

Adopt or innovate: Understanding technological responses to cap-and-trade

Raphael Calel

May 2018

Centre for Climate Change Economics
and Policy Working Paper No. 332
ISSN 2515-5709 (Online)

Grantham Research Institute on
Climate Change and the Environment
Working Paper No. 296
ISSN 2515-5717 (Online)

The Centre for Climate Change Economics and Policy (CCCEP) was established by the University of Leeds and the London School of Economics and Political Science in 2008 to advance public and private action on climate change through innovative, rigorous research. The Centre is funded by the UK Economic and Social Research Council. Its second phase started in 2013 and there are five integrated research themes:

1. Understanding green growth and climate-compatible development
2. Advancing climate finance and investment
3. Evaluating the performance of climate policies
4. Managing climate risks and uncertainties and strengthening climate services
5. Enabling rapid transitions in mitigation and adaptation

More information about the Centre for Climate Change Economics and Policy can be found at: www.cccep.ac.uk.

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training. The Institute is funded by the Grantham Foundation for the Protection of the Environment and the Global Green Growth Institute. It has six research themes:

1. Sustainable development
2. Finance, investment and insurance
3. Changing behaviours
4. Growth and innovation
5. Policy design and evaluation
6. Governance and legislation

More information about the Grantham Research Institute on Climate Change and the Environment can be found at: www.lse.ac.uk/grantham.

This working paper is intended to stimulate discussion within the research community and among users of research, and its content may have been submitted for publication in academic journals. It has been reviewed by at least one internal referee before publication. The views expressed in this paper represent those of the authors and do not necessarily represent those of the host institutions or funders.

Adopt or Innovate:

Understanding technological responses to cap-and-trade

Raphael Calel*^{1,2}

¹McCourt School of Public Policy, Georgetown University

²Grantham Research Institute on Climate Change and the Environment, London School of Economics

Abstract

Environmental regulations have consistently been found to spur innovation in ‘clean’ technologies, with one significant exception. Past cap-and-trade programs have encouraged adoption of existing pollution control technologies, but had little effect on innovation. Several explanations have been offered, including secondary market failures and a lack of polluter sophistication. In this paper I argue that it likely has more to do with the state of the technologies. Using a newly constructed panel of British companies, I show that the European carbon market—the world’s largest cap-and-trade program—has, contrary to past experience, encouraged innovation rather than adoption. I discuss how these contrasting findings can be reconciled, and the implications for planned reforms.

*I am grateful to George Akerlof, Robert Akerlof, Antoine Dechezleprêtre, Corado Di Maria, Sam Fankhauser, Meredith Fowle, Kenneth Gillingham, Cameron Hepburn, Michael Hoel, Harry Holzer, Arik Levinson, Carmen Marchiori, Jeffrey Mayer, and Karsten Neuhoff for their comments on earlier drafts of this paper, as well as seminar participants at Yale University, Georgetown University, and the CESifo Area Conference on Energy and Climate Economics. I owe special thanks to Antoine Dechezleprêtre for assisting in the compilation of patent data, and to David Popp for sharing files allowing me to partially replicate his work. This research would not have been possible without the assistance of the UK Environment Agency and Data Services, as well as funding provided by the UK’s Economic and Social Research Council (ESRC), the Jan Wallander and Tom Hedelius Foundation, and the Ciriacy-Wantrup Fellowship. This work contains statistical data from the Office of National Statistics which is Crown copy-right and reproduced with the permission of the Controller of HMSO and the Queens Printer for Scotland. The UK Data Service agrees that the attached outputs are non-disclosive, and cannot be used to identify a person or organisation. The original data creators, depositors or copyright holders, the funders of the Data Collections (if different) and the UK Data Service at the UK Data Archive bear no responsibility for the further analysis or interpretation of the data. This work uses research datasets which may not exactly reproduce National Statistics aggregates. A full replication archive can be accessed through the UK Data Service, with their permission. To learn more, go to <https://www.ukdataservice.ac.uk/get-data/how-to-access/accessecurelab>.

1 Introduction

The Paris Agreement began a new era of climate regulation, with over 100 countries committing themselves to putting a price on carbon. Cap-and-trade programs in Europe, North America, Asia, and Oceania already regulate 20% of global greenhouse gas emissions, and this figure may more than double once the Paris commitments are fully implemented (World Bank; Ecofys, 2017).

But will companies respond to these cap-and-trade regulations in all the ways that policy makers hope? One important motivation for these programs has been their potential for spurring low-carbon innovation (European Commission, 2005). A recent roadmap for decarbonisation suggests that R&D for clean technologies needs to increase by an order of magnitude by 2030 (Rockström et al., 2017), and indeed, the ‘induced innovation hypothesis’ suggests that environmental regulations, especially market-based regulations like cap-and-trade, will encourage innovation. Yet when cap-and-trade programs have been used in the past—to regulate SO₂, NO_x, lead, and ozone-depleting substances—they achieved emissions reductions mainly by encouraging adoption of off-the-shelf abatement solutions, without significant development of new technologies. Such an emphasis on adoption, coupled with a lack of innovation, would present a serious challenge if it were repeated with carbon markets.

This has generated interest in the technological responses to early carbon market experiments like the European Union’s Emissions Trading System (EU ETS), which has been in operation since 2005 (European Commission, 2005; Schleich and Betz, 2005; Gagelmann and Frondel, 2005; Grubb et al., 2005). But no clear answer has yet emerged from the young literature on the EU ETS. I present new evidence that, contrary to the pattern observed under cap-and-trade programs in the past, the EU ETS: (1) likely *has not* encouraged widespread adoption of off-the-shelf low-carbon technologies, but (2) it *has* encouraged low-carbon patenting and R&D spending. I argue that these findings can be reconciled with responses to past cap-and-trade program by considering how a price on emissions interacts with the set of abatement technologies.

I have assembled a new data set on British companies that links administrative data with business surveys and regulatory records. This provides the most comprehensive database yet on company-level CO₂-intensity, low-carbon patenting, and low-carbon R&D expenditures. By measuring both adoption and innovation together, I am able to describe company responses in greater detail than previous studies. These new data also allow me,

for the first time, to investigate a number of alternative response mechanisms—other than ‘induced innovation’—that have been proposed in reaction to earlier research.

To identify the effects of the EU ETS, I exploit two features of its implementation in Britain. First, only the companies that operate at least one ‘large’ plant are required to comply with the EU ETS, even though the *companies* that operate smaller plants may be just as large, just as polluting, and just as innovative. This means that I can in principle recover a quasi-experimental estimate of the EU ETS’s effects on adoption and innovation by comparing initially similar companies that have recently become subject to different regulatory regimes (List et al., 2003; Greenstone and Gayer, 2009). Second, Britain phased in the EU ETS over a few years, while phasing out a number of domestic climate change regulations (Bowen and Rydge, 2011). The transition was implemented in such a way that companies that were previously subject to the same regulations became subject to different ones. This provides an opportunity to disentangle the EU ETS’s effects from the effects of other carbon pricing policies, something that has been impossible elsewhere.

The lack of innovation in past cap-and-trade programs, and its presence in the EU ETS, is not as surprising as it might first appear. Past cap-and-trade programs were implemented against a backdrop of relatively mature pollution control technologies, which made adoption an attractive strategy. But carbon markets are now being rolled out in a different technological landscape. The current state of low-carbon technologies makes innovation more attractive today. I argue that the empirical findings from past and present cap-and-trade programs can be reconciled once we take account of how the technology sets differ.

2 Definitions and measurement

The ‘induced innovation’ hypothesis is typically described very simply in the empirical literature: a rise in the relative price for a given factor of production spurs greater innovation aimed at economising on the use of that factor. Following Porter’s (1991) environmental reframing of this hypothesis, a considerable theoretical literature has emerged on how environmental regulations affect technological change (Palmer et al., 1995; Ambec et al., 2013), especially in the context of climate change mitigation (Acemoglu et al., 2012; Aghion et al., 2016). But even though there is now a substantial literature examining the drivers

of technology adoption¹ and innovation, economists often disagree about what precisely constitutes each one. The best I can do here is to explain how I use these terms.

Adoption

The envelope of currently available pollution control solutions are summarised by the marginal abatement cost curve (shown in figure 1). It reflects the cost of abating emissions by installing off-the-shelf abatement technologies, as well as from using well-known pollution management strategies like fuel substitution. We can even think of output reduction as just another management strategy, although it is usually so expensive at the margin that it is not relevant for cap-and-trade programs in practice. Adoption moves a company *along* the marginal abatement cost curve; after adopting the cheapest abatement technology, it is left with options with a higher marginal abatement cost.

A cap-and-trade program encourages adoption by raising the opportunity cost of emissions for regulated companies. The company that faces a marginal emissions price p , and would therefore have to pay pE_0 , stands to save an amount equal to the shaded area **A** in figure 1 if it adopts the cheaper abatement options instead. But the policy provides no such incentive for companies that are not subject to the cap. Even if the emissions price is passed on to them in the form of higher product prices, it would create an incentive for unregulated companies to adopt technologies that rely less on pollution-intensive intermediate goods, but it would not affect their incentive to adopt technologies that reduce their own direct emissions.

A second effect, on the cost-side, could further widen the adoption-gap between regulated and unregulated companies. If emissions allowances are distributed to polluters free of charge (as they often are), this can create windfall profits. This windfall reduces the need to raise external funds to finance technology adoption, and can therefore lower the marginal cost of capital (Cramton and Kerr, 2002).

When the set of germane abatement strategies is small, it may be possible to observe adoption directly. More generally, though, adoption can be inferred from its consequences; companies are able to produce more output per unit of emissions. The absolute level of emissions is likely to be a poor indicator of adoption because output generally rises and falls for reasons unrelated to adoption decisions, such as the business cycle. Emissions-intensity

¹The term ‘diffusion’ is sometimes used, when the unit of analysis is the technology rather than the technology user.

will better proxy adoption decisions as long as emissions are locally proportional to output.

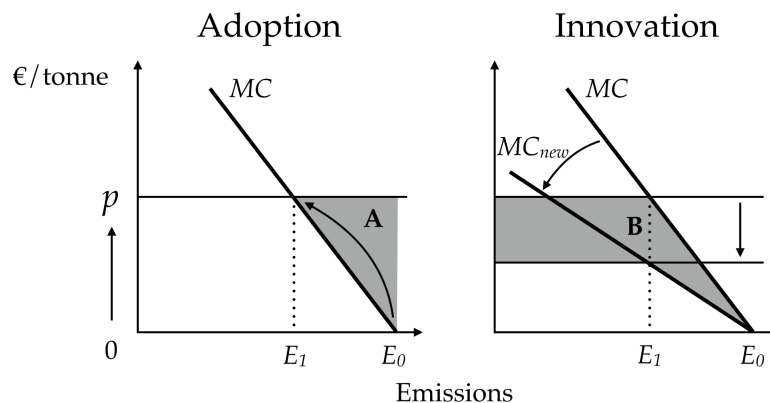


Figure 1: Incentives for adoption and innovation. In the absence of a regulatory constraint, a company will produce emissions E up to point where the marginal benefit from emissions goes to zero, E_0 . The curve labeled MC shows the marginal cost of reducing emissions below E_0 . A company that becomes regulated under a cap-and-trade program is faced with paying p to obtain allowances at the margin, but can reduce its compliance cost by adopting all abatement technologies with a marginal cost below the allowance price p . The shaded areas **A** shows the gain from doing so. Innovation increases the availability of abatement technologies, illustrated as a pivot of the MC curve in the right-hand panel (formally, a continuum of abatement technologies with uniformly distributed marginal abatement costs on $[0, a]$ gives rise to a linear marginal abatement cost curve, where the slope is decreasing in the availability of abatement technologies). Since the technology frontier shifts for all companies, innovation reduces aggregate demand for allowances and the allowance price p falls. The shaded area **B** shows the reduction in a regulated company's compliance cost that result from this innovation. For a standard discussion of the incentives to adopt and innovate under cap-and-trade, see Requate (2005).

Innovation

Innovation expands the technology set and creates opportunities for greater abatement at a given marginal cost. This *shifts* the marginal abatement cost curve inwards (figure 1). A company regulated under a cap-and-trade program can expect to earn a surplus **B** if it innovates, which comes partly from the lower abatement costs and partly from the fall in the emissions price that occurs when the new technology comes into general use.

