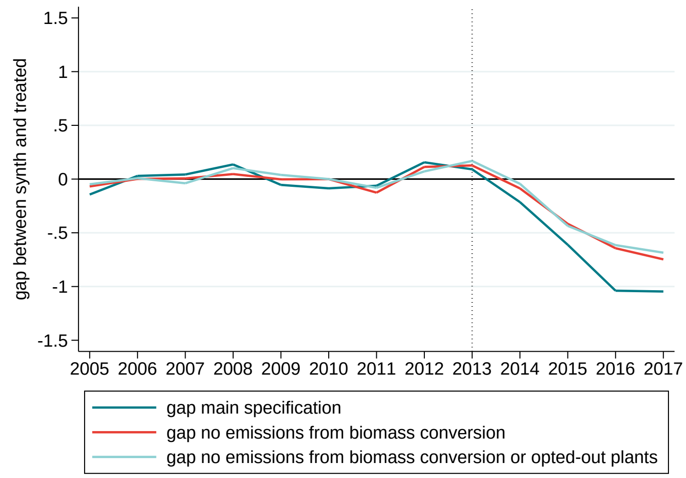

(b) Emission gap between treated and synthetic

Note: Each year, the variable of per capita emissions corresponds to the sum of CO_2e verified emissions from power installations subject to the EU ETS, divided by the average country population that year. The vertical line is set in 2013, the date for the start of the policy. The synthetic UK for the main specification is made of six countries: Ireland (53%), Slovakia (17.5%), Italy (15%) Finland(8.7%), Poland (5.7%), and Denmark (0.1%). For the specification without opted-out emissions and emissions from plants converted to biomass, the synthetic UK is made of The Netherlands (38.6%), Italy (35.4%), Slovakia (22.9%), Spain (1.7%) and Denmark (1.3%); for the specification without emissions from plants converted to biomass, it is made of Italy (65%), Ireland (15.8%), Slovakia (8.1%), the Netherlands (4.5%), Poland (4.2%) and Finland (2.4%).

4.4 Inference

In this section, I argue that my results are driven by the causal impact of the Carbon Price Floor by measuring (1) how likely it is to find an effect of the same magnitude as what I find when I apply the method before 2013 (“in time” placebo); (2) how likely it is to find an effect the magnitude from what I find when I apply the method to other countries (permutation test). These tests are run for the main specification[20].

[20]Placebo tests run for the specification without emissions from biomass conversion and LCP opted out plants provide similar results, although the magnitude of the difference between the UK and the other countries is smaller (results available upon request).

"In-time" placebo One way to check that the results observed are indeed caused by the 2013 policy is to assume that a similar policy was implemented at another date prior from 2013, apply the same method to generate a synthetic UK, and check that the UK and synthetic UK have similar per capita emissions before and after this artificial intervention date. Figure 10 shows what happens assuming that the CPF was implemented in 2010 rather than 2013. Once again, the synthetic UK closely resembles the UK emission trajectory before 2010. As expected, there is no significant gap between treated and synthetic UK in 2011 and 2012.

Figure 10: Gap between treated and synthetic UK when the CPF is assumed to be implemented in 2010

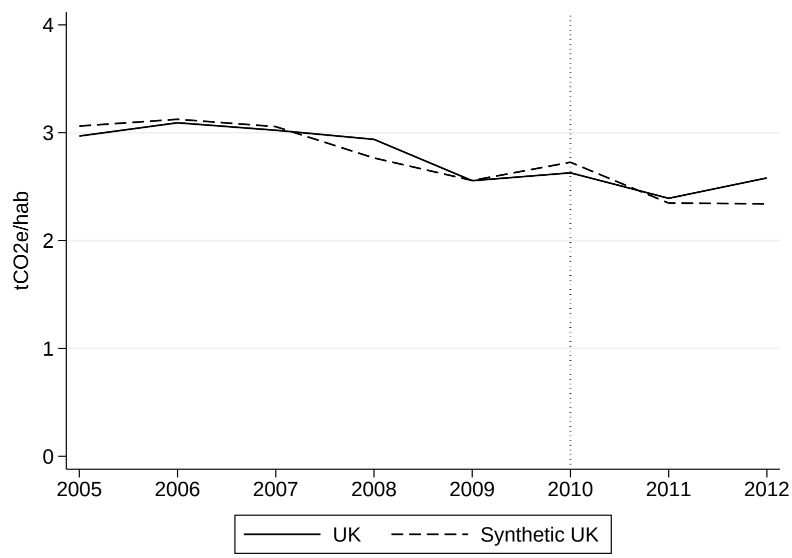

Note: Each year, the variable of per capita emissions corresponds to the sum of CO_2e verified emissions from power installations subject to the EU ETS, divided by the average country population that year. Here predictors are averaged over the 2005-2009 period, except for the coal-gas price ratio and its square, averaged over 2007-2010

Permutation test The permutation test consists in building a synthetic counterfactual for each country of the donor pool, and test whether any of them experiences a drop in per capita emissions (relative to its synthetic counterpart) as large as the UK. If it was the case for many of them, it would make it less plausible that the observed path for the UK and synthetic UK

Figure 11: Permutation test

(a) All countries

(b) Without the two countries with a large MSPE

Note: In both figures, Czech Republic and France are not included: for these countries it is impossible to find a stable diagonal V matrix. In the right-hand side the countries excluded are those with an MSPE 10 times higher than the UK (Denmark and Finland)

is indeed caused by the CPF. Figure 11a shows the gap between the treated and synthetic country for the UK and all the other countries in the donor pool, alternatively assumed to be the treated country. For the two countries with the highest and lowest per capita emissions (Czech Republic and France respectively), it is impossible to find a convex combination of "control" countries replicating the pre-2013 emissions, so these country are not included. For others, the pre-2013 fit is poor, with a high pre-treatment MSPE. Comparing the post-2013 emission gap between the UK and these countries is not very meaningful since the conditions for a good synthetic control are not met. Hence Figure 11b drops Denmark and Finland, the two countries with a pre-2013 MSPE 10 times greater than the UK[21]. On that figure, the UK clearly stands out as having the largest decrease in per capita emissions after 2013.

To illustrate the difference in the magnitude of pre- and post-2013 gap between the UK and the other permutations, I compute the ratio of post to pre-MSPE for all countries, where

[21] As mentioned in section 4.1, Denmark and Finland have a high variability in emissions, likely explained by the large inter-annual variations in renewable sources available for electricity generation.

we should observe an unusually high ratio for the UK. Figure 12 shows that the UK ratio is indeed the largest, and is about 5 times higher than Austria, the country with the second highest ratio. To make the parallel between this method and traditional hypothesis testing, we can interpret the result as the estimated probability to observe an effect as large as the one observed for the UK under a random permutation of the intervention on our data by dividing the countries having a higher ratio by the total number of countries(Abadie et al., 2010). In this case the UK has the highest ratio amongst the 14 countries, so the associated probability is $1/14 = 7.14\%$, the lowest probability one can hope to obtain with a sample of 14 countries.

