Environmental Policy and Directed Technological Change: Evidence from the European carbon market*

Raphael Calel[1] and Antoine Dechezleprêtre[1,2]

[1]Grantham Research Institute on Climate Change and the Environment, London School of Economics
[2]Centre for Economic Performance, London School of Economics

March 2012

Abstract

The European Union Emissions Trading Scheme (EU ETS) has aimed to encourage the development of low-carbon technologies by putting a price on carbon emissions. Using a newly constructed data set that links 8.5 million European companies with their patenting history and their regulatory status under EU ETS, we investigate the hypothesis that the EU ETS has encouraged development of low-carbon technologies. Exploratory data analysis reveals a rapid increase in low-carbon patenting activities at the EPO since 2005, especially among EU ETS regulated companies during the Scheme's second phase. Naive estimates obtained by comparing EU ETS and non-EU ETS firms suggest that the Scheme may be responsible for up to 30% of the increase in low-carbon patenting of regulated companies. However, more refined estimates that combine matching methods with difference-in-differences provide evidence that the EU ETS has not impacted the direction of technological change. This finding appears to be robust to a number of stability and sensitivity checks. While we cannot completely rule out the possibility that the EU ETS has impacted only large companies for which suitable unregulated comparators cannot be found, our findings suggest that the EU ETS so far has had at best a very limited impact on low-carbon technological change.

JEL: O31, Q55, Q58

Keywords: Directed technological change, EU Emissions Trading Scheme, Policy evaluation

---

*E-mail: `r.calel@lse.ac.uk` and `a.dechezlepretre@lse.ac.uk`. Raphael Calel is grateful for funding provided by the ESRC and the Jan Wallander and Tom Hedelius Foundation. Antoine Dechezleprêtre gratefully acknowledges the support of ADEME and of the ESRC under the ESRC Postdoctoral Fellowship Scheme (award no: PTA-026-27-2756). The authors would also like to thank Xavier Vollenweider for excellent research assistance.

1 Introduction

The European Union Emissions Trading Scheme (EU ETS) was launched in 2005 and is today the world's largest carbon market. Under the scheme, around 11,000 power stations and industrial plants in 30 countries are allocated tradable emissions permits, covering 40% of the EU's total greenhouse gas emissions.[1] As the main instrument of the EU's policy to mitigate climate change, the EU ETS was primarily intended to reduce carbon emissions. However, when regulated firms expect to face a higher price on emissions relative to other costs of production, this provides them with an incentive to make operational changes and investments that reduce the emissions intensity of their output. The "induced innovation" hypothesis, dating back to Sir John Hicks (1932) and restated in the context of environmental policy by Porter (1991), suggests that part of this new investment will be directed toward developing and commercializing new emissions-reducing technologies. The EU ETS can therefore be expected to spur development of new low-carbon technologies. This vision has been articulated many times by EU policy makers, who envisage the EU ETS to be a driving force of the transition to a low-carbon economy.

Over the past few decades a considerable theoretical literature has developed the induced innovation hypothesis, in particular in the context of climate change mitigation (Goulder and Schneider, 1999; van der Zwaan et al., 2002; Popp, 2004, 2006a; Gerlagh, 2008; Acemoglu et al., 2012). The impact of environmental policy on technological change may be the greatest determinant of the long-run cost of emissions abatement, and hence, perhaps one of the most important criteria on which to judge its success (Kneese and Schultze, 1975). In light of this, there is an ongoing research effort to empirically understand and quantify the link between environmental policies and technological change (see Popp et al., 2009, Popp, 2010, and Ambec et al., 2010, for recent surveys). Our study contributes to this literature.