Innovation is difficult to distill into a single metric, which means there is an advantage to being able to look at multiple measures together (Martin, 1996; Hagedoorn and Cloudt, 2003; Lanjouw and Schankerman, 2004; Arvanitis et al., 2010). I use both patenting and R&D spending to try to get a clearer empirical signal. Patenting has been used extensively as a measure of innovation output (Popp, 2002, 2006; Johnstone et al., 2010), while R&D expenditures are used to measure innovation inputs (Chavas et al., 1997; Hirshleifer et al., 2013; Arora et al., 2017). Thanks to efforts undertaken by the European Patent Office

and the former UK Department of Energy and Climate Change, I can also separately identify patents and R&D spending that concern low-carbon technologies.² These measures cannot fully capture all aspects of innovation, of course, but they provide useful proxies for companies' efforts and success in expanding the technology set (see OECD (2009) and OECD (2015) for surveys).

The distinction between regulated and unregulated innovators is an important one, but less straightforward than it is for adoption. A regulated innovator will not only earn the surplus \mathbf{B} , but also earn some average fraction of the surplus $\beta \in [0, 1]$ from other polluters that purchase or license their new technology. A standard way to model the appropriability problem is to assume that β is strictly less than one. For a cap-and-trade program that covers N companies, then, the additional rent will be $\beta(N - 1)\mathbf{B}$. The total gain from innovation for a regulated company is therefore $S_R = \mathbf{B}(1 + \beta(N - 1))$. The innovation surplus is subtly different for a company that isn't itself regulated under the cap. Because it does not stand to benefit directly from using the technology, its innovation surplus is only $S_U = \mathbf{B}\beta N$, which is strictly less than S_R for any $\beta < 1$.

This stylized framework highlights two principled reasons why cap-and-trade programs would create greater incentives for innovation among regulated companies: (1) $\beta < 1$ and (2) systematic differences in estimates of \mathbf{B} . The first condition, that $\beta < 1$, follows from the fact that an innovator cannot completely appropriate the gains from a new technology (Milliman and Prince, 1989; Fischer et al., 2003). Transactions costs, asymmetric information, less-than-total market power, and the fact that some of the gains from innovation come about indirectly through falling allowance prices, all dissipate the rents that an inventor can extract. The appropriability problem might even be particularly severe for clean technologies, where technological spillovers are typically larger (Dechezlepretre et al., 2014). Moreover, introducing a new abatement technology into a plant can require substantial in-house development, which a third-party innovator may not be able to provide

²The European Patent Office categorises all patents according to the European patent classification (ECLA), which includes a recently developed class pertaining to “technologies or applications for mitigation or adaptation against climate change”, or “low-carbon technologies” for short, also known as the “Y02” class. It is the result of an unprecedented effort by patent examiners specialised in each technology, with the help of external experts, to develop a tagging system for all patents ever filed at the EPO that are related to climate change mitigation. The Y02 class provides the most accurate tagging of climate change mitigation patents available today and is becoming the international standard for clean innovation studies. Importantly, the Y02 class is consistently applied to patents filed both before and after the EU ETS was introduced. See Veefkind et al. (2012) for more details. For more details on DECC's low emissions R&D survey, see Department of Energy & Climate Change and Office for National Statistics (2012).

effectively. If only the ‘research’ portion of \mathbf{B} is appropriable but not the ‘development’ portion, this further reduces β . This would be consistent with evidence that internal R&D is more valuable to companies (Arora et al., 2017).

There is also evidence supporting the second condition, that regulated and unregulated companies have systematically different estimates of the innovation surplus \mathbf{B} . Environmental regulations increase the salience of pollution management within regulated companies compared to unregulated ones, leading them to be more aware of the costs of regulation (Martin et al., 2012), and more attuned to the gains from innovation (Martin et al., 2011).

On the cost-side, free emissions allowances can generate windfall profits for regulated companies (as discussed with respect to adoption), which might reduce the marginal cost of innovation (Cramton and Kerr, 2002). Together, higher marginal benefit along with the lower marginal cost imply that a cap-and-trade program creates greater incentives for innovation for regulated companies (illustrated in figure 2).




Figure 2: Optimal innovation for regulated and unregulated companies: MS_R and MS_U illustrates marginal innovation surplus for a regulated and an unregulated company that are otherwise identical. The cap-and-trade program drives up the marginal benefit by $(1 - \beta)$ for every level of innovation. Their corresponding marginal innovation costs are illustrated by MC_R and MC_U . If the cap-and-trade program creates windfall profits for the regulated companies, this can reduce their marginal innovation costs.

Whether or not these specific mechanisms operate in practice, evidence strongly suggests that the incentive to innovate provided by environmental policies is in fact substantially smaller for unregulated companies (Popp, 2006; Dechezleprêtre and Glachant, 2014; Calel and Dechezleprêtre, 2016). Whatever the reason for this, it would lead one to expect a cap-and-trade program would drive a wedge between the incentives of regulated and unregulated companies to innovate.

So what can we expect?

The ‘induced innovation’ hypothesis is typically described very simply in the empirical literature: a rise in the relative price for a given factor of production spurs greater innovation aimed at economising on the use of that factor. In this section I have presented a stylised framework that illustrates how this hypothesis fits into the broader context of directed technological change. A cap-and-trade program can induce both adoption and innovation, and it is easy to visualise how a particular emissions cap might result in relatively more adoption or innovation depending on the particular set of current abatement options (captured by the shape of the marginal abatement cost curve) and potential technological innovations (captured by the nature of the inward shift of the curve).

This simple framework does predict one important empirical pattern, though. If a cap-and-trade program encourages adoption or innovation, it will do so more strongly for companies regulated under the cap than those that are not. The point is relatively straightforward for adoption, but it applies also to innovation. A cap-and-trade program can encourage unregulated companies to innovate as well, in which case a comparison between the two groups of companies will tend to produce conservative estimates for the program’s total effect. But nonetheless, for companies with similar levels of patenting and R&D spending prior to the new regulation—which ensures they balance the benefits and costs of innovation at the same level—the wedge introduced by a cap-and-trade program provides a way to empirically identify whether the EU ETS is having an impact.

Finally, by looking at three corresponding measures of general adoption and innovation—labour-intensity, other patent filings, and other R&D spending—I am also able to see the direction of technological change and whether adoption and innovation in low-carbon technologies have crowded-out investments elsewhere. When taken together, these six measures provide a richer understanding of whether and how the EU ETS has altered the incentives for adoption and innovation.

3 Literature review

The success or failure of environmental regulations often hinge on how they affect technological choices. The understanding that they create incentives for innovation is often critical to build political support for environmental regulations in the first place (European Commission, 2005, 2012). Once the policies are implemented, the technological choices

become important since clean innovation creates exceptionally large knowledge spillovers (Dechezlepretre et al., 2014), and since new clean technologies may substantially reduce the long-run cost of pollution abatement (Kneese and Schultze, 1975; Jaffe et al., 2003; Stavins, 2007). Successful technological innovation can also feed back and inspire more ambitious environmental policies (Schmidt and Sewerin, 2017).

Environmental regulations have consistently been found to spur innovation in clean technologies, but the historical evidence reveals a more complicated picture with respect to cap-and-trade programs.

Evidence from past cap-and-trade programs

Empirical research has found broad support for the environmental ‘induced innovation hypothesis’ (Jaffe and Palmer, 1997; Newell et al., 1999; Berman and Bui, 2001; Popp, 2002; Brunnermeier and Cohen, 2003), but when it comes to cap-and-trade programs specifically, responses have tended to emphasise adoption over innovation. The Acid Rain Program—a major cap-and-trade program for SO₂ emissions from US coal plants—encouraged companies to adopt readily available flue-gas desulfurization (FDG) technology (or ‘scrubbers’) and a strategy of purchasing coal with a lower sulfur content (sometimes known as ‘blending’) (Schmalensee et al., 1998). As we can see in panel A.1 of figure 3, the number of scrubber installations shot up right when the Acid Rain Program was introduced, and panel A.2 shows how the share of low-sulfur coal kept rising.³ Although there was a relative increase in patenting for higher-efficiency scrubbers after the 1990 Clean Air Act Amendments were passed (Popp, 2003), the efficiency of newly installed scrubbers actually went into decline after the surge of installations in the mid-1990s (B.1). Overall, patenting for low-sulfur technologies declined substantially after the cap-and-trade program was introduced (Taylor, 2012), as shown in B.2. The majority of improvements in the performance and capital costs of scrubbers occurred well before the Acid Rain Program came into effect (Taylor et al., 2003, 2005), and the Program provided no obvious boost to other low-emission technologies for coal combustion either (Hanemann, 2010; Taylor, 2012). Furthermore, it may be that some of the US patents for SO₂-control technologies were not

³The graphs collected in figure 3 are synthesised from the empirical literature to illustrate its main conclusions. The conclusions themselves, however, are based on a much more comprehensive and rigorous assessment of available data. I should also note, unsurprisingly, that there is no universally agreed upon line that distinguishes ‘adoption’ from ‘innovation.’ Sometimes ‘blending’ is referred to as innovation, and sometimes as adoption. I consider it to be adoption, in line with the definition offered in the previous section.

protecting new inventions but rather technologies imported from other countries that did not have cap-and-trade programs for SO_2 (Dekker et al., 2012).

Studies of other cap-and-trade programs tell a similar story. The NO_X Budget Trading Program spurred adoption of selective catalytic reduction (SCR) technology (A.3 and Fowlie et al., 2012), but patenting for SCR and other NO_X -control technologies fell dramatically (B.3 and Taylor, 2012). The EPA's market-based phase-down of leaded petroleum stimulated a spurt of adoption of the main substitute technology, pentane-hexane isomerization (A.4), but the cost of the equipment did not come down markedly (B.4), nor were any noteworthy new technologies introduced (Kerr and Newell, 2003). When markets for chlorofluorocarbons (CFCs) were set up in the late 1980s, an economically competitive replacement (hydrochlorofluorocarbons, or HCFCs) already existed (Gorman and Solomon, 2002). It was relatively easy for most industries to modify their production processes, and where technological advances were necessary and did occur, they can be attributed to investments that took place well before the CFC trading program was in place (Falkner, 2008). These cap-and-trade programs were implemented at different times, for different industries and pollutants, requiring different abatement technologies, yet the collected evidence shows a remarkably consistent pattern of companies favouring adoption over innovation.

The emphasis on adoption over innovation in past cap-and-trade programs has been viewed as something of an economic puzzle, especially in light of claims that market-based environmental regulations will provide especially large incentives for innovation (Downing and White, 1986; Jaffe et al., 1995). One explanation is that explicit emissions prices address the environmental externality in a targeted way, but do nothing to remedy innovation-related market failures (Jaffe et al., 2005). Others have emphasised non-rational behaviour of companies (Kreutzer, 2006; Fazekas, 2008) and high transactions costs (Stavins, 1995; Naegele, 2015) as obstacles preventing companies from participating actively and responding to the incentives these programs theoretically provide. There is circumstantial evidence for each of these explanations, and if correct, they suggest a similar pattern of 'adoption-over-innovation' would be expected in future cap-and-trade programs. The stylised framework discussed in the previous section, however, provides another interpretation—one based in the particular characteristics of those abatement technologies, and one that may not extend this pattern to future cap-and-trade programs.

A. Adoption

B. Innovation




Figure 3: Adoption and innovation in past cap-and-trade programs. Panels A.1, A.3, B.2, and B.3 are adapted from Taylor (2012). A.2 is adapted from EIA (2009). B.1 is replicates and extends Popp (2003) using information from the EIA-767 forms (as in Popp, I show a 3-year moving average). A.4 and B.4 are adapted from Kerr and Newell (2003). These graphs are only meant to illustrate the main findings from the empirical literature, but the conclusions themselves are based on a much more comprehensive and rigorous assessment of available data.

To be clear, the lack of innovation has not been a major problem in past programs, and may even have been desirable. But a repetition of this pattern would be concerning today, where cap-and-trade programs play a large and growing role in climate change

policy around the world (Zechter et al., 2016). Since current technologies are insufficient to make deep cuts in industrial greenhouse gas emissions at low cost, we need these new cap-and-trade programs to stimulate innovation (Stavins, 2007). A recent roadmap for decarbonisation suggests that R&D for clean technologies needs to increase by an order of magnitude by 2030 (Rockström et al., 2017). It is therefore an important question whether carbon markets, contrary to past cap-and-trade programs, will contribute to this low-carbon transformation.