Figure 12: Ratio of post to pre-MSPE

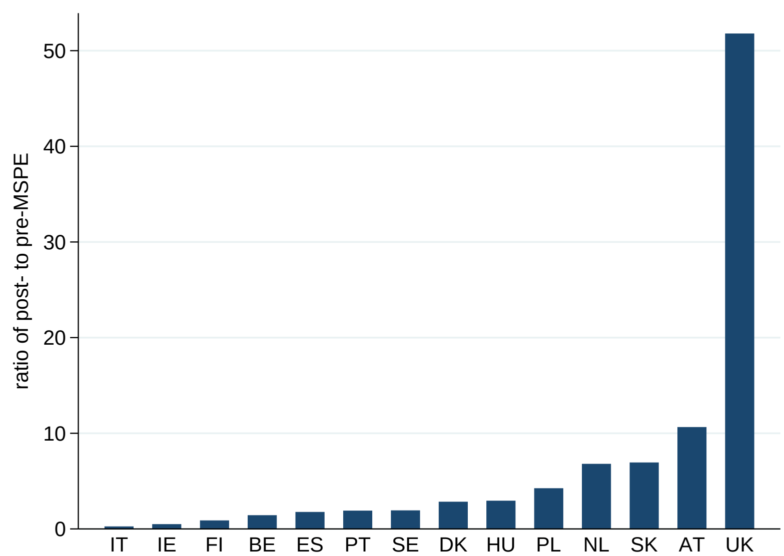

Note: Czech Republic and France are not included: for these countries it is impossible to find a convex combination of countries replicating pre-2013 emissions.

4.5 Discussion of the results

I find that the UK Carbon Price Floor resulted in a total cumulated abatement of between 106 and 185 MtCO_2e over the 2013-2017 period. The corresponding abatement for the 2013-2016 period lies between 61 MtCO_2e and 116 MtCO_2e. In contrast, Abrell et al. (2019) find a smaller total of 26.1 $MtCO_2e$ abated over the same period. The difference in our results is most likely explained by the difference in the data coverage and methodologies used. In terms of data coverage, Abrell et al. only consider coal- and gas-fired power plants producing electricity for the main grid, and running for the whole period considered, from 2009 to 2016. In contrast I consider all kind of power plants (using coal, gas, but also oil and petroleum gas), including those closing or opening during the period considered, and including those producing electricity for use on an industrial site. In terms of methodology, Abrell et al. look specifically at the impact of the CPF on fuel switching while I capture the other channels via which the CPF might impact emissions. For this reason, they actually conclude their paper saying that their estimate "should best be viewed as providing a lower-bound empirical estimate of the environmental effectiveness of the UK carbon tax" (Abrell et al. (2019), p41). Notably, the 61 $MtCO_2e$ estimate accounts for production re-allocation from coal-based generation to not only gas-based generation, but also to more efficient coal-fired power plants or to renewable sources of energy. The 116 $MtCO_2e$ estimate also accounts for plant closures and biomass conversion assumed to be caused by a mixture of the Carbon Price Floor and other policies. Applying the same emission decomposition as in section 4.3 to the profile of UK and synthetic UK emissions for the lower bound estimate, I can roughly decompose the abatement of 61 $MtCO_2e$ into two components. The first component represent the emission reduction from plants present in the market the whole period throughout (which should approximate the abatement found in Abrell et al), and the second component represents the emission reduction from plants exiting or entering the market during the period (which is not captured by Abrell et al). Simple back-of-the-envelope calculations[22] suggest that the first component corresponds to roughly 31 $MtCO_2e$ of abatement, which is broadly in line with Abrell et al.'s findings (31 $MtCO_2e$ is close to their estimate of 26.1 $MtCO_2e$ of fuel switching between plants running for the whole period).

[22]Drawing on the diff-in-diff methodology, I estimate for each component the double difference between the average pre-treatment and post-treatment emissions, between the UK and synthetic UK

5 Potential threats to the identification strategy

For the synthetic control method to identify the causal impact of the intervention, candidate units for the synthetic control group should not be affected by the intervention. In the case of the CPF, there are two channels via which the policy could spill over to other European countries' power sectors: directly via an increase in their net electricity exports to the UK; and indirectly via the effect of the CPF on demand for carbon allowances on the carbon market and the subsequent effect on EUA prices (the so-called waterbed effect). In this section I address these two concerns and provide empirical evidence that their role is limited.

5.1 Risk of spillover via increased electricity imports

In section 2.2, I already mentioned electricity net imports as one of the mechanisms through which the CPF could have lowered UK emissions. In section 4.1, I showed that UK electricity imports did not increase significantly after 2013 *compared to other European countries.* Here, I approach electricity imports as a specific challenge to my identification strategy, which is based on a comparison with other European countries: any increase in UK electricity imports caused by the CPF will increase the UK trading partners' CO_2e emissions. For the same reasons exposed in section 4.1, the magnitude of this increase should be small. However, the risk for increased net imports to contaminate the control unit might still be substantial given the composition of the synthetic UK: Ireland and the Netherlands, two of the three UK trading partners for electricity, have a combined weight of 62%. Absent the CPF, the UK might have exported more electricity to these two countries, which might have themselves produced less electricity and have lower emissions. If this scenario is true, the synthetic UK per capita emissions are overestimated. I address this concern by estimating a lower bound of Ireland's and the Netherlands' emissions in the absence of the CPF, and a corresponding lower bound for the emissions of the synthetic UK.

I proceed as follows: first, using Eurostat annual data on electricity trade, I calculate UK net electricity imports on the one hand[23], and Ireland's and the Netherlands' net electricity exports to the UK on the other hand[24]. Trade data for 2017 are not yet available on Eurostat, so I restrict this analysis to the 2013-2016 post-treatment period.

In a second step, I compute the increase in UK net imports and Ireland's and the Netherlands' net exports in the post-treatment period relative to the 2005-2012 mean. To obtain an upper bound of the spillover, I assume that all the post-2012 variation in UK net imports (respectively other countries' net exports) is caused by the CPF.

Third, I calculate the corresponding exported/imported emissions[25]. I do so multiplying the UK's net electricity imports/Ireland's and the Netherlands' net electricity exports with its emission intensity. For a given period, a country's emission intensity is defined as its total power sector CO_2e emissions divided by its gross electricity production. Finally, I define a modified per capita emission variable that re-allocates the post-treatment variation in emissions to the country where the electricity is consumed. For the UK, it corresponds to the UK observed power sector emissions *plus* its imported emissions. For Ireland and the Netherlands, it corresponds to their observed power sector emissions *minus* its exported emissions. For all other countries, the modified per capita emission variable stays the same.

Finally, I plot the UK and synthetic UK emission paths obtained with this modified variable.[26]

Figure 13 overlays the emission path and emission gap for 1) the UK and synthetic UK with the original per capita emission variable, and 2) the UK and synthetic UK with the modified per capita emission variable. As expected, the "modified" UK emission path is higher than the original one, but the difference between the two is very small. The "modified" synthetic UK emission path is actually higher than the original one in 2013 and 2014 because

[23]which only come from France, Ireland and the Netherlands due to the physical constraints exposed above

[24]Given that France is not part of the synthetic UK, the calculation is not done for France

[25]In terms of terminology, I consider that net electricity imports (respectively exports) correspond to an export (respectively imports) of emissions.

[26]The composition of the synthetic UK is the same since the only modified values, after 2013, do not enter the optimization program.