Evidence from the EU ETS

The EU ETS was the first large-scale cap-and-trade program for carbon emissions, regulating nearly half of the EU’s greenhouse gas emissions. There has been a great deal of interest in the private sector’s technological responses to the program ever since it launched in 2005 (European Commission, 2005; Schleich and Betz, 2005; Gagelmann and Frondel, 2005; Grubb et al., 2005), but no clear answer has yet emerged.

One set of studies have suggested that a large portion of the emissions reductions can be accounted for by fuel switching in the power sector (Delarue et al., 2008, 2010). When power companies bring less polluting gas-fired plants online before coal-fired ones as demand ramps up, this changes the average fuel-mix in favour of natural gas and reduces the carbon-intensity of output. This is essentially a story of adoption. More recent studies, however, have hinted at the exact opposite result. Studies of regulated companies in Sweden (Widerberg and Wråke, 2009; Löfgren et al., 2013), Norway (Klemetsen et al., 2016), Lithuania (Jaraite-Kažukauske and Di Maria, 2016) and Ireland (Jaraite and Maria, 2011) have all struggled to find evidence that the EU ETS has encouraged adoption and any reduction in the CO₂-intensity of output.⁴

There are contradictory results about the innovation-response as well. A handful of early case-studies and expert interviews found that companies were not making resources available for development of new low-carbon technologies (Hoffmann, 2007; Tomás et al., 2010; Anderson et al., 2011), and these findings were bolstered by evidence from patent data with wider geographical and sectoral coverage (Aghion et al., 2009). Yet, a more recent theme in the literature is that low-carbon innovation appears to be on the rise,

⁴These findings do not contradict evidence that EU ETS regulated companies have reduced the absolute level of emissions of regulated companies in Germany (Petrick and Wagner, 2014), France (Wagner et al., 2014), and Norway (Klemetsen et al., 2016). In fact, the 2008 recession might contribute to just such reductions even in the absence of adoption.

where we can measure it, and EU ETS companies in particular appear to be investing in low-carbon innovation for the longer term (Martin et al., 2011; Calel and Dechezleprêtre, 2016).

How shall I summarise these findings? One reading of the literature casts the EU ETS as a driver of significant adoption, with little or no effect on innovation. This would be a repetition of a pattern familiar from many previous cap-and-trade programs. If we place greater emphasis on the more recent studies, though, we might conclude the pattern has been exactly reversed; the EU ETS as a driver of innovation, without much effect on adoption. But ultimately, it is difficult to compare estimates from different countries and sectors, obtained with different methodologies. And of course, no single study has tried to look at both adoption and innovation together, so it is an even more tenuous exercise to infer how companies have responded across these different margins.

To help adjudicate between the conflicting accounts of the EU ETS's effects, one would ideally compile a data set recording multiple comparable measures of adoption and innovation at the company-level. Unfortunately, these data are rarely collected by the agencies with the power to do so, and even when they are, access to administrative company-level data sets is restricted so that they cannot be merged across national borders in Europe. The research reviewed above has dealt with these constraints in different ways. Calel and Dechezleprêtre (2016) studied only patenting, since comprehensive and consistent records are available from the European Patent Office. Martin et al. (2011), aiming to get a richer measure of innovation activities, instead conducted their own telephone interviews with company managers. But instead their measures are more difficult to compare with other studies, and they suffer from low response rates. This paper aims for a middle ground between these two approaches: focusing on a single country, Britain, which is one of Europe's largest polluters and which collects an unusually large amount of relevant company-level data. This will allow me to study six key measures of adoption and innovation for as large a sample as is available anywhere: CO₂-intensity, labour-intensity, low-carbon patenting, other patenting, low-carbon R&D spending, and total R&D spending. These different measures can more easily be compared across studies, and yet they combine to provide a rich understanding of how companies have responded to the EU ETS.

4 Adoption, innovation, and the EU ETS in Britain

Data

Britain houses about 1,200 of the roughly 12,000 plants regulated under the EU ETS across Europe. A total of 445 British companies operate EU ETS-regulated plants. But the EU ETS was rolled out gradually in Britain—it was first introduced in 2005, but some plants were given temporary exemptions that expired in 2007 and 2008 (a feature I exploit later). The darkening shading in figure 4 illustrates the growing number of regulated companies over time. The left panel of figure 4 shows substantial economy-wide gains in CO₂-intensity in this period. From 2000 to 2010, the CO₂-intensity of British industry fell by 25%. But it is not clear if we should attribute these gains to the EU ETS. There is no obvious discontinuity or break in the trend around the time the EU ETS was implemented, and the decline in CO₂-intensity was closely matched by declines in labour-intensity, which suggests the reduction in CO₂-intensity may just reflect economy-wide efficiency improvements.

Measures of innovation exhibit a rather different pattern (although it is somewhat muted by the logarithmic scale in figure 4). The number of low-carbon patents (middle panel) was growing at 9% per year in the early part of the decade, but the growth rate more than doubled in 2005, a change that was sustained for the next few years before eventually settling back to 9% by the end of the decade. There is a clear structural break in 2005 (F -test, $p = 0.003$). Meanwhile, the number of patents filed to protect other technologies changed more slowly and steadily throughout the period, and with no apparent trend-break. As a result, the share low-carbon patents rose from 3% in 2000 to 13% in 2010.

Corporate R&D spending seems to exhibit a similar pattern as patenting (right panel of figure 4). What few estimates we have also show that corporate low-carbon R&D spending rose substantially after the EU ETS was implemented, though official statistics are entirely lacking before 2005. Corporate low-carbon R&D spending quadrupled between 2006 and 2012, corresponding to an annual growth rate of more than 25%.⁵ Total corporate R&D spending grew at a more modest rate of 3% per year, although it grows much faster in 2006 and 2007 (15% and 10% respectively). Most of this growth spurt vanishes if we take the low-carbon R&D figures at face value and net them out from the total figures, which

⁵An alternative source sometimes used to estimate corporate low-carbon R&D spending is the BERD survey compiled by Eurostat. It reports corporate R&D spending by the energy sector starting in 2007, and excludes much of corporate low-carbon R&D and includes much that would not be classified as low-carbon. Yet for comparison, they record a 70% increase from 2007 to 2012, corresponding to a more modest but still impressive annual growth rate of 11%.

suggests that the faster growth may primarily be a consequence of rising low-carbon R&D spending.




Figure 4: Trends in efficiency and innovation. **Left panel:** Plots the average number of tonnes of CO₂ emitted and people employed by British industry per £1,000 of Gross Value Added, obtained from the Office of National Statistics. **Middle panel:** Plots the annual counts of patents protecting low-carbon and other technologies, obtained from PATSTAT. **Right panel:** Plots total corporate R&D spending and that part specifically directed toward developing low-carbon technologies. Estimates of total corporate R&D spending are obtained from Eurostat, while estimates of corporate low-carbon R&D spending are collected from Wiesenthal et al. (2009), Wiesenthal et al. (2012), Corsatea et al. (2015), and Fiorini et al. (2016), four studies that use a broadly consistent methodology to produce single-year estimates. The four estimates are connected by a dashed line to make it easier to see the trend, but they do not strictly form a single time-series. The darkening grey shading indicates the growing number of companies that were regulated under the EU ETS over time.

These aggregate statistics provide an informative starting point, but it would be naive to attribute increases in low-carbon patents and R&D spending to the EU ETS without accounting for other factors that affect adoption and innovation. For instance, the price of oil and the amount of government-funded R&D investments, both known drivers of low-carbon innovation, also increased over the same period (see figure 5). To isolate the effects of the EU ETS, we have to distinguish between the outcomes of those companies that became regulated under the EU ETS and unregulated companies.

I have constructed a new data set for this purpose, which merges several restricted-access micro-data sets held by UK Data Service—the Business Structure Database (Office for National Statistics, 2012), the Quarterly Fuels Inquiry (Department of Energy & Climate Change, 2011), the UK Innovation Survey (Department for Business, Innovation & Skills et al., 2012), the Business Enterprise Research and Development survey (Office for National Statistics, 2011), and the DECC Low Emissions R&D Survey (Department of

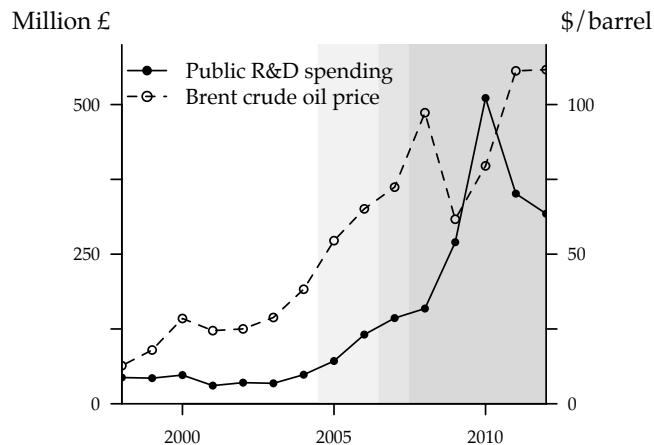


Figure 5: Potential confounders. Figures for annual public R&D spending are taken from International Energy Agency, and annual average oil prices are from the BP Statistical Review. As in figure 4, the shading indicates the increasing number of companies that became regulated under the EU ETS as the program was rolled out in the UK.

Energy & Climate Change and Office for National Statistics, 2012)—with the EPO Worldwide Patent Statistical Database, as well as regulatory databases gathered from the UK’s Environment Agency, Department of Energy & Climate Change, and the European Commission. Aside from standard company characteristics such as revenue and employment, this combined data set records all six of our key measures of adoption and innovation at the company-level.

The full data set covers over six million companies active in Britain between 2000 and 2012, although data are missing for many years and variables since many companies in the data set were not active for the whole period, and because of the widely varying scope and coverage of the underlying data sources. The DECC Survey, for instance, was a one-off survey conducted for the year 2008, so variables measured in this survey can only be measured post-treatment. The database also identifies companies regulated under Britain’s domestic climate change policies (discussed below), as well as 445 companies operating 966 British plants regulated under the EU ETS.⁶

Table 1 summarises the company characteristics and data availability during the pre-EU ETS period. The EU ETS companies were hardly representative of the business community

⁶This data set excludes 242 British EU ETS regulated plants that were operated by hospital trusts, city councils, the Ministry of Defence, and other non-firm entities. The behaviour of these entities is beyond the scope of this study.

at large even before the EU ETS was implemented. They had on average 650 times greater annual revenues than the average British company and 260 times as many employees. Companies in the electricity and gas sector (7%) as well as in manufacturing (49%) are likewise heavily overrepresented among future EU ETS companies. Data on efficiency and R&D spending are much more limited, but what we do have indicate that would-be EU ETS companies emitted nearly four times as much CO₂ per £1,000 of revenue they earned, and employed only half as much labour. They filed patents at substantially higher rates than the rest of the business community, and invested much more heavily in R&D. The fact that these companies were had substantial capacity for innovation even before the EU ETS lends further support to the expectation that they would respond to the EU ETS by innovating in-house.

Table 1: Average annual company characteristics before the EU ETS, 2000-2004

	EU ETS companies (<i>N</i> = 445)			Non-EU ETS companies (<i>N</i> = 6,022,188)		
	μ	σ	% missing	μ	σ	% missing
Company basics						
Revenues ¹ (£1,000)	770,417	2,795,245	18	1,191	94,127	39
Employees ¹	2,349	8,081	18	9	282	39
Electricity and Gas ¹	0.07	0.25	0	<0.01	0.01	0
Manufacturing ¹	0.49	0.50	0	0.06	0.23	0
Agriculture & Services ¹	0.23	0.42	0	0.63	0.48	0
Efficiency						
CO ₂ -intensity ^{1,2*}	0.42	0.62	78	0.11	0.21	99
Labour-intensity ¹	0.01	0.01	20	0.02	0.02	43
Patenting						
Low-carbon ³	0.11	1.13	0	<0.01	0.01	0
Total ³	2.61	16.05	0	<0.01	0.12	0
R&D spending						
Low-carbon ⁴ (£1,000)	-	-	100	-	-	100
Total ^{5,6**} (£1,000)	20,978	74,437	67	789	1,405	99

Data sources: ¹ Business Structure Database, ² Quarterly Fuels Inquiry, ³ EPO Worldwide Patent Statistical Database, ⁴ DECC Low Emissions R&D Survey, ⁵ UK Innovation Survey, ⁶ Business Enterprise Research and Development survey. * Computed by dividing CO₂ emissions (from the Quarterly Fuels Inquiry) by revenue (from the Business Structure Database), where both are observed. ** Computed by taking the UKIS estimate whenever available, and imputing the BERD estimate when only it is available. Alternative procedures for combining the UK Innovation Survey and Business Enterprise Research and Development estimates do not yield substantially different results.

It is hard to tell whether the EU ETS is the cause of systematic differences in outcomes when the two groups of companies were so different to start with. But if we could identify

a subset of companies for which treatment was plausibly randomised—groups of EU ETS and non-EU ETS companies that appeared similar *ex ante*—the comparison might yet provide credible estimates of the EU ETS’s effects.

Matching design

To estimate the EU ETS’s effects, I would like to compare the outcomes of regulated companies with those of historically similar companies that did not become regulated under the EU ETS. The design of the EU ETS gives me reason to believe I might be able to find such companies.

In order to control administrative costs, the EU ETS covers only sufficiently ‘large’ plants. For instance, only combustion plants with a historical annual thermal input in excess of 20 MWh are regulated, only ceramics factories with a capacity in excess of 75 tonnes per day, only paper mills with a production capacity exceeding 20 tonnes per day, and so on. Companies with smaller plants are not regulated under the EU ETS, although the companies themselves might be just as large, just as polluting, and just as innovative as those covered by the regulations. Because regulatory status is determined at the plant-level, but efficiency and innovation are company-level characteristics, I can in principle compare companies of similar size, with similar initial levels of input-intensities, similar resources available for research, and similar patenting histories, but which have recently fallen under different regulatory regimes.

But in Britain, as in many other places, the EU ETS was one of several carbon pricing policies, so how can I be sure that I am estimating the effect of the EU ETS and not of those other policies that might correlate with EU ETS treatment? Usually it is impossible to disentangle the effects, but the implementation of the EU ETS in Britain has some unique features that allow me to isolate the effect of the EU ETS.

When the EU ETS launched in 2005, Britain already had three carbon pricing policies in place: the Climate Change Levy (CCL), Climate Change Agreements (CCAs), and the UK Emissions Trading Scheme (UK ETS). The important feature of these policies, for my purposes, is that plants that were already participating in the UK ETS in 2005 could defer entry into the EU ETS until 2007, and any plants covered by CCAs could defer until 2008. Consequently, a substantial proportion of British companies did not become regulated under the EU ETS until 2007 or 2008.⁷ But there were discrepancies between

⁷See Bowen and Rydge (2011) for a more comprehensive review of British carbon pricing policies.

the regulatory thresholds for these domestic policies and the EU ETS. Participation in a CCA therefore does not perfectly predict EU ETS membership. The same applies to participation in the UK ETS and a company's CCL-bill. This means that when a given company enters into the EU ETS, there will in principle be another company that shares the same regulatory history but remains outside of the EU ETS and instead continues to be regulated under the counterfactual carbon policy. I can therefore compare ETS companies and non-ETS companies with identical histories of prior regulation, thus isolating the effects of the EU ETS. To the extent that the effect of the EU ETS interacts with general economic conditions, this gradual roll-out also lets me average the treatment effect across states of the world in a way that is not possible when all companies become regulated at once.

I start by trying to find one or more unregulated companies that were similar to an EU ETS company up until the eve of regulation. I drop 42 ETS companies for which there are no close matches, though I will later show that this omission does not alter the conclusions in substance. The matched sample includes 403 ETS companies and 446 non-ETS companies. I take the average when I find several close matches, so that ultimately I have the weighted equivalent of one non-ETS company for every ETS company.⁸

Even in a true randomised experiment, one expects treated and control groups to be significantly different in some dimensions (for 5% of covariates at the 5% significance level, and so on). Table 2 summarises the differences for the matched set, and shows a high degree of covariate balance. The matched ETS companies are somewhat larger than their matches, although the mean and median differences are small relative to the standard deviation. The final two columns report the results of tests for the equivalence of ranks and means, and I can confidently reject the hypothesis that the groups are statistically significantly different by either measure. I have stated 'difference' as the null hypothesis, and only reject it in favour of 'equivalence' if there is sufficient evidence that the difference

⁸The matched set was constructed in three steps: (1) exact and coarsened exact matching is used to quickly discard a large number of unregulated companies that are very poor candidates for matching (Iacus et al., 2009), (2) genetic matching is used on the remaining set to find a set of matches that maximises balance across all covariates (Sekhon, 2007), and (3) calipers are imposed to exclude EU ETS companies where the best available match is poor, which drops 42 EU ETS companies and their matches. The appendix reports results that include the 42 EU ETS companies that are omitted from the main analysis. This unconventional matching algorithm was used, rather than simple propensity score matching say, because it managed to achieve a higher degree of balance for a larger set of regulated companies. But the nature of the matching algorithm is ultimately irrelevant (except for replication purposes), since the results are conditional on the matched set. The validity of the results only depend on the similarity between the matched companies.

is within the equivalence range (conventionally 0.2 standard deviations). A small p -value is therefore evidence of equivalence. This is a more stringent test of covariate balance than merely failing to reject the null of equivalence (Hartman and Hidalgo, 2014).

The sectoral distribution (at 3-digit level) is exactly the same across groups. The baseline levels of efficiency, patenting, and R&D spending also appear similar across the two groups.⁹ ETS companies are slightly more likely to have participated in a CCA and the UK ETS; they also tended to have ever-so-slightly larger CCL-bills and receive marginally more government support for R&D, but these differences are neither substantively nor statistically significant. In addition to matching on observed values, companies are matched on the pattern of missing values to account for the possibility that data is not missing at random.

The quantile-quantile plots in figure 6 further shows that it is not just the central tendencies that are similar, but that the whole distributions of the efficiency, patenting, and R&D variables are similar across the two groups of companies. Figure 7, which tracks the outcomes of the matched companies between 2000 and 2012, also reveals that these variables do not exhibit different pre-treatment trends.

Even though the matched companies are observationally equivalent, there is never a guarantee of similarity with respect to unobserved characteristics. Nevertheless, the fact that ETS assignment is determined by the capacities of individual plants does provide some assurance. There is no body of evidence showing that company-level adoption and innovation activities are correlated with the capacities of individual plants, conditional on company-level characteristics. Moreover, as the robustness tests will show later, the matching appears to have achieved balance even with respect to some variables that were not explicitly matched for. This provides another reason to believe that the matched set approximates the statistical conditions of a randomised experiment. These regulated and unregulated companies are statistically indistinguishable from each other on the eve of regulation, and it is as if regulatory status was assigned by a coin-flip within each matched pair. Next, I compare their outcomes.

⁹I have used the *share* of low-carbon patents to proxy low-carbon R&D spending, since the latter is only observed in the treatment period. This is in line with other studies that have tried to infer company-level corporate low-carbon R&D spending (Wiesenthal et al., 2009, 2012; Corsatea et al., 2015; Fiorini et al., 2016).

Table 2: Equivalence tests for matched EU ETS ($N = 403$) and non-EU ETS companies ($N = 446$)

	Median difference	Mean difference	Equivalence range	Signed-rank test p -value	Paired t -test p -value
Company basics					
Revenues (£1,000s)	10,367	152,942	$\pm 461,220$	$<.001$	$<.001$
Employees	49	473	± 796	$<.001$.011
Economic sector (3-digit)	<i>Exact match</i>	–	–	–	–
Efficiency					
CO ₂ -intensity	-0.001	0.031	± 0.093	.015	.025
Labour intensity	-0.001	-0.001	± 0.003	.004	.005
Patenting					
Low-carbon patents	0	0.003	± 0.029	.642 ^{LP}	.002
Other patents	0	-0.003	± 0.635	.049	$<.001$
R&D spending					
Low-carbon patent share	0	0.003	± 0.007	.783 ^{LP}	.028
R&D spending (£1,000s)	65	-1,938	$\pm 8,975$	$<.001$.005
Regulatory history					
CCA participation	0	0.112	–	–	–
UK ETS participation	0	0.025	–	–	–
CCL bill (£1,000)	12.719	11.034	± 34.671	.044	.024
R&D support (£1,000)	0.394	11.901	$\pm 1,655.956$	$<.001$	$<.001$


The first two columns report the median and mean differences between matched EU ETS and non-EU ETS companies. Economic sector is exactly matched at the 3-digit level. CCA and UK ETS participation are both coded as 1 for participants and 0 otherwise, so the median difference tells me whether the median is a concordant or discordant pair, and the mean difference is the difference in proportions. The empirical distributions of EU ETS and non-EU ETS companies are judged to be substantively equivalent if the difference lies within the ‘equivalence range’ reported in the third column, specified as ± 0.2 standard deviations of the pooled sample (Cochran and Rubin, 1973; Ho and Imai, 2006). The fourth and fifth report the p -values for Wilcoxon’s signed-rank test and the paired t -test of the null hypothesis that the difference between distributions lies outside the equivalence range. *LP* indicates low power. The signed-rank statistic depends only on discordant pairs, so the effective sample size is small for covariates with a large proportion of zeros, in this case patent filings and R&D spending. The test can then fail to reject, at the 5% level, both that the difference is inside and outside the equivalence range. The paired t -test treats concordant pairs as evidence against the hypothesis of difference, and will therefore tend to reject it in favour of equivalence when a covariate has a large proportion of zeros. I report the results from both tests, although favour the signed-rank test for patenting and R&D spending, as concordant zero-pairs do not seem to provide much evidence against the hypothesis of difference. Fisher’s exact test was performed for CCA and UK ETS participation, and in neither case could an odds-ratio of one be rejected (numbers not reported for disclosure reasons).

5 Results

Main results

Matched ETS and non-ETS companies followed similar trends in the pre-treatment period (see figure 7). The most dramatic feature of figure 7 is the divergence of patenting activities after the EU ETS came into effect. Figure 7 also reveals that EU ETS companies have

Figure 6: Comparison of matched EU ETS and non-EU ETS companies



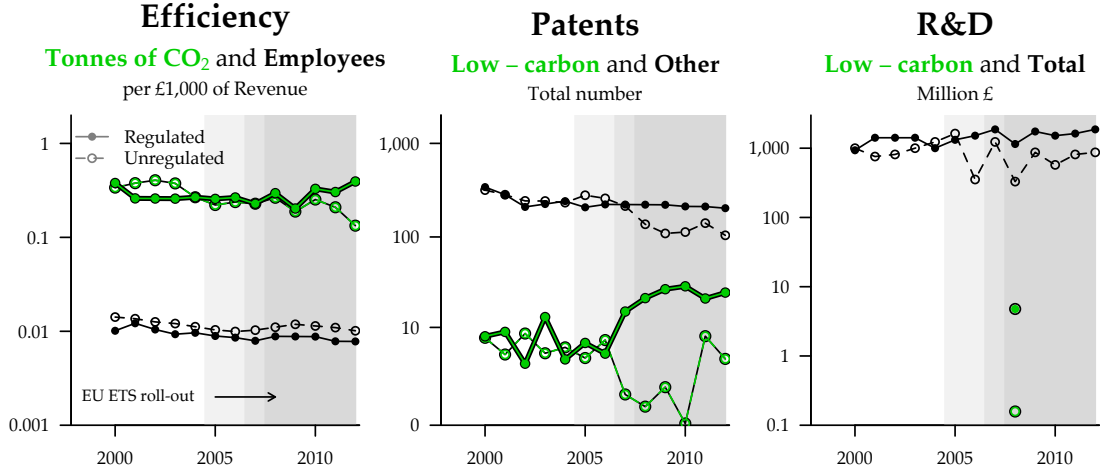
The panel titled **Efficiency** displays the empirical quantile-quantile (e-QQ) plot for average annual CO₂-intensity (in green) and labour intensity (in black) before entry into the EU ETS. The points would all be on the 45° line if the distributions of the EU ETS and non-EU ETS companies were exactly the same. The panel titled **Patenting** shows the e-QQ plot for the average annual number of low-carbon patents filed (in green) and other patents (in black) over the same period. The panel titled **R&D spending** shows the e-QQ plot for the average annual R&D spending, which cannot be separated into low-carbon and other. Scales are logarithmic. To comply with UK Data Service disclosure rules, the each point represents a percentile-pair rather than a company-pair. The company-pair version of this graph is available to view through the UK Data Service.

invested more in R&D in the post-treatment period. It is less clear that the CO₂-intensity or labour-intensity has responded to the EU ETS. Adoption of off-the-shelf technologies and managerial solutions have in past cap-and-trade programs offered the fastest route to comply. If the EU ETS was driving widespread adoption of such abatement solutions, one would expect a reduction in the CO₂-intensity of EU ETS companies early on in the post-treatment period, but in fact CO₂-intensity follows parallel trends for the most part. A gap opens up only at the end of the study period, and in the opposite direction from what one would expect. At first glance, then, it appears that the EU ETS may have induced greater innovation, but has not encouraged adoption.