Figure 13: Synthetic control method accounting for imports spillovers

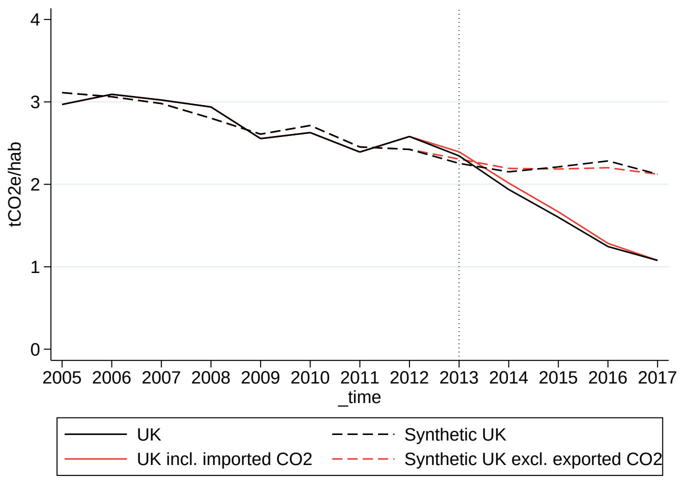

(a) Absolute per capita emissions

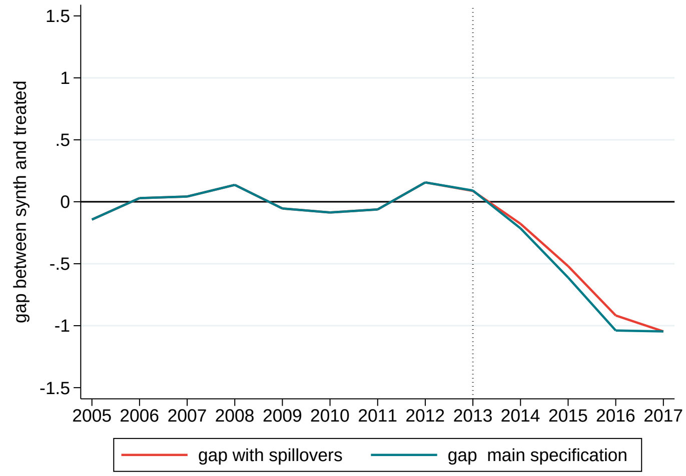

(b) Emission gap between treated and synthetic

Note: Each year, the variable of per capita emissions corresponds to the sum of CO_2e verified emissions from power installations subject to the EU ETS, divided by the average country population that year. The vertical line is set in 2013, date where the CPF is introduced. The synthetic UK is made of five countries: Ireland (51.3%), Slovakia (25.7%), the Netherlands (11.1%), Finland(6.3%) and Czech Republic (5.5%).

Ireland's net exports to the UK are lower in 2013 and 2014 compared to the average 2005-2012 period. It becomes lower in 2015 and 2016, reflecting the increase in Ireland's and the Netherlands' net exports to the UK.

The overall estimate of CO_2e abatement is smaller by only 16 MtCO_2e with this modified emission variable (total cumulated abatement of 169 MtCO_2e vs. 185 in the main specification). This result has two implications: first, it confirms that the risk of spillover is low and does not pose a severe threat to the identification strategy. Second, it is informative in itself and can be seen as the impact of the CPF on UK CO_2e emissions from a consumption-based rather than a production- based accounting perspective[27].

[27]This is the case since all emissions caused by electricity produced in other countries but consumed in the UK are re-allocated to the UK, and under the conservative assumption that the variation in electricity imports after 2013 is entirely imputable to the CPF.

5.2 Risk of spillover via a waterbed effect

The CPF might also spill over to other countries in an indirect way via the so-called waterbed effect. In general, the waterbed effect designates the mechanism via which, under a common emission cap, any emission reduction in a given country only leads to an emission increase elsewhere (Böhringer et al., 2008; Goulder and Stavins, 2011; IPCC et al., 2014; Perino, 2018). The risk of such a waterbed effect has been mentioned by some authors prior to the introduction of the CPF (Berghmans and Sartor, 2011; IPCC et al., 2014). Theoretically, it would work as follows: the CPF decreases demand for EUA permits from UK installations subject to a higher carbon price. On the EU ETS market as a whole, the shift to the left of the demand curve can only be compensated by a price decrease since the supply is fixed and perfectly inelastic (emission cap). Individual installations subject to the ETS but not the CPF can buy more allowances and emit more, with the aggregate emission reduction being unchanged.

A full empirical investigation of the existence of a waterbed effect in the case of the Carbon Price Floor is beyond the scope of this paper. However, for two reasons the risk of waterbed effect sould not be a severe threat to my identification.

First, I am only concerned with the risk of waterbed effect *within* the power sector, since my counterfactual is made of power installations. A waterbed effect materializing as an increase in CO_2e emissions from installations other than power generators would not be an issue for my estimation (although it could be an issue from the point of view of the overall efficiency of the policy).

Second, given the context of the EU ETS in the recent years, there are at least two reasons to think that the waterbed effect as a whole must have been limited. The first reason is that one can doubt whether the CPF induced a substantial decrease in EUA prices, given that UK power installations represent only a small share of total EU ETS emissions (only 8.8% in 2012). The second reason is that the waterbed effect depends on theoretical assumptions that may well not have hold at the period where the CPF was implemented. Indeed, the waterbed effect only happens if the cap constrains emissions. In the case of the EU ETS, the annual aggregate demand for allowances has been below the EU-wide cap since 2008, with a surplus of 1.8 bn allowances at the end of phase II (Ellerman et al., 2016). The likely reason why prices have not fallen to zero is that installations have been allowed to bank allowances for future use since phase II. Such increase in banking has been shown to be a rational behavior for agents facing a smoothly declining cap such as in the EU ETS (Ellerman et al., 2016) (unpublished paper). It means that even if the CPF induced a price decrease on the EU carbon market, triggering an increase in permit purchases, these permits were likely banked for future use rather than transformed in emissions from other countries' power installations. Finally, these additional banked allowances might well be cancelled through the recent set of ETS reforms, more precisely the Market Stability Reserve. The MSR is a new feature of the ETS implemented in 2019, that withdraws and inject allowances from and to the bank of unused allowances according to a quantitative formula. Crucially, allowances placed in the MSR can also be cancelled when they exceed a certain threshold. As evidenced by Quemin and Trotignon (2018), the MSR can act as a patch to curb some excess supply from past shocks or policies. With its ability to cancel allowances, the MSR can also retroactively and temporarily puncture any waterbed effect created by past overlapping policies as evidenced by Perino (2018): intuitively, the more allowances are banked for future use, the more the MSR is likely to exceed the threshold where it starts cancelling allowances.

In short, while I can not rule out the existence of a waterbed effect inducing non-UK power installations to emit more after 2013, the reasons exposed above suggest that this effect is probably very small relative to these installations' total emissions. I can actually calculate what the effect would be, assuming that 100% of the emission decrease in the UK translated in emission increases elsewhere. I assume that the observed decrease in UK emissions has translated in an increase by the same amount of emissions in other countries' power sector emissions. I estimate the amount of emissions attributable to the waterbed effect for each country in the donor pool, assuming that the share re-allocated depends

Figure 14: Synthetic control method assuming full waterbed effect

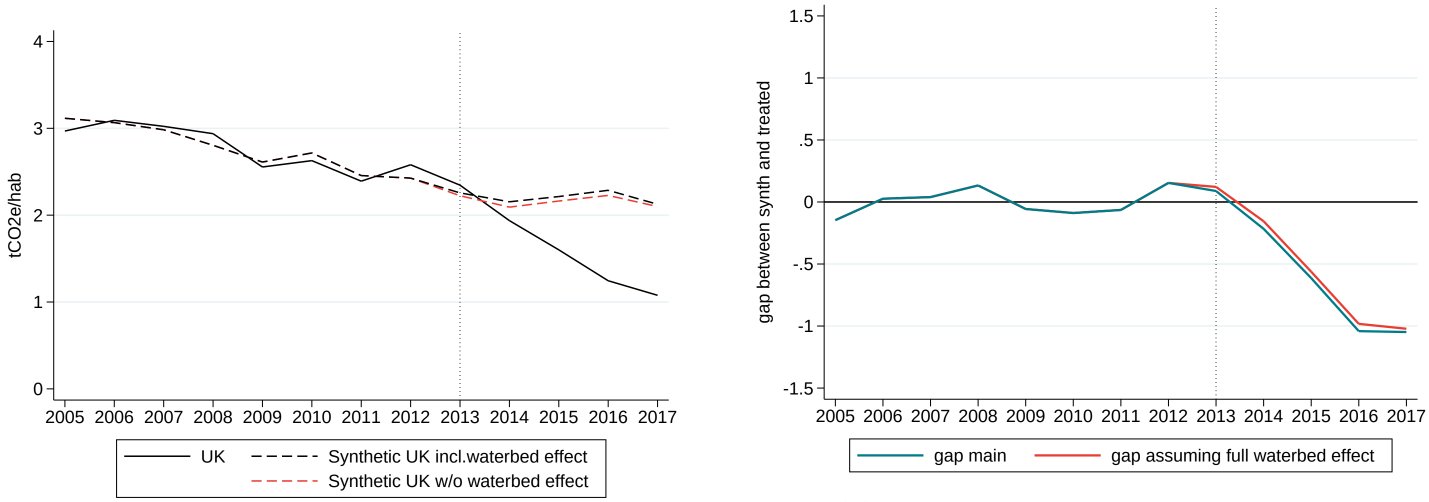

(a) Absolute per capita emissions

(b) Emission gap between treated and synthetic

Note: Each year, the variable of per capita emissions corresponds to the sum of CO_2e verified emissions from power installations subject to the EU ETS, divided by the average country population that year. The vertical line is set in 2013, date where the CPF is introduced. The synthetic UK is made of five countries: Ireland (51.3%), Slovakia (25.7%), the Netherlands (11.1%), Finland(6.3%) and Czech Republic (5.5%).

on each country's share in 2012 power sector ETS emissions (excluding the UK). I then re-calculate the emission path for the UK and synthetic UK [28]. The results are visible on figure 14. The overall estimate of CO_2e abatement is smaller by only 14 MtCO_2e (total cumulated abatement of 171 MtCO_2e vs. 185 in the main specification).

6 Sensitivity analysis

In this section I present three sensitivity analyses: the first assumes an anticipation effect in 2011, the second tests the sensitivity of the method to choosing alternative predictors, and the third tests the method when including two excluded countries back in the donor pool.

[28]Like in the preceding sensitivity analysis, the composition of the synthetic UK is the same since the values are only modified after 2013, a period not included in the optimization program

6.1 Synthetic Control Method assuming anticipation effect between 2011 and 2013

As explained in section 4.2, since the CPF was announced in March 2011, power plants may have anticipated its introduction in the two years between announcement and actual implementation by using their high-emitting inputs more in 2011 and 2012. In particular, plants subject to the LCP opt-out regime may have been willing to produce more these years than they would have done under normal conditions, in order to get rid of their coal stocks before coal-fired generation gets highly taxed. To account for this effect, I backdate the intervention to 2011 and apply again the SCM - making sure that the predictor values are averaged only until 2010. Figure 15a shows the resulting synthetic UK and new emission gap. The gap between synthetic and treated UK becomes larger in 2012 with the UK emitting more than the synthetic UK. This suggests an anticipation effect, presumably coming from the LCP opted-out plants (responsible for the 2012 emission peak, see Figure 8a). Figure 15b overlays the two emission gaps with treatment assumed to start in 2013 (main specification) and 2011 (this specification). Given the higher UK emissions in 2012 compared to the synthetic UK, the resulting total abatement is lower than in the main specification and amounts to 125 MtCO_2e. By 2017, emissions are still lower by 42% compared to the synthetic UK.

Figure 15: Synthetic control method when treatment is assumed to start in 2011, when the CPF was announced

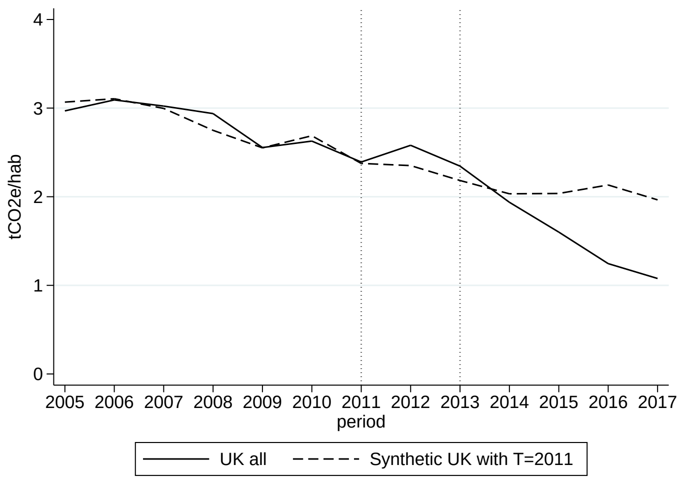

(a) Absolute per capita emissions

(b) Emission gap between treated and synthetic

Note: For each period, the gap is the difference between per capita emissions in the UK and synthetic UK. The vertical lines are set in 2011 and 2013, alternative dates assumed for the treatment start

6.2 Synthetic Control Method using alternative predictors

I also test the sensitivity of the method to the set of predictors used. I test four alternative specifications. The data sources used to build the new variables are described in appendix B. In the first alternative specification, I use the annual number of heating degree days instead of the per capita residual load. This variable approximates the demand for energy needed for heating. In the EU, it is measured as the number of days of the year where the average temperature is below a reference temperature of 15.5°C (under which energy for heating is needed), times the difference between this reference temperature and the temperature of the day. Compared with the average annual temperature used in other papers, this variable better accounts for the need for power generation related to low temperatures. Since gas- and coal-fired generation are often used to cover peak power demand, this variable is expected to correlate with the emission intensity of the power sector. In the second and third alternative specifications, I consider that the potential for fuel switching may depend on a country's total fossil fuel capacity and each fuel's (gas and coal) capacity. I therefore add as a predictor the annual country-level per capita combustion capacity (second specification), and the annual country-level per capita coal and gas capacity[29] (third specification). Finally, in the fourth alternative specification, I account for the trend in the available capacity for renewables, more specifically wind and solar, just before the introduction of the CPF. This trend can be considered as a proxy for the "business as usual" growth in renewable capacity and renewable production, which would occur absent the policy. I choose to focus on the two years between 2010 and 2012: 2012 is the last year of pre-treatment; 2010 is the year where the Europe 2020 strategy was implemented by the EU, setting a target for the share of renewables in final energy consumption. We might expect this announcement to be followed by a growth in renewable capacity in all European countries.

Figure 16 shows the resulting synthetic UK and emission gap for the main specification and for these four alternative specifications. The results are barely unchanged. The specification with the number of degree days is the one giving the lowest cumulated abatement, 15% lower than the main specification.

[29]For coal and gas capacity, I do not have the exact value and build a proxy. See B.