Efficiency

I can leverage the matched design of the sample to provide more rigorous statements about treatment effects. If the matching was successful and the EU ETS had no effect, one would expect there to be no systematic difference between ETS and non-ETS companies in the post-treatment period. After all, they were statistically indistinguishable beforehand, and the only thing that has happened is that I have arbitrarily labeled half of them as ‘ETS companies.’ Formally, let me write the treated-minus-control difference as $\delta_n = T_n - C_n$ for

Figure 7: Adoption and innovation of matched EU ETS and non-EU ETS companies



Left panel plots average annual tonnes of CO₂ and number of employees that matched companies used per £1,000 of revenues they generates. The grey shading indicates how many of the EU ETS companies were regulated under this program at different points in time. **Middle panel** plots the total number of patents filed annually to protect low-carbon and other technologies, respectively. **Right panel** plots total annual R&D spending, and for the year 2008, the total R&D spending directed toward developing low-carbon technologies. The colours in the titles provide an implicit legend: the green lines refer to low-carbon measures while the black lines refer to their complements. To deal with the fact that non-EU ETS companies filed zero low-carbon patents in 2010, the scale of middle panel is adjusted to reflect that I have added 1 to all counts before taking the logarithm.

pair $n = 1, \dots, N$. One would expect the δ s to be positive as often as they were negative, the sign depending only on our labelling. If the EU ETS caused treated companies to reduce their CO₂-intensity by some proportion τ , however, there would be an excess of negative δ s compared to what one would expect by chance. I can then recover the EU ETS's effects on CO₂-intensity by finding the value $\hat{\tau}$ in the expression $\delta_n = T_n(1 + \hat{\tau}) - C_n$ that restores this known null-distribution of δ s. I follow the standard practice of using Wilcoxon's signed-rank statistic to measure distributional similarity, and the Hodges-Lehmann point estimate $\hat{\tau}$ is the value that maximises the p -value, or equivalently, minimises the difference (Rosenbaum, 2013). This estimation procedure is very parsimonious, requiring only the assumption that matched companies had an equal *ex ante* chance of becoming regulated.

Earlier studies have struggled to find evidence that the EU ETS has encouraged adoption of abatement technologies and strategies (Löfgren et al., 2013; Widerberg and Wråke, 2009; Klemetsen et al., 2016; Jaraite-Kažukauske and Di Maria, 2016; Jaraite and Maria, 2011). Figure 7 suggests that ETS companies experienced an *increase* in CO₂-intensity over the period, while non-ETS companies experienced a reduction. Indeed, the point es-

timate reported in table 3 indicates that the EU ETS caused a typical regulated company to increase the CO₂-intensity of its output by 16.1% compared to historically similar unregulated companies. For a fixed revenue, this would imply that the EU ETS may have caused an additional 6,000 tons of CO₂ to be emitted per company per year. These estimates are nowhere near statistically significant, however. Together with the suspiciously late timing of the divergence of CO₂-intensities seen in figure 7, this should not be read as strong evidence that the EU ETS has caused an increase. Rather, it adds to the number of studies looking for reductions in CO₂-intensity in the EU ETS and turn up no evidence of it. While I cannot reject the hypothesis of a reduction in CO₂-intensities, the general absence of evidence is now starting to look more like evidence that the EU ETS has not yet spurred widespread adoption.

Table 3: Matching estimates of the effects of the EU ETS

	Hodges-Lehmann point estimate	<i>p</i> -value
Adoption (proportion)		
CO ₂ -intensity	0.161	0.588
Labour intensity	-0.166	0.010
Patenting (number of patents)		
Low-carbon patents	0.415	0.079
Other patents	0.130	0.167
R&D spending (£1,000s)		
Low-carbon R&D	200.000	0.046
Total R&D	514.000	<0.001

The first column reports the location shift parameter that maximises the *p*-value of Wilcoxon’s signed rank statistic, or put another way, that minimises the difference in mean ranks. This will correspond to the median treatment effect for symmetric distributions. The accompanying *p*-values are computed under the null hypothesis that the EU ETS has had no effect.

The apparent lack of adoption could, in principle, be explained if ETS companies were slower at adopting efficiency-enhancing technologies in general, for some unobserved reason. This alternative explanation seems implausible, however, given that ETS companies reduced their labour-intensity faster than their unregulated counterparts.¹⁰ This instead casts the lack of reductions in CO₂-intensity in a more negative light.

¹⁰As a matter of economic theory, there is no difficulty explaining why a regulation that increases the relative price of one factor of production would encourage increased efficiency in another. For instance, efficiency improvements in the latter factor may be technologically complementary, or cheaper to implement per unit of total cost reduction. Whether a policy like the EU ETS has a disproportionate impact on CO₂-intensity is therefore an empirical question.

These substance of these findings is not altered by including the 42 ETS companies that were omitted from the main results (see appendix). These omitted companies are of course quite different from the other companies in our sample—generally larger, more polluting, more innovative—and were omitted for this reason by design. Yet, our inability to find good matches for these companies does not mean they do not matter. Their average CO₂-intensity was roughly two-and-a-half times greater than the 403 matched companies before EU ETS started, and subsequently fell by nearly 40%. Including the 42 omitted EU ETS companies and the best matches available does not change the signs of my estimates (see appendix). Indeed, even if I were to attribute the whole of their 40% decline to the EU ETS, this would largely cancel out any positive impact, but would not enough to reverse the sign.

Patenting

The EU ETS may not have encouraged widespread adoption of abatement technologies, but could still have encouraged investments to invent and develop new low-carbon technologies that haven't been adopted yet. ETS companies collectively filed 38 low-carbon patents in the pre-treatment period and 147 after they became regulated. Unregulated companies started out filing 31 low-carbon patents, which later fell to 26. As reported in table 3, I estimate that the EU ETS has encouraged regulated companies to file an additional 0.415 low-carbon patents per year.¹¹ The effective sample size is relatively small since most companies never file low-carbon patents, but the point estimate is remarkably close to the 2 additional patents over a 5-year period reported by Calel and Dechezleprêtre (2016). The 90% confidence interval narrowly excludes zero, (0.09, 11.89). These numbers imply that ETS companies collectively would have filed 29 (10, 138) fewer low-carbon patents if they had not been regulated, which translates into the EU ETS having caused low-carbon patenting to increase by 25% compared with the counterfactual. Even if I only attributed

¹¹The multiplicative model of the treatment effect used for efficiency earlier is problematic here since most companies do not file any patents or spend anything on R&D, which means the proportional difference is generally undefined. Instead I use an additive model that is censored at zero, which non-parametrically adjusts for the fact that most companies have a zero propensity to file patents (Rosenbaum, 2009; Keele and Miratrix, 2015; Calel and Dechezleprêtre, 2016). In effect, it is a non-parametric Tobit estimator, relying only on the assumption that the chance of treatment is the same for matched companies. The process is exactly the same as before, except that τ is estimated using the following formula instead:

$$\delta_n = \begin{cases} \max(T_n - \hat{\tau}, 0) - C_n & \text{if } \hat{\tau} \geq 0 \\ T_n - \max(C_n - \hat{\tau}, 0) & \text{otherwise.} \end{cases}$$

10 low-carbon patents to the EU ETS (the lower end of the 90% confidence interval), it would imply a 7% increase in their low-carbon patenting.

Meanwhile, patenting for other technologies fell in both groups, albeit slightly faster for the unregulated companies. ETS companies collectively filed 1,612 non-low-carbon patents over the whole regulatory period, and a treatment effect of 0.13 (-0.10, 0.43) implies that they would instead have filed 1,550 (1,455, 1,663) such patents had they escaped regulation. This translates into a 4% (-3%, 11%) increase in patenting for these technologies among ETS companies. If I focus on the lower end of the confidence intervals, I have to admit the possibility that the EU ETS may have increased low-carbon patenting only modestly among regulated companies, but at the expense of a greater reduction in patenting for other technologies. The best estimates, though, indicate that the EU ETS has encouraged patenting among regulated companies overall, and most especially for low-carbon technologies. This is evidence of directed technological change, and also mitigates concerns that low-carbon innovation has crowded out other innovation efforts (Popp and Newell, 2012). As discussed earlier, to the extent that the EU ETS has encouraged patenting among the unregulated companies as well, the total effect of the EU ETS on innovation will be even greater than my estimates suggest.

At the cost of producing a somewhat less credible counterfactual comparison, I can incorporate the 42 omitted ETS companies into our analysis (see appendix). These companies were much more innovative to start with—one of the reasons that good matches could not be found for them—and if the EU ETS caused these companies to reduce their patenting it could in principle cancel out the positive effect reported above. As it happens, though, the number of low-carbon patents filed by these companies rose by 60% after the EU ETS came into effect. If I re-estimate the treatment effect using the best available controls for these 42 companies, the effect size actually increases to 0.68 (0.30, 12.83) additional low-carbon patents per company per year (see appendix). Including the best available matches for the previously omitted companies also increases the estimated treatment effect to 0.74 (0.20, 1.49) extra *other* patents per company per year (see appendix). In this case, including the 42 omitted ETS companies actually tends to strengthen the earlier conclusions.

R&D spending

To summarise so far, I find no evidence that the EU ETS has reduced CO₂-intensity, but do find that it has boosted low-carbon patenting. An optimistic observer may take these

estimates to indicate that, while ETS companies can quite easily cope with regulations in the short term, they are nevertheless developing and bringing to market new low-carbon technologies that will become increasingly valuable in the future when climate change regulations are expected to grow more stringent. A cynical observer, on the other hand, might interpret the same estimates as evidence that the EU ETS is only encouraging companies to patent a larger share of low-carbon technologies in their research pipelines to capitalise on the regulator’s passing commitment to reducing CO₂ emissions, while in actuality doing nothing more than usual to reduce their CO₂-intensity. Neither my estimates above nor the estimates in the literature can adjudicate between these interpretations. An analysis of investment in low-carbon R&D spending, though, would provide a clue. R&D expenditures could tell me whether the additional patenting is merely superficial reclassification or reflects genuine investments in innovation.

Unfortunately I do not have a great deal of information to go on; data on low-carbon R&D spending was collected only for the year 2008, and records the low-carbon R&D spending of only 27 of the ETS and 21 of the non-ETS companies in my matched sample.¹² Of the responding companies, the proportion reporting that they were undertaking any low-carbon R&D was substantially higher among EU ETS companies as compared to non-EU ETS companies ($OR > 1$, $p = 0.064$).¹³ For a typical innovator in 2008, I estimate that the EU ETS increased low-carbon R&D spending by £200,000 (£65,000, £656,000).¹⁴ This translates into a large proportional increase, similar in magnitude to what I found for low-carbon patenting.¹⁵ Even when compared to total low-carbon R&D spending reported in

¹²The DECC low-carbon R&D survey was sent to a sample of roughly 4,500 of the largest innovators in the UK, and had a response rate of just over 50%. The sampling frame does not overlap well with the set of matched ETS and non-ETS companies. Only 83 ETS companies and 42 non-ETS companies in my matched sample were surveyed, and only 27 of the ETS and 21 of the non-ETS companies responded. The difference in response rates is just about significant at conventional levels; $OR = 0.485$, $p = 0.079$.

¹³Disclosure rules preclude public reporting of the exact odds ratio, but the numbers are available through the UK Data Service.

¹⁴As with patenting, I estimate a censored additive treatment effect model to account for the fact that some companies have zero propensity to conduct R&D.