Figure 16: Sensitivity analysis: using alternative set of predictors

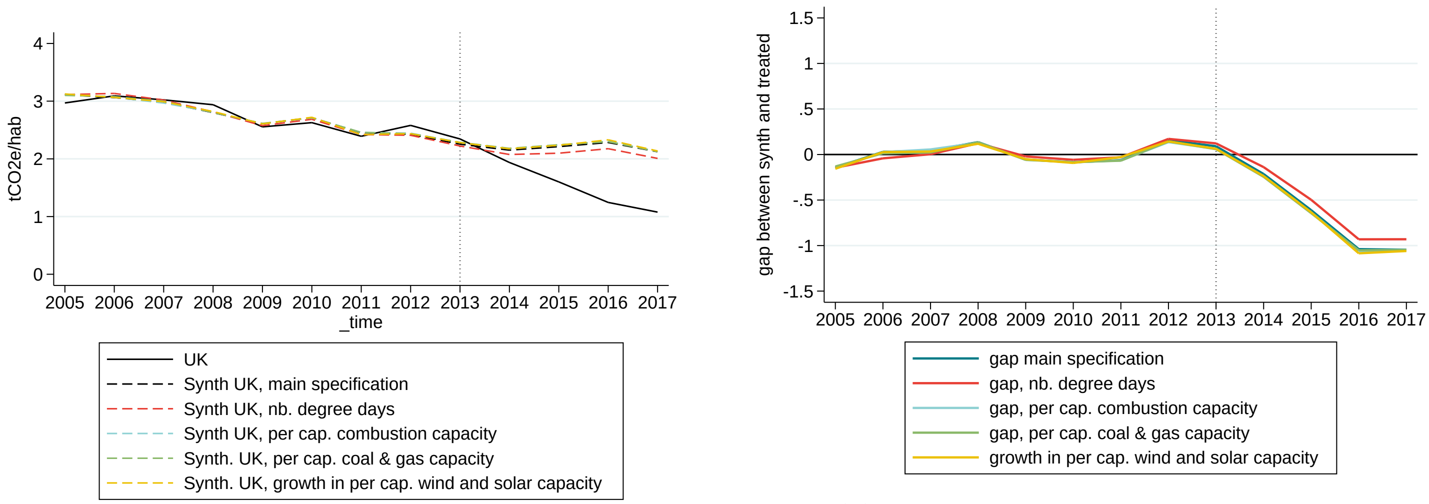

(a) Absolute per capita emissions

(b) Emission gap between treated and synthetic

Note: Each year, the variable of per capita emissions corresponds to the sum of CO_2e verified emissions from power installations subject to the EU ETS, divided by the average country population that year. The vertical line is set in 2013, the date for the start of the policy. The synthetic UK derived from the main specification is made of five countries: Ireland (51.3%), Slovakia (25.7%), the Netherlands (11.1%), Finland(6.3%) and Czech Republic (5.5%). The synthetic UK derived from the specification with the number of degree days is made of of six countries: Ireland (52.6%), Slovakia (24.3%), Italy (15%), Finland (8.7%), Poland (5.7%), and Denmark (0.1%). The synthetic UK derived from the specification with per capita combustion capacity is made of six countries: Ireland (53%), Slovakia (17.5%), the Netherlands (10.3%), Finland (5.8%), Poland (5%), and Czech Republic (2%). The synthetic UK derived from the specification with per capita gas and coal capacity is made of seven countries: Ireland(48.2%), Slovakia (20.9%), the Netherlands (12.4%), Finland(7.6%), Poland(6.5%), Spain (4.4%), and Denmark(0.1%). The synthetic UK derived from the specification with the growth in per capita solar and wind capacity is made of seven countries: Ireland(58%), France (16%), the Netherlands (10.4%), Finland(6.4%), Poland(5.6%), Slovakia (2.1%), and Czech Republic (1.5%).

6.3 Synthetic Control Method including Greece and Germany in the donor pool

Finally, I test whether including Greece and Germany in the donor pool changes the result. Figure 17 shows that the emission gap is very similar to the initial one.

Figure 17: Sensibility analysis: not excluding Greece and Germany in the donor pool

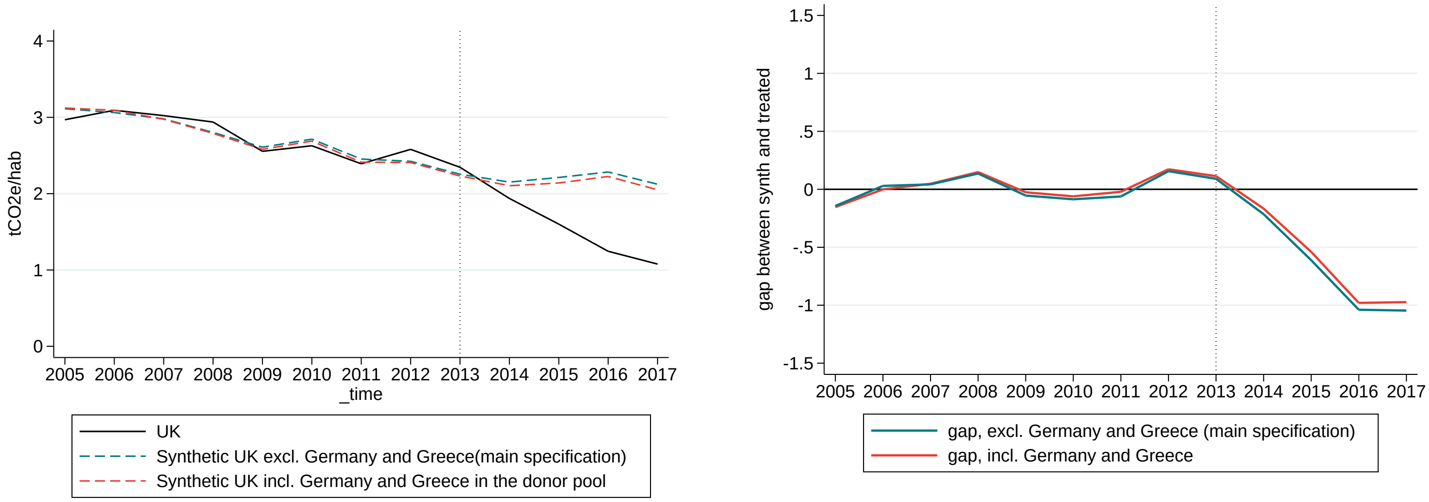

(a) Absolute per capita emissions

(b) Emission gap between treated and synthetic

Note: Each year, the variable of per capita emissions corresponds to the sum of CO_2e verified emissions from power installations subject to the EU ETS, divided by the average country population that year. The vertical line is set in 2013, the date for the start of the policy. The synthetic UK for the main specification is made of five countries: Ireland (51.3%), Slovakia (25.7%), the Netherlands (11.1%), Finland(6.3%) and Czech Republic (5.5%). The synthetic UK for the specification including Greece and Germany in the donor pool is made of six countries: Ireland (58%), Slovakia (25.5%), Finland(7.7%), the Netherlands (3.6%), Poland (3.1%), and Czech Republic (2%).

7 Conclusion

This paper estimates the impact of a carbon tax introduced in the UK power sector in 2013 by comparing the UK power sector's emission path to the power sector's emission path of a weighted combination of non-treated countries. I find that introducing a carbon price support to the EU ETS resulted in a decrease of CO_2e emissions of between 106 and 185 MtCO_2e in the UK over the 2013-2017 period. In 2017, emissions in the UK were almost halved compared to what would have happened absent the CPF. These results are robust to small variations in the donor pool and set of predictors used. Importantly, the estimates derived from the synthetic control method include all the different channels through which the CPF may have impacted per capita emissions: demand, trade and the domestic emission intensity of the power sector. Descriptive evidence suggest that the latter channel contributed the most to abatement.

While this paper adopts an estimation strategy that allows to take into account all these channels, it also has some limitations: first, it cannot estimate in a causal fashion the relative contributions of the three channels exposed in section 2.2. Second, having only one treated unit and fifteen units in the donor pool is not sufficient to build confidence intervals as done in Gobillon and Magnac (2016) and Isaksen (2018), and the only way to draw inference is via a permutation test using 14 countries. This limitation is however imposed by the nature of the policy evaluated - the policy applies only to one country, the natural donor pool would have a maximum of 27 countries (the number of EU countries without the UK).

From the mere point of view of abatement, the CPF policy can be considered as a successful policy: the estimated abatement represents between 45% and 79% of the abatement necessary to achieve targets of the 2nd carbon budget. It might be tempting to extrapolate from these results the likely impact of a carbon price floor introduced elsewhere. Yet it is important to keep in mind the specificities of the United Kingdom that may affect their external validity: the UK's relative isolation from other electricity markets limits the risk of carbon leakage, while countries in continental Europe are more interconnected. Furthermore, the context may have been particularly favourable to the introduction of a CPF in 2011, a period where some major coal-fired power plants were already deemed to shut down because of the LCP directive. In this context, companies operating multiple power plants may have been keen on investing in low-carbon generations and shown little resistance to the introduction of the CPF, more so than in the absence of such constraint. An evidence for this is that UK power companies supported the CPF (Hirst, 2018), in strong contrast with the opposition raised by a similar policy proposal in France (Newbery et al., 2018).

This work opens interesting avenues for future research related to the impact of carbon pricing in the power sector: first, it would be interesting to examine individual plants' behaviours and better understand the different decarbonisation channels from a plant's perspective. Second, this paper focuses on estimating the impact of the CPF on abatement, but other outcomes could be considered, such as the impact of the CPF on electricity prices and the resulting impact for electricity-intensive industries. Finally, the large decrease in coal burning driven by the CPF might have led to ancillary health benefits in the form of improved local air quality, which could be investigated.

References

Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association*, 105(490):493–505.

Abadie, A., Diamond, A., and Hainmueller, J. (2015). Comparative Politics and the Synthetic Control Method. *American Journal of Political Science*, 59(2):495–510.

Abadie, A. and Gardeazabal, J. (2003). The Economic Costs of Conflict: A Case Study of the Basque Country. *American Economic Review*, 93(1):113–132.

Abrell, J., Kosch, M., and Rausch, S. (2019). How Effective was the UK Carbon Tax?—A Machine Learning Approach to Policy Evaluation. *Unpublished Working Paper*.

Andersson, J. (2017). Cars, carbon taxes and CO2 emissions. *Centre for Climate Change Economics and Policy Working Paper*, (No. 238).

Ares, E. and Delebarre, J. (2016a). The Carbon Price Floor. Technical report, House of Commons Library.

Ares, E. and Delebarre, J. (2016b). The Carbon Price Floor. Technical report, House of Commons Library - Briefing Paper Number CBP05927.

BEIS (Department for Business, Energy & Industrial Strategy) (2019). UK Coal statistics.

Berghmans, N. and Alberola, E. (2013). The Power Sector in Phase 2 of the EU ETS: Fewer CO2 Emissions but just as much Coal. *Climate Report. paris CDC Climate research*, (42).

Berghmans, N. and Sartor, O. (2011). Carbon Price Flaw? The impact of the UK's CO2 price support on the EU ETS.

Böhringer, C., Koschel, H., and Moslener, U. (2008). Efficiency losses from overlapping regulation of EU carbon emissions. *Journal of Regulatory Economics*, 33(3):299–317.

Brown, G. (2017). British power generation achieves first ever coal-free day. *The Guardian*.

Ellerman, A. D., Marcantonini, C., and Zaklan, A. (2016). The European Union Emissions Trading System: Ten Years and Counting. *Review of Environmental Economics and Policy*, 10(1):89–107.

Ellerman, A. D. and McGuinness, M. (2008). CO2 Abatement in the UK Power Sector: Evidence from the EU ETS Trial Period. Working Paper.

European Commission (2001). Large Combustion Plants Directive (2001/80/EC).

Evans, S. (2019). Analysis: UK electricity generation in 2018 falls to lowest level since 1994.

Gloriant, S. (2018). Une Evaluation Quantifiee de la Taxe Carbone Française. *Information et Débats, Chaire Economie du Climat, Dauphine Université*.

Gobillon, L. and Magnac, T. (2016). Regional Policy Evaluation: Interactive Fixed Effects and Synthetic Controls. *The Review of Economics and Statistics*, 98(3):535–551.

Goulder, L. H. and Stavins, R. N. (2011). Challenges from State-Federal Interactions in US Climate Change Policy. *American Economic Review*, 101(3):253–257.

Hintermann, B. (2016). Pass-Through of CO2 Emission Costs to Hourly Electricity Prices in Germany. *Journal of the Association of Environmental and Resource Economists*, 3(4):857–891.

Hirst, D. (2018). Carbon Price Floor (CPF) and the price support mechanism. Technical Report Number 05927, House of Commons Library.

HM Revenue & Customs (2017). Excise Notice CCS 1/6: a guide to carbon price floor. Technical report.

IPCC (2014). Climate Change 2014: Impacts, Adaptation, and Vulnerability. Technical report.

IPCC, Somanathan, E., Sterner, T., Sugiyama, T., Chimanikire, D., Dubash, N. K., Essandoh-Yeddu, J., Fifita, S., Goulder, L., Jaffe, A., Labandeira, X., Managi, S., Mitchell, C., Montero, J. P., and Teng, F. (2014). National and Sub-national Policies and Institutions. In *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Edenhofer, O., R. Pichs-Madruga, Y. Sokona, E. Farahani, S. Kadner, K. Seyboth, A. Adler, I. Baum, S. Brunner, P. Eickemeier, B. Kriemann, J. Savolainen, S. Schlömer, C. von Stechow, T. Zwickel and J.C. Minx (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.*

Isaksen, E. T. (2018). Have International Pollution Protocols Made a Difference? *Grantham Research Institute on Climate Change and the Environment Working Paper No. 310,Available at SSRN 3317675*.

Kirat, D. and Ahamada, I. (2011). The impact of the European Union emission trading scheme on the electricity-generation sector. *Energy Economics*, 33(5):995–1003.

Le Hir, P. (2019). La France pourrait brûler encore un peu de charbon jusqu'en 2024. *Le Monde*.

Lee, K. and Melstrom, R. T. (2018). Evidence of increased electricity influx following the regional greenhouse gas initiative. *Energy Economics*, 76:127–135.

Martin, R., Muûls, M., and Wagner, U. J. (2016). The Impact of the European Union Emissions Trading Scheme on Regulated Firms: What Is the Evidence after Ten Years? *Review of Environmental Economics and Policy*, 10(1):129–148.

Newbery, D., Reiner, D., and Ritz, R. (2018). When is a carbon price floor desirable? Technical Report 1833, Faculty of Economics, University of Cambridge.

OFGEM (2013). Electricity interconnectors.

Perino, G. (2018). New EU ETS Phase 4 rules temporarily puncture waterbed. *Nature Climate Change*, 8:262–264.

Quemin, S. and Trotignon, R. (2018). Intertemporal Emissions Trading and Market Design: an Application to the EU ETS. *Centre for Climate Change Economics and Policy Working Paper*.

Rocha, M., Hare, B., Yanguas Parra, P., Roming, N., Ural, U., Ancygier, A., Cantzler, J., Sferra, F., Li, H., and Schaeffer, M. (2016). Implications of the Paris Agreement for Coal Use in the Power Sector. Technical report, Climate Analytics.

Sato, M., Singer, G., Dussaux, D., and Lovo, S. (2019). International and sectoral variation in industrial energy prices 1995–2015. *Energy Economics*, 78:235–258.

Staffell, I. (2017). Measuring the progress and impacts of decarbonising British electricity. *Energy Policy*, 102:463–475.

Van den Bergh, K. and Delarue, E. (2015). Quantifying CO2 abatement costs in the power sector. *Energy Policy*, 80:88–97.

Vaughan, A. (2018). UK government spells out plan to shut down coal plants. *The Guardian*.

World Bank and Ecofys (2018). State and Trends of Carbon Pricing 2018.

Zachmann, G. and Hirschhausen, C. (2008). First evidence of asymmetric cost pass-through of EU emissions allowances: Examining wholesale electricity prices in Germany. *Economics Letters*, 99(3):465–469.