¹⁵UK Data Service disclosure rules unfortunately prevent public reporting of the observed average low-carbon R&D spending and consequently the proportional increase. However, I can provide bounds based on publicly disclosed figures. My estimate implies that ETS companies would have spent on average £119,704 (£69,037, £141,000) had they not been regulated. If only one of the 27 matched EU ETS companies with data was responsible for all of this low-carbon R&D, it would have spent £3,232,008 in 2008 had it not been regulated. Against this, £200,000 corresponds to a 6% increase. At the other extreme, if each of the 27 would have spent £119,704 in the absence of the EU ETS, but actually spent an additional £200,000 (totalling an extra £5.4 million), this would imply a 167% increase.

the DECC survey—£66 million—the effect is non-negligible.¹⁶ If 2008 was a typical post-treatment year, this would quickly add up to many additional millions spent on low-carbon R&D.

For additional evidence I can look at total R&D spending, which I can observe over the whole study period and for a much greater number of companies. Unless low-carbon R&D was crowding out other R&D activities, Total.assets.th.EUR would expect the effect on total R&D to be at least as large as what I estimate for low-carbon R&D. Conversely, if the effect on total R&D spending was very much larger, it would suggest either implausible levels of crowding-in or substantial mis-classification of low-carbon innovation, either of which would cast doubt on the research design. As it happens, I estimate that the EU ETS was responsible for an additional £514,000 (251,000, £1,002,000) per company and year, which is neither implausibly small nor implausibly large if the estimate for low-carbon R&D is correct. It is no less than the effect estimated for low-carbon R&D, while translating into a proportional increase of 3.2% (1.7%, 5.3%) of total R&D spending. These estimates suggest that the EU ETS has had a disproportionate positive impact on low-carbon R&D—at least twice as large as the overall increase, in proportional terms—without cannibalising resources intended for other R&D activities. If anything, they suggest a mild crowding-in consistent with my estimates of increased patenting activities. The signs and magnitudes of these estimates also match up with my findings on patenting, which suggests that I am picking up genuine investments in innovation. Moreover, the substance of these conclusions are all robust to the inclusion of the 42 omitted ETS companies (see appendix).

Challenges

My estimates suggest that the EU ETS has encouraged innovation, but not to adopt already available abatement technologies as a means to achieve immediate efficiency improvements. It is unfortunate that available data often are inadequate to generate small p -values, but the consistency of signs and magnitudes across different outcomes, different specifications, and across different studies does provide other reasons for confidence in this general pattern. Nevertheless, my findings do appear to contradict a historical pattern in cap-and-trade programs, and such contrary results should be viewed with added skepticism. While I do not have the data needed to pursue an exhaustive set of robustness tests, I can try

¹⁶Following the same logic as the previous footnote, the proportional effect can be bounded between 0.3% and 8.9%.

to answer a few potential challenges to the interpretation I have offered: (1) government R&D support (an omitted variable), (2) switching to electricity (a form of measurement error), (3) exit of ETS companies (creating selection bias), and (4) innovation by non-EU ETS companies (omitted outcome).

Omitted variable: Government R&D support

Perhaps the most conspicuous omitted variable is post-treatment government R&D support. Earlier, figure 5 showed that there was an increase in public R&D spending that coincided with the roll-out of the EU ETS. Although the matching design ensures that companies received similar pre-treatment levels of support, it is possible that ETS companies received more support in subsequent years. In that case, ETS companies might have spent more on R&D without any encouragement from the EU ETS itself.

In fact, though, ETS companies did not receive more government R&D support in the years after they became regulated. ETS and non-ETS companies received virtually identical levels of R&D support prior to the EU ETS (roughly £850,000 per company per year). After the EU ETS came into effect, support for non-ETS companies fell slightly to an average of £753,000 per company per year, but much more dramatically for ETS companies, down to an average of £102,000 per company per year. Moreover, there is no significant difference in the odds that matched ETS and non-ETS companies received government R&D support specifically targeted at low-carbon R&D projects. In light of this, it is difficult to sustain the argument that differences in R&D spending previously attributed to the EU ETS were in fact due to ETS companies disproportionately benefitting from greater government support.

Leakage: Switching to electricity

The EU ETS regulates direct emissions only. A company is required to surrender emissions allowances from a coal-burning generator (direct emissions), say, but not for electricity purchases (indirect emissions). On average, only about one third of the carbon price was passed on in the form of higher electricity prices in Britain (Fezzi and Bunn, 2010), so it would have been relatively cheap to comply with the EU ETS by outsourcing energy production (except for electric utilities, of course, which are a small minority of regulated companies). Depending on whether one views outsourcing of energy production as a legitimate form of ‘adoption’ or not, the CO₂-intensities may systematically misrepresent adoption behaviour.

In fact, the proportion of total emissions from electricity consumption did not change markedly among ETS companies (see left panel of figure 8). Regulated and unregulated companies track each other closely after the EU ETS came into effect. What is more, they align closely even before 2005, despite not having been matched on this variable. These findings provide circumstantial evidence of the validity of the previous estimate on CO₂-intensity, and of the success of the matching more generally.




Figure 8: Exploring mechanisms. **Left panel:** Plots the average share of total emissions (direct plus indirect) that from electricity consumption. Indirect emissions are calculated by multiplying electricity purchases with the average emissions intensity of grid-delivered electricity. **Right panel:** Plots the number of ETS and non-ETS companies that cease operations in a given year.

Selection: Exit of EU ETS companies

The pattern of results I have found could potentially be explained by changes in the composition of the matched sample after the EU ETS was implemented. If the EU ETS drove the least innovative regulated companies to exit in greater numbers, this would inflate average patenting and R&D spending among the regulated companies without actually increasing their innovation activities.

There are three main reasons to be skeptical of this explanation. First, because the least innovative companies were also less efficient than the average, one would expect this increased exit to give rise to the appearance of increased efficiency, which I do not find. Second, figure 7 showed that the differences in patenting and R&D are partly a consequence of declining investments among unregulated companies. For the exit-hypothesis to explain this pattern, I would need to make the unlikely supposition that the *most* innovative

unregulated companies exited the sample also. Finally, as shown in the right panel of figure 8, ETS companies were actually less likely to exit than non-ETS companies ($OR = 0.43$; $p < 0.001$), and there were no obvious pre-treatment differences among the exiting companies. This type of non-random selection is therefore unlikely to explain my main findings.

Omitted outcome: Innovation by non-EU ETS companies

Another challenge to interpreting my findings is that they capture only the EU ETS's differential impacts on comparable regulated and unregulated companies. I have focused on this difference as a way to identify the program's effects, but this means I cannot see what additional encouragement the EU ETS may have provided for unregulated companies to innovate. Two studies of the EU ETS have attempted to estimate this spillover effect using patent data. Cael and Dechezleprêtre (2016) impose minimal econometric structure, and struggle to identify any such spillovers. (Miller, 2014) makes stronger identifying assumptions, and finds that, in aggregate, the spillover might be as big as the direct effect itself, which would suggest that a doubling of my estimated totals might be appropriate at most. Considering how much general innovation happens outside of these EU ETS companies, this spillover actually seems rather small. Studies of other environmental regulation also confirm that the spillovers are not generally large (Popp, 2006). In the present study, it is possible that the matched unregulated companies responded to the EU ETS, but the resulting bias is likely to be small since I generally do not observe innovation increasing among the matched unregulated companies.

More broadly, the concern may be that the 'top innovators,' often not regulated polluters themselves nor likely to be matched to them, stand ready to respond to policy changes that create demand for certain new technologies. If this spillover effect dominates, my estimates may tell only a small part of the story. To look at this, figure 9 plots the total low-carbon patenting for EU ETS companies, as well as for the 100 most prolific non-ETS patent filers. The increase in low-carbon patenting after 2005 is roughly similar across both groups (though the timing differs slightly). If all of the increase among the 'top innovators' were attributable to the EU ETS, the aggregate direct and spillover effects would be of comparable magnitude, in line with Miller's (2014) estimates.




Figure 9: Top innovators: Total granted low-carbon patents, by filing year, for EU ETS regulated companies as well as the 100 non-ETS companies that filed the most patents in the years 2000 through 2004.

6 Discussion

One of the principal objectives of cap-and-trade programs like the EU ETS is to achieve abatement by encouraging companies to innovate and become more efficient (Stavins, 2007; European Commission, 2005, 2012). Historically, though, such programs have primarily encouraged adoption of existing abatement technologies rather than investments in innovation. This paper extends recent investigations of the technological response to the EU ETS, concluding that the EU ETS may have reversed the pattern seen in past cap-and-trade programs: regulated companies *have not* widely adopted technologies that reduce their CO₂-intensity, yet they *have* increased their patenting and R&D spending. There is evidence of directed technological change, as low-carbon patenting and low-carbon R&D spending have increased faster than for other technologies (roughly 20-25% versus 3-5%). What is more, these estimates tell only of the direct effect of the EU ETS on regulated companies, and there is suggestive evidence that the spillover to unregulated innovators might be almost as large. The EU ETS is not the only conceivable explanation for these patterns, of course, but it does appear to stand up to scrutiny reasonably well.

I offer two brief reflections on these findings. The first concerns the reason for the possible reversal of the historical response to cap-and-trade programs. New findings that appear to contradict historical patterns should always be viewed with skepticism. The present study is no exception. Nevertheless, such a reversal is not as surprising as one

might first think. Previous cap-and-trade programs have generally targeted pollutants and polluters that had technologies available that could achieve the regulator's long-term abatement targets at reasonable cost. The US phase-out of leaded petroleum, for instance, was specifically aimed at encouraging adoption of existing refining technologies (Kerr and Newell, 2003). By the time tradeable quotas for CFCs were set up, an economically competitive replacement already existed (Gorman and Solomon, 2002; Falkner, 2008). When the Acid Rain Program launched, the technology to scrub sulfur from exhaust was already widely available, and on its own could have achieved greater emissions cuts than the Acid Rain Program required (Schmalensee et al., 1998). When relatively mature abatement technologies already exist, new technologies have less chance of creating major savings in compliance costs, and innovators are more likely to work to existing technical standards and focus on making marginal improvements. Under such conditions, the only real question facing companies is *when* to install the existing abatement technologies. It is not surprising that one then observes increased adoption coupled with a lack of innovation.

But the same technological pre-conditions did not exist in the EU ETS, nor do they in carbon markets more generally. The abatement potential of low-cost strategies such as fuel switching is limited (Delarue et al., 2008), which means that achieving longer-term decarbonisation goals requires new abatement technologies. Indeed, a recent roadmap indicates that substantial technological development is required in the coming decades (Rockström et al., 2017). But then why adopt off-the-shelf abatement technologies that may soon become obsolete anyway, especially as the emissions caps are not strictly binding in the short term? Why not instead try to invest and develop new low-carbon technologies that will reduce abatement costs more dramatically by the time future emissions caps become more stringent? These dynamics shift resources from adoption to innovation (Aghion and Howitt, 1992; Mauritzen, 2014), and squares well with survey evidence showing that, conditional on current stringency, EU ETS companies that expect higher future carbon prices are also more likely to innovate (Martin et al., 2009). A proper reading of the induced innovation hypothesis emphasises this interaction of the emissions price and the technology set. One does not need to invoke secondary market failures, high transactions costs, lack of polluter sophistication, or other complicating factors, to explain the lack of innovation in past cap-and-trade programs.

If this explanation is correct, it changes how one interprets assessments of carbon markets relative to other cap-and-trade programs. The increases in patenting and R&D signal that new low-carbon technologies are being developed. As these technologies become

more widely available and used, one would expect abatement costs to decline. At the same time, though, carbon markets may provide weaker incentives for adoption because of the nature of these technologies, and this may result in excess build-up of greenhouse gas concentrations in the near-term. Future ambition may well encourage longer-term directed technological change, but greater near-term stringency may also be needed to avoid the cumulative costs of too little adoption.

The second reflection concerns the constraints that this and other studies face when trying to answer questions about the effects of cap-and-trade programs. As I have noted, the evidence presented here and elsewhere does show a consistent pattern, but is ultimately quite weak by conventional statistical measures. This reflects an unfortunate reality that, even after a decade of the world's largest cap-and-trade experiment, the necessary data to learn about the program's effects simply are not available to be analysed. Despite painstaking efforts to obtain the best and most comprehensive data available, the effective sample sizes are often small and several of my estimates hover around conventional thresholds of statistical significance. This is an important finding in its own right. Imprecise estimates too often end up in the file-drawer even when they reflect our best understanding, but these findings provide critical evidence of the need for improved data collection (DeLong and Lang, 1992; Ioannidis and Doucouliagos, 2013). Although one may be limited in what can be inferred from imprecise estimates for now, this reminds us—just as China readies itself to launch an even larger carbon market—that the first step in improving regulatory design is to improve data collection practices. More deliberate collection of data on key outcomes at the start of new programs, as well as greater data sharing and integration across regulators and national statistical offices, provide low-cost opportunities that would dramatically enhance our understanding of the effects of future carbon markets.