Appendix A. Method used to identify power installations in the EUTL plant-level emission data

By default, the EUTL gives the broad sector category of all participants, but there is no specific category for the generation sector. The UK-based think-tank Sandbag kindly provided a database with total verified emissions data for 2008-2016 supplemented with a variable identifying all power plants. This identification has been performed internally by Sandbag based on (1) a file circulated by the European Commission in 2014, containing a list of individual participants and their associated NACE rev2 code (including a specific category for Production of electricity), and (2) some additional matching for participants joining the scheme after 2014, based on a top-down approach and desk-based research. For the verified emissions variable, the data provided by Sandbag are the same as the raw data retrieved from the EUTL.

Since Sandbag data is based on the 2008-2016 period, they do not provide information on whether or not an installation covered by the EU ETS is a power plant if the installation closed before 2008 or was opened in 2017. Using information on the date where a given installation first and last appeared in the data, I identified no installations opened in 2017, but 1,624 closed before 2008. I was able to match most of them to a dataset built by a consortium of European universities and hosted by the Florence School of Regulation, listing participating installations until 2013 with their NACE rev2 sectoral classification[30]. I identified 209 additional power plants with the NACE rev2 code. For the remaining installations, I flagged as power plants the most obvious ones, which name had "power station" in it or its equivalent in one of the European languages. After this additional matching, the power plant status is missing for only 457 installations, with only 113 of them having non-zero CO_2e emissions for at least one period.

[30]This dataset is called "Accounts to Firms Matching" and can be downloaded on this website: http://fsr.eui.eu/climate/ownership-links-enhanced-eutl-dataset-project/. The dataset was, according to the FSR website, part of a cooperation between several European Universities

Appendix B. Data sources for per capita emissions predictors

This appendix describes how the predictors used for the synthetic control method were built.

Coal-to-gas price ratio

This ratio is not readily available as a harmonized time series for all countries for the electricity generation sector. Sato et al. (2019) produce harmonized sector-specific energy prices and taxes data for 48 countries in the world, but their data stop in 2011 and do not include the power generation sector. I therefore build my own country-specific gas and coal price series for the generation sector by combining Eurostat and IEA sources.

Coal prices For coal, I use annual trade data for imported coal from Comext, the official EU trade statistics publicly available on Eurostat. More precisely, I combine volume and price data for all subcategories of coal that may be used for coal generation (i.e anthracite, codes HS27011110 and HS 27011190 before 2011 and HS 27011100 after 2011 ; bituminous coal, code HS 27011290 ; coal, code HS 27011900). I obtain average nominal unit prices for imported coal, expressed in Euro. Data for Denmark is only available for 2013 and 2014, so I fill the gaps by applying the growth rates from the closest non-missing data source, in this case the IEA nominal coal price index for industry.

The reason for choosing the Comext trade data over other available sources such as the IEA energy price statistics is their completeness for the time and geographic coverage that interests me. The Comext data also have drawbacks: they encompass coal purchased by all economic agents (rather than only power generators), only include imported coal (and not coal produced domestically), and ds not have taxes. For several reasons, I do not think that these drawbacks create a severe bias in the data series: first, I restrict the type of coal to the subcategories of coal most likely used for electricity generation, in order to be closer to the price paid for coal in the generation sector. Second, even if I do not have price data on coal produced domestically, marginal prices of traded and domestic coal should be about the same, if importing coal is efficient from a generator's perspective. Finally, not having taxes should not create too much of a distortion compared to the prices actually paid by generators. Indeed, the generation sector seems to benefit from fuel tax exemptions in most countries: a comparison between the IEA energy prices series in the electricity generation sector excluding and including taxes (for the 10 European countries where the data are available) indicates that no tax is applied to gas or coal used in this sector.

As a consistency check, I compare the obtained coal price series using trade data with IEA price series in the electricity generation sector for the few countries where both data are available. I find that the two series are very close in magnitude and time evolution.

Gas prices For gas, the main source of data is Eurostat. Eurostat provides average wholesale prices in euro paid for gas by industrial customers, for six consumption bands (I1..I6). Gas-fired power plants subject to the ETS are likely to fall amongst the largest bands. Using another database with information on the amount of input fuel used by large combustion plants (the LCP database), I find that electricity generators using gas as an input use on average 2,500,000 GJ, which falls into Eurostat's I5 consumption band. I use the price series in euro excluding VAT and other recoverable taxes and levies but including other taxes, to account for country-level differences in fuel taxation.

Where the data are incomplete or missing for this band, the values are imputed from other bands: when it is incomplete, the gaps are filled by imputing values from the consumption band below, I4 (for Ireland and Croatia), or the IEA gas price data (for Greece); when the value is completely missing for I5, the I4 values are used for all the years instead (for Luxembourg).

One drawback of this data source is that the categories (consumption bands) and methodology changed in 2007, which makes it difficult to build a consistent series of coal/gas price ratio before 2007. For this reason, I only use the 2007-2016 period for this variable.

I then combine coal and gas price series to build coal-to-gas price ratio for all European countries over the 2007-2016 period. Since coal Comext data are expressed in kg while Eurostat gas data are in kWh, coal data are converted into kWh in two steps: first, IEA country-specific kg-to-toe conversion factors for steam coal used in industry are applied[31]. For the few missing countries, the same conversion factor as the average for the non-missing countries is applied[32]. A common toe-to-kWh standard conversion factor ($1kWh = 8.6 * 10^{(-5)} toe$) is then applied to give coal price in euro per kWh.

Reassuringly, when comparing my price ratios obtained with my data sources with price ratios from national statistical institutes sources[33] or other sources like IEA, the ratio have about the same magnitude and path over time (with a margin of +/-10%).

Lignite resources

Data on lignite resources in Europe come from the industry association Euracoal (European Association for Coal and Lignite; Source: https://euracoal.eu/info/euracoal-eu-statistics/) I build a lignite dummy variable equal to 1 for countries with lignite resources greater than 0.5 Gt in 2012, and 0 otherwise. The variable is equal to one for Germany, Poland, Hungary, Greece, Czech Republic, and Bulgaria.

Residual load per capita

I define residual load as the difference between electrical energy available for Final consumption taken from Eurostat, and generation from renewables and nuclear power plants. The latter term is calculated as the sum of total net electricity production from six renewable sources (hydro, tide wave and ocean, solar PV, solar thermal and wind) and total net electricity production from nuclear power plants, including conventional plants, auto producers,

[31] I prefer to use the conversion factor for industry than for electricity production, since the data used for coal reflects coal imported for both industry and electricity generation

[32] the average does not include the few countries specified to use brown coal for electricity generation - which has a conversion factor half to three times as low as regular coal

[33] e.g, in the UK the Department for Business, Energy and Industrial Strategy publishes such data each quarter

and co-generation plants[34] This variable is then divided by the average population by country given by Eurostat.

Emissions from opted-out plants in 2009

The amount of CO_2e emissions coming from the installations who opt out from the LCP directive in 2005 comes from manually linking the EUTL dataset with data reported by the installations subject to the LCP directive (data version 3.1). The latter is publicly available on the European Environmental Agency's website. For each installation subject to the LCP directive, a variable indicates whether the installation is included in the opt-out regime. Since there is no common identifier between the EU ETS and LCP data, the 172 opted out installations located in the UK or in a country from the donor pool were manually matched to the EUTL emission data. Using information on the plant name and location, available in both datasets, I managed to match all LCP plants to a EU ETS installations except for a Finnish plant and a Polish one. The final variable is obtained by aggregating CO_2e emissions from these opted-out plants in the power sector at the country-level.