References

- Acemoglu, D., Aghion, P., Bursztyn, L., and Hémous, D. (2012). The Environment and Directed Technical Change. *The American Economic Review*, 102(1):131–166.
- Aghion, P., Dechezleprêtre, A., Hémous, D., Martin, R., and Reenen, J. V. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1):1–51.
- Aghion, P. and Howitt, P. (1992). A Model of Growth Through Creative Destruction. *Econometrica*, 60(2):323–351.
- Aghion, P., Veugelers, R., and Serre, C. (2009). Cold start for the green innovation machine. *Bruegel Policy Contribution*, (12):1–12.

- Ambec, S., Cohen, M. A., Elgie, S., and Lanoie, P. (2013). The Porter Hypothesis at 20: Can Environmental Regulation Enhance Innovation and Competitiveness? *Review of Environmental Economics and Policy*, 7(1):2–22.
- Anderson, B., Convery, F., and Maria, C. D. (2011). Technological change and the EU ETS: the case of Ireland. *IEFE Working Paper Series*, (43).
- Arora, A., Belenzon, S., and Sheer, L. (2017). Back to Basics: Why do Firms Invest in Research? Working Paper 23187, National Bureau of Economic Research. DOI: 10.3386/w23187.
- Arvanitis, S., Donzé, L., and Sydow, N. (2010). Impact of Swiss technology policy on firm innovation performance: an evaluation based on a matching approach. *Science and Public Policy*, 37(1):63–78.
- Berman, E. and Bui, L. T. M. (2001). Environmental regulation and labor demand: evidence from the South Coast Air Basin. *Journal of Public Economics*, 79(2):265–295.
- Bowen, A. and Rydge, J. (2011). Climate-change policy in the United Kingdom. *OECD Economic Survey of the United Kingdom*, 886.
- Brunnermeier, S. B. and Cohen, M. A. (2003). Determinants of environmental innovation in US manufacturing industries. *Journal of Environmental Economics and Management*, 45(2):278–293.
- Calel, R. and Dechezleprêtre, A. (2016). Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market. *Review of Economics and Statistics*, 98(1):173–191.
- Chavas, J.-P., Aliber, M., and Cox, T. L. (1997). An Analysis of the Source and Nature of Technical Change: The Case of U.S. Agriculture. *Review of Economics and Statistics*, 79(3):482–492.
- Cochran, W. and Rubin, D. (1973). Controlling bias in observational studies: A review. *Sankhyā: The Indian Journal of Statistics, Series A*, 35(4):417–446. Dedicated to the Memory of P. C. Mahalanobis.
- Corsatea, T., Fiorini, A., Georgakaki, A., and Lepsa, B.-N. (2015). Capacity Mapping: R&D investment in SET-Plan technologies. Technical Report JRC95364, Publications Office of the European Union.
- Cramton, P. and Kerr, S. (2002). Tradeable carbon permit auctions: How and why to auction not grandfather. *Energy Policy*, 30(4):333–345.
- Dechezleprêtre, A. and Glachant, M. (2014). Does Foreign Environmental Policy Influence Domestic Innovation? Evidence from the Wind Industry. *Environmental and Resource Economics*, 58(3):391–413.
- Dechezleprêtre, A., Martin, R., and Mohnen, M. (2014). Knowledge spillovers from clean and dirty technologies.
- Dekker, T., Vollebergh, H. R., de Vries, F. P., and Withagen, C. A. (2012). Inciting protocols. *Journal of Environmental Economics and Management*, 64(1):45–67.
- Delarue, E., Ellerman, A., and D’haeseleer, W. (2010). Short-term CO₂ abatement in the European power sector: 2005-2006. *Climate Change Economics*, 1(2):113–133.
- Delarue, E., Voorspools, K., and D’haeseleer, W. (2008). Fuel switching in the electricity sector under the EU ETS: review and prospective. *Journal of Energy Engineering*, 134(2):40–46.
- DeLong, J. B. and Lang, K. (1992). Are all economic hypotheses false? *Journal of Political Economy*, 100(6):1257–1272.

- Department for Business, Innovation & Skills, Office for National Statistics, and Northern Ireland Department of Enterprise, Trade & Investment (2012). *UK Innovation Survey, 1994-2010: Secure Data Service Access [computer file]*. Colchester, Essex: UK Data Archive [distributor], 2nd edition. SN: 6699.
- Department of Energy & Climate Change (2011). *Quarterly Fuels Inquiry, 1993-2008: Secure Data Service Access [computer file]*. Colchester, Essex: UK Data Archive [distributor]. SN: 6898, <http://dx.doi.org/10.5255/UKDA-SN-6898-1>.
- Department of Energy & Climate Change and Office for National Statistics (2012). *DECC Low Emissions RD Survey, 2008: Secure Data Service Access [computer file]*. Colchester, Essex: UK Data Archive [distributor]. SN: 6894, <http://dx.doi.org/10.5255/UKDA-SN-6894-1>.
- Downing, P. B. and White, L. J. (1986). Innovation in pollution control. *Journal of Environmental Economics and Management*, 13(1).
- EIA (2009). Coal Market Module Data Tables: US Coal Production by Coal Market Module Supply Region and Mine Type, 1978-2007. Supplemental Tables to the Annual Energy Outlook 2010. Technical report, Energy Information Administration.
- European Commission (2005). EU action against climate change: EU emissions trading—an open scheme promoting global innovation. http://ec.europa.eu/environment/climat/pdf/emission_trading2_en.pdf.
- European Commission (2012). Emissions trading: annual compliance round-up shows declining emissions in 2011. Press Release IP/12/477.
- Falkner, R. (2008). *Business power and conflict in international environmental politics*. Palgrave Macmillan.
- Fazekas, D. (2008). Hungarian Experiences with the EU ETS. *Unpublished working paper available at: [http://www.cdclimat.com/IMG/pdf/33b - Hungarian_experience_with_EU_ETS_carbon_market_implications_Fazekas.pdf](http://www.cdclimat.com/IMG/pdf/33b_Hungarian_experience_with_EU_ETS_carbon_market_implications_Fazekas.pdf)*.
- Fezzi, C. and Bunn, D. (2010). Structural Analysis of Electricity Demand and Supply Interactions*. *Oxford Bulletin of Economics and Statistics*, 72(6):827–856.
- Fiorini, A., Georgakaki, A., Lepsa, B.-N., Pasimeni, F., and Salto Saura Maria, L. (2016). Estimation of Corporate R&D investment in Low-Carbon Energy Technologies: A Methodological Approach. Technical Report JRC101552, Joint Research Centre of the European Commission.
- Fischer, C., Parry, I., and Pizer, W. (2003). Instrument choice for environmental protection when technological innovation is endogenous. *Journal of Environmental Economics and Management*, 45(3):523–545.
- Fowlie, M., Holland, S. P., and Mansur, E. T. (2012). What Do Emissions Markets Deliver and to Whom? Evidence from Southern California’s NO_x Trading Program. *The American Economic Review*, 102(2):965–993.
- Gagelmann, F. and Frondel, M. (2005). The impact of emission trading on innovation—science fiction or reality? *European Environment*, 15(4):203–211.
- Gorman, H. and Solomon, B. (2002). The Origins and Practice of Emissions Trading. *Journal of Policy History*, 14(3):293–320.
- Greenstone, M. and Gayer, T. (2009). Quasi-experimental and experimental approaches to environmental economics. *Journal of Environmental Economics and Management*, 57(1):21–44.

- Grubb, M., Azar, C., and Persson, U. (2005). Allowance allocation in the European emissions trading system: a commentary. *Climate Policy*, 5(1):127–136.
- Hagedoorn, J. and Cloudt, M. (2003). Measuring innovative performance: is there an advantage in using multiple indicators? *Research Policy*, 32(8):1365–1379.
- Hanemann, M. (2010). Cap-and-trade: a sufficient or necessary condition for emission reduction? *Oxford Review of Economic Policy*, 26(2):225–252.
- Hartman, E. and Hidalgo, F. D. (2014). What’s the Alternative?: An Equivalence Approach to Balance and Placebo Tests. *UC Berkeley, Department of Political Science Working paper*.
- Hirshleifer, D., Hsu, P.-H., and Li, D. (2013). Innovative efficiency and stock returns. *Journal of Financial Economics*, 107(3):632–654.
- Ho, D. E. and Imai, K. (2006). Randomization Inference With Natural Experiments. *Journal of the American Statistical Association*, 101:888–900.
- Hoffmann, V. H. (2007). EU ETS and Investment Decisions: The Case of the German Electricity Industry. *European Management Journal*, 25(6):464–474.
- Iacus, S., King, G., and Porro, G. (2009). Causal inference without balance checking: Coarsened exact matching. *Retrieved September*.
- Ioannidis, J. and Doucouliagos, C. (2013). What’s to Know About the Credibility of Empirical Economics? *Journal of Economic Surveys*, 27(5):997–1004.
- Jaffe, A., Newell, R., and Stavins, R. (2005). A tale of two market failures: Technology and environmental policy. *Ecological Economics*, 54(2-3):164–174.
- Jaffe, A. B., Newell, R. G., and Stavins, R. N. (2003). Chapter 11 Technological change and the environment. In Karl-Göran Mäler and Jeffrey R. Vincent, editor, *Handbook of Environmental Economics*, volume Volume 1 of *Environmental Degradation and Institutional Responses*, pages 461–516. Elsevier.
- Jaffe, A. B. and Palmer, K. (1997). Environmental Regulation and Innovation: A Panel Data Study. *The Review of Economics and Statistics*, 79(4):610–619.
- Jaffe, A. B., Peterson, S. R., Portney, P. R., and Stavins, R. N. (1995). Environmental Regulation and the Competitiveness of U.S. Manufacturing: What Does the Evidence Tell Us? *Journal of Economic Literature*, 33(1):132–163. ArticleType: research-article / Full publication date: Mar., 1995 / Copyright © 1995 American Economic Association.
- Jaraite, J. and Maria, C. D. (2011). Did the EU ETS Make a Difference? A Case Study of Ireland and Lithuania. *Presented at EARE 2011, cited with author permission*.
- Jaraite-Kažukauske, J. and Di Maria, C. (2016). Did the EU ETS Make a Difference? An Empirical Assessment Using Lithuanian Firm-Level Data. *The Energy Journal*, 37(1).
- Johnstone, N., Haščič, I., and Popp, D. (2010). Renewable energy policies and technological innovation: Evidence based on patent counts. *Environmental and Resource Economics*, 45(1):133–155.
- Keele, L. and Miratrix, L. (2015). Randomization inference for outcomes with clumping at zero. *Working paper*, pages 1–36.

- Kerr, S. and Newell, R. G. (2003). Policy-Induced Technology Adoption: Evidence from the U.S. Lead Phasedown. *The Journal of Industrial Economics*, 51(3):317–343.
- Klemetsen, M. E., Rosendahl, K. E., Jakobsen, A. L., and others (2016). The impacts of the EU ETS on Norwegian plants’ environmental and economic performance. Technical report.
- Kneese, A. V. and Schultze, C. (1975). Pollution, Prices, and Public Policy. (*Brookings Institution, Washington, DC*).
- Kreutzer, J. (2006). Cap and Trade: A Behavioral Analysis of the Sulfur Dioxide Emissions Market. *New York University Annual Survey of American Law*, 62:125.
- Lanjouw, J. O. and Schankerman, M. (2004). Patent Quality and Research Productivity: Measuring Innovation with Multiple Indicators. *The Economic Journal*, 114(495):441–465.
- List, J. A., Millimet, D. L., Fredriksson, P. G., and McHone, W. W. (2003). Effects of Environmental Regulations on Manufacturing Plant Births: Evidence from a Propensity Score Matching Estimator. *Review of Economics and Statistics*, 85(4):944–952.
- Löfgren, Å., Wråke, M., Hagberg, T., and Roth, S. (2013). The Effect of EU-ETS on Swedish Industry’s Investment in Carbon Mitigating Technologies. *University of Gothenburg Working Papers in Economics*, (565).
- Martin, B. R. (1996). The use of multiple indicators in the assessment of basic research. *Scientometrics*, 36(3):343–362.
- Martin, R., de Preux, L., and Wagner, U. (2009). The impacts of the Climate Change Levy on business: evidence from microdata. *CEP Discussion Paper*, (917).
- Martin, R., Muûls, M., de Preux, L. B., and Wagner, U. J. (2012). Anatomy of a paradox: Management practices, organizational structure and energy efficiency. *Journal of environmental economics and management*, 63(2):208–223.
- Martin, R., Muuls, M., and Wagner, U. (2011). Climate change, investment and carbon markets and prices – evidence from manager interviews. *Climate Strategies, Carbon Pricing for Low-Carbon Investment Project*.
- Mauritzen, J. (2014). Scrapping a Wind Turbine: Policy Changes, Scrapping Incentives and Why Wind Turbines in Good Locations Get Scrapped First. *Energy Journal*, 35(2):157–181.
- Miller, S. (2014). Indirectly Induced Innovation: Consequences for Environmental Policy Analysis.
- Milliman, S. and Prince, R. (1989). Firm incentives to promote technological change in pollution control. *Journal of Environmental Economics and Management*, 17(3):247–265.
- Naegele, H. (2015). Offset Credits in the EU ETS : A Quantile Estimation of Firm-Level Transaction Costs. *DIW Berlin Discussion Paper*.
- Newell, R., Jaffe, A., and Stavins, R. (1999). The Induced Innovation Hypothesis and Energy-Saving Technological Change. *Quarterly Journal of Economics*, 114(3):941–975.
- OECD (2009). OECD Patent Statistics Manual. Technical report, OECD.
- OECD (2015). *Frascati Manual 2015*. Organisation for Economic Co-operation and Development, Paris.