Number of degree days

I use the annual data on the number Degree days produced by Eurostat ("cooling and heating degree days by country - annual data" series).

Per capita capacity for combustible fuels, gas and coal

I use annual data on electricity production capacities for combustible fuels by technology and operator, produced by Eurostat. The technology variable has five values: Turbine, Combined Cycle, Steam, Internal Combustion and Other (For some countries, 100% of the capacity falls into the 'Other' category). The operator variable has two values: main producers and auto-producers. For the total capacity for combustible fuels, I use the total, including both main producers and auto-producers. Capacities distinct for coal and gas are not readily available. Instead I allocate the different technologies to coal and gas based on which technology dominates for each of the two fuels. 'Turbine' and 'Combined Cycle' is assumed to be gas capacity, while 'Steam' and 'Other' is assumed to be coal capacity. The technology 'Internal Combustion' is assumed to be mainly for oil and is not included.

Growth in per capita wind and solar capacity

I use annual data on Net electrical maximum capacity for wind and Net electrical maximum capacity for solar, produced by Eurostat. I add up the two and compute the average annual growth rate between 2010, the year where the Europe 2020 strategy was adopted (including the target of increasing the share of renewable energy in final energy consumption to 20% by 2020), and 2012.

[34] geothermal, biomass and waste are not included since they are available on demand

Appendix C. Common support for the distribution of predictors for the UK and countries from the donor pool

Figure 18: Distribution of characteristics for the UK (in red) and donor pool (in grey)

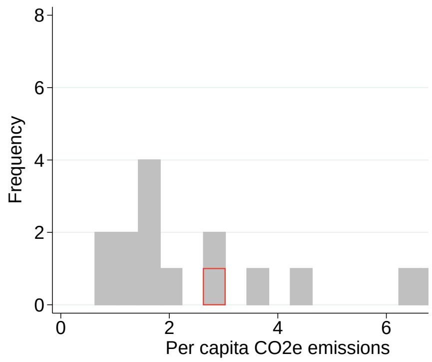

(a) Per capita CO_2e emissions

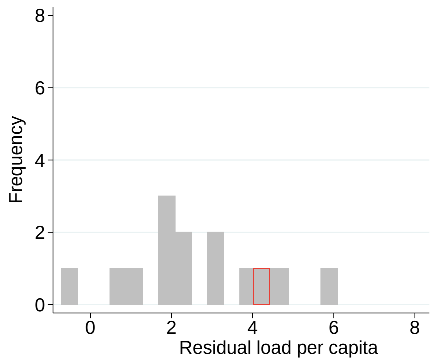

(b) Residual load per capita

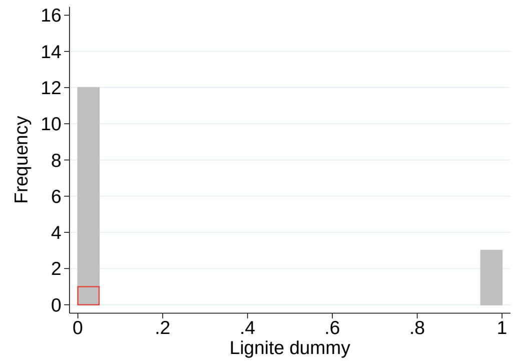

(c) Lignite dummy

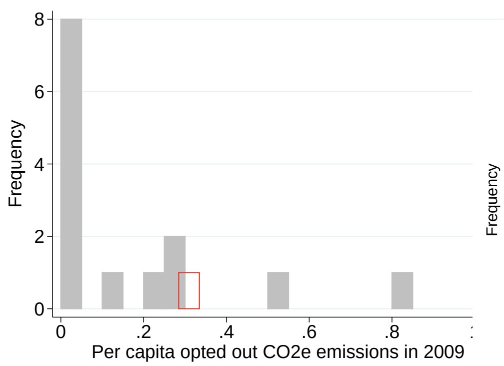

(d) Opted-out emissions per capita (2009)

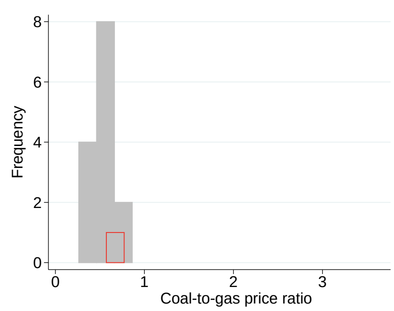

(e) Coal-to-gas price ratio

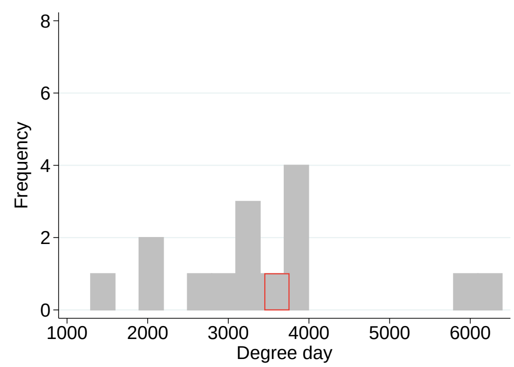

(f) Number of degree days

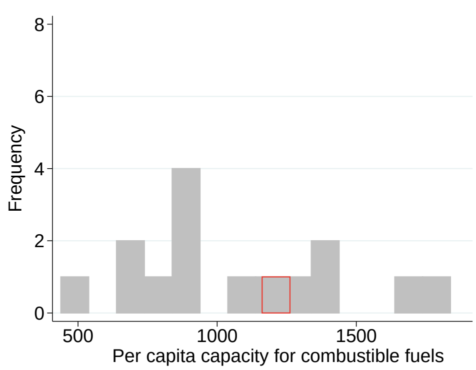

(g) Per capita capacity for combustible fuels

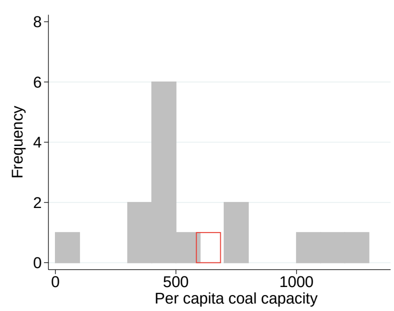

(h) Per capita coal capacity

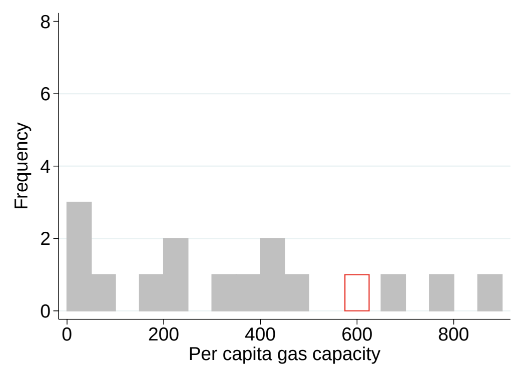

(i) Per capita gas capacity

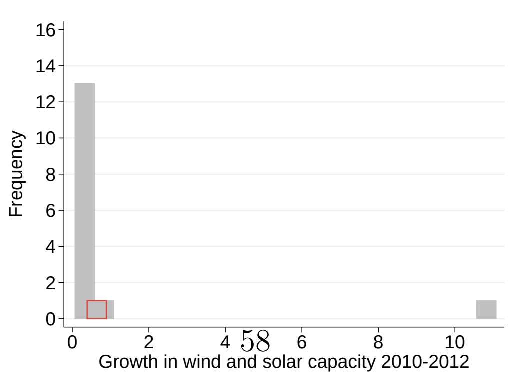

(j) Average annual growth in solar and wind capacity