- Office for National Statistics (2011). *Business Expenditure on Research and Development, 1994-2008: Secure Data Service Access [computer file]*. Colchester, Essex: UK Data Archive [distributor]. SN: 6690, <http://dx.doi.org/10.5255/UKDA-SN-6690-1>.
- Office for National Statistics (2012). *Business Structure Database, 1997-2011: Secure Data Service Access [computer file]*. Colchester, Essex: UK Data Archive [distributor], 3rd edition. SN: 6697.
- Palmer, K., Oates, W. E., and Portney, P. R. (1995). Tightening Environmental Standards: The Benefit-Cost or the No-Cost Paradigm? *Journal of Economic Perspectives*, 9(4):119–132.
- Petrack, S. and Wagner, U. J. (2014). The Impacts of Cap-and-Trade on Industry: Evidence from the European Carbon Market and German Manufacturing Plants. *Kiel Working Papers*, (1912).
- Popp, D. (2002). Induced innovation and energy prices. *The American Economic Review*, 92(1):160–180.
- Popp, D. (2003). Pollution control innovations and the Clean Air Act of 1990. *Journal of Policy Analysis and Management*, 22(4):641–660.
- Popp, D. (2006). International innovation and diffusion of air pollution control technologies: the effects of NO_x and SO₂ regulation in the US, Japan, and Germany. *Journal of Environmental Economics and Management*, 51(1):46–71.
- Popp, D. and Newell, R. (2012). Where does energy R&D come from? Examining crowding out from energy R&D. *Energy Economics*, 34(4):980–991.
- Porter, M. E. (1991). Essay: America’s green strategy. *Scientific American*, 264(3).
- Requate, T. (2005). Dynamic incentives by environmental policy instruments—a survey. *Ecological economics*, 54(2-3):175–195.
- Rockström, J., Gaffney, O., Rogelj, J., Meinshausen, M., Nakicenovic, N., and Schellnhuber, H. J. (2017). A roadmap for rapid decarbonization. *Science*, 355(6331):1269–1271.
- Rosenbaum, P. R. (2009). *Design of Observational Studies*. Springer.
- Rosenbaum, P. R. (2013). *Observational Studies*. Springer Science & Business Media.
- Schleich, J. and Betz, R. (2005). Incentives for energy efficiency and innovation in the European Emission Trading System. In *Proceedings of the 2005 ECEEE Summer Study—What Works and Who Delivers? Mandelie*, pages 1495–1506.
- Schmalensee, R., Joskow, P., Ellerman, A., Montero, J., and Bailey, E. (1998). An interim evaluation of sulfur dioxide emissions trading. *The Journal of Economic Perspectives*, 12(3):53–68.
- Schmidt, T. S. and Sewerin, S. (2017). Technology as a driver of climate and energy politics. *Nature Energy*, 2(6):nenergy201784.
- Sekhon, J. (2007). Multivariate and propensity score matching software with automated balance optimization: The matching package for R. *Journal of Statistical Software*, 10(2):1–51.
- Stavins, R. (2007). A US cap-and-trade system to address global climate change. *Regulatory Policy Program Working Paper*, (RPP-2007-04).

- Stavins, R. N. (1995). Transaction costs and tradeable permits. *Journal of Environmental Economics and Management*, 29(2):133 – 148.
- Taylor, M. R. (2012). Innovation under cap-and-trade programs. *Proceedings of the National Academy of Sciences*, 109(13):4804–4809. PMID: 22411797.
- Taylor, M. R., Rubin, E. S., and Hounshell, D. A. (2003). Effect of Government Actions on Technological Innovation for SO₂ Control. *Environmental Science & Technology*, 37(20):4527–4534.
- Taylor, M. R., Rubin, E. S., and Hounshell, D. A. (2005). Regulation as the mother of innovation: The case of SO₂ control. *Law & Policy*, 27:348.
- Tomás, R., Ribeiro, F. R., Santos, V., Gomes, J., and Bordado, J. (2010). Assessment of the impact of the European CO₂ emissions trading scheme on the Portuguese chemical industry. *Energy Policy*, 38(1):626–632.
- Veefkind, V., Hurtado-Albir, J., Angelucci, S., Karachalios, K., and Thumm, N. (2012). A new EPO classification scheme for climate change mitigation technologies. *World Patent Information*, 34(2):106–111.
- Wagner, U., Muûls, M., Martin, R., and Colmer, J. (2014). The Causal Effects of the European Union Emissions Trading Scheme: Evidence from French Manufacturing Plants. *Unpublished paper*.
- Widerberg, A. and Wråke, M. (2009). The Impact of the EU Emissions Trading System on CO₂ Intensity in Electricity Generation. *University of Gothenburg Working Papers in Economics*, (361).
- Wiesenthal, T., Leduc, G., Haegeman, K., and Schwarz, H.-G. (2012). Bottom-up estimation of industrial and public R&D investment by technology in support of policy-making: The case of selected low-carbon energy technologies. *Research Policy*, 41(1):116–131.
- Wiesenthal, T., Leduc, G., Schwarz, H.-G., and Haegeman, K. (2009). R&D investment in the priority technologies of the European Strategic Energy Technology Plan. *Joint Research Centre. Brussels: European Commission*.
- World Bank; Ecofys (2017). Carbon pricing watch 2017. <https://openknowledge.worldbank.org/handle/10986/26565>.
- Zechter, R., Kerr, T., Kossoy, A., Peszko, G., Oppermann, K., Ramstein, C., Prytz, N., Klein, N., Lam, L., Wong, L., Blok, K., Neelis, M., Monschauer, Y., Nierop, S., Berg, T., Ward, J., Kansy, T., Kemp, L., Vadheim, B., and Kingsmill, N. (2016). State and Trends of Carbon Pricing 2016. Technical report, World Bank.

For Online Publication

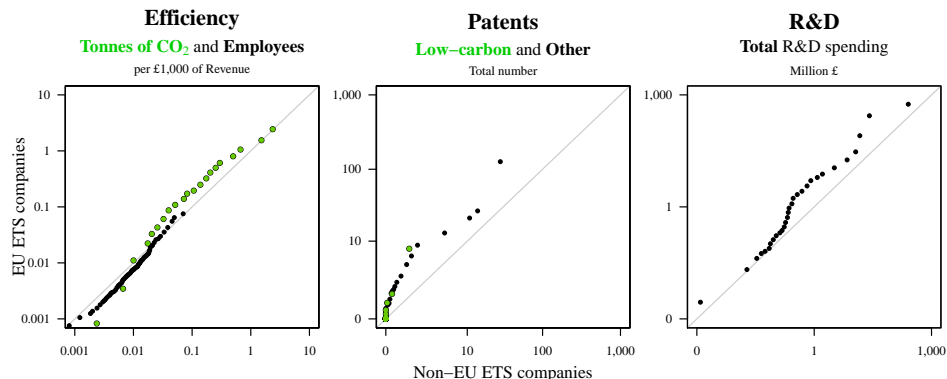
Table 4 and figure 10 shows the pre-EU ETS covariate balance for a matched sample of 445 EU ETS companies and 473 non-EU ETS companies. The distribution of covariates is unbalanced in several important dimensions. The regulated are substantially larger than their unregulated counterparts, though the differences are generally not statistically significant. The statistical evidence is stronger that the two groups of companies have different pre-treatment CO₂-intensities, patenting, R&D spending, and regulatory histories. These differences are driven by the 42 EU ETS companies that were omitted from the main analysis, along with their matches. While treatment status and outcomes is plausibly random for the smaller sample, the full sample appears to include regulated companies that would not plausibly have been left unregulated in any incarnation of the EU ETS. This makes it difficult to think of a reasonable basis for constructing empirical counterfactuals for these companies, which is why they were omitted from the main analysis. I repeat the analysis here including these companies for completeness.

Table 4: Equivalence tests for all EU ETS ($N = 445$) and matched non-EU ETS companies ($N = 473$)

	Median difference	Mean difference	Equivalence range	Signed-rank test p -value	Paired t -test p -value
Company basics					
Revenues (£1,000s)	16,198	311,030	±543,983	<.001	<.001
Employees	70	1,192	±1,216	<.001	0.473–
Economic sector (3-digit)	<i>Exact match</i>	–	–	–	–
Efficiency					
CO ₂ -intensity	0.035	0.157	±0.103	.531	.892
Labour intensity	-0.001	-0.001	±0.002	.002	.003
Patenting					
Low-carbon patents	0	0.099	±0.160	.994 ^{LP}	.127
Other patents	0	1.780	±2.284	.010	.243
R&D spending					
Low-carbon patent share	0	0.009	±0.009	.999 ^{LP}	.423
R&D spending (£1,000s)	147	7,091	±13,267	<.001	.126
Regulatory history					
CCA participation	1	0.120	–	–	–
UK ETS participation	0	0.025	–	–	–
CCL bill (£1,000)	23,413	37,362	±33,215	.620	.622
R&D support (£1,000)	1.001	-188.782	±2,047.122	<.001	<.001

Figure 11 and table 5 show the results for the full sample of matched companies. The

Figure 10: Comparison of all matched EU ETS and non-EU ETS companies



trends and estimates are qualitatively consistent with those reported in the main analysis, although the exact magnitudes and p -values generally suggest there is more evidence against reductions in CO₂-intensity and that the effects on innovation may have been even larger. As I have outlined earlier, however, the interpretation of these results is problematic, and they are reported here primarily for completeness.

Figure 11: Adoption and innovation of all matched EU ETS and non-EU ETS companies

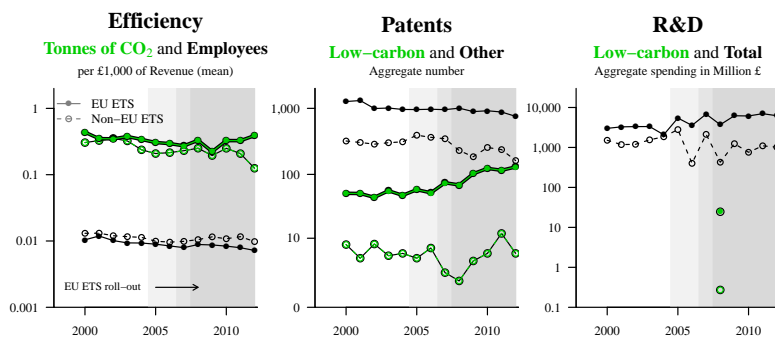


Table 5: Matching estimates of the effects of the EU ETS

	Hodges-Lehmann point estimate	<i>p</i> -value
Efficiency (proportion)		
CO ₂ -intensity	0.404	0.008
Labour intensity	-0.078	0.660
Patenting (number of patents)		
Low-carbon patents	0.680	<0.001
Other patents	0.740	<0.001
R&D spending (£1,000s)		
Low-carbon R&D	521.000	0.181
Total R&D	1,024.000	<0.001