In this paper we investigate the impact of the EU ETS on low-carbon technological change in the first 5 year of the Scheme's existence. Previous studies have found that most of the induced innovation response is observed in the first 5 years following the introduction of a new environmental policy (Popp, 2002). We examine a newly constructed data set that records firms' regulatory status with respect to the EU ETS, basic firm characteristics, and patenting activities. The new low-carbon patent classification recently developed by the European Patent Office (EPO) allows us to precisely identify emissions reduction technologies. Our data set covers over 8.5 million firms across 22 EU

[1] 24 countries were included from the beginning. 6 countries have joined since then.

countries. We identify over 5'500 firms that are directly regulated under the EU ETS, accounting for nearly 80% of EU ETS covered emissions and installations. The data stretches back to before 2005, so that we are able to compare unregulated and would-be regulated firms both before and after the EU ETS was implemented.

The EU ETS offers a unique opportunity to investigate the impact of environmental policy on technological change. To be sure, the EU ETS is the first and largest environmental policy initiative of its kind anywhere in the world, which makes it an interesting large-scale case study. But more importantly, in order to control administrative costs the EU ETS was designed to cover only large installations. Firms operating smaller installations are not covered by EU ETS regulations. Because environmental regulations create stronger incentives for innovation among the regulated firms (Milliman and Prince, 1989; Fischer et al., 2003), we can detect the impact of the EU ETS on technological change by comparing regulated companies with otherwise similar but unregulated companies. Using both exploratory data analysis and matching methods that enable us to control for firm heterogeneity and for factors that affect both regulated and unregulated firms (input prices, sector- and country-specific policies, etc.), we provide the first comprehensive empirical assessment of the impact the EU ETS on low-carbon technological change.

Exploratory data analysis reveals a rapid increase in low-carbon patenting activities since 2005. This increase seems to have disproportionately affected EU ETS regulated companies, especially during the more stringent second trading phase that started in 2008. Naive estimates obtained by comparing EU ETS and non-EU ETS firms suggest that the Scheme may be responsible for up to 30% of the increase in low-carbon patenting of regulated companies in the first 5 years of the EU ETS compared to the 5 years preceding the new regulations. More refined estimates that combine matching methods with difference-in-differences, however, provide evidence that the EU ETS has not impacted the direction of technological change. This finding appears to be robust to a number of stability and sensitivity checks. While we cannot completely rule out the possibility that the EU ETS has impacted only large companies for which suitable unregulated comparators cannot be found, our findings suggest that the EU ETS so far has had at best a very limited impact on low-carbon technological change.

The EU ETS is expected to remain an integral part of the EU's strategy for building a low-carbon Europe (European Commission, 2011). Our results suggest, however, that the EU ETS in its current form might not be providing strong enough incentives for low-carbon technological change. This may have important policy implications for the EU's low-carbon strategy going forward, as well as other regulatory carbon market programs now being implemented in New Zealand, the North-Eastern United States, Australia, and elsewhere.

The paper proceeds as follows. Section 2 surveys the evidence on environmental policy and directed technological change, especially in the context of emissions trading. Evidence from the US Acid Rain Program and early studies of the EU ETS help us develop expectations of how the EU ETS is likely to have impacted technological change. We also discuss to what extent our patent data enables us to measure technological change. In section 3 we familiarize ourselves with our newly constructed data set, and use the data to begin unpacking the characteristics of low-carbon technological change. In section 4 we turn our eye to estimating the impact of the EU ETS, and section 5 summarizes the evidence and offers some concluding observations.

2 Detecting directed technological change

2.1 Previous empirical studies

Several studies have found evidence that environmental policy does impact the direction of technological change (Lanjouw and Mody, 1996; Brunnermeier and Cohen, 2003; Popp, 2002, 2003, 2006b; Arimura et al., 2007; Lanoie et al., 2007). Popp (2006b) finds an almost immediate patenting response to domestic clean air regulations in the US, Germany, and Japan. Johnstone et al. (2010) find that renewable energy patents have increased dramatically as national and international climate change policies have multiplied.

When examining the impact of emissions trading specifically the conclusions are more modest. Most studies concern the US Clean Air Act's Acid Rain Program, launched in 1995. Early estimates suggested nearly half of the emissions reductions were achieved by installing scrubber technology, and the remainder by switching to coal with a lower sulphur content (Schmalensee et al., 1998). The scrubber technology existed before 1995, but had in many instances not been economically viable. The "innovation" resulting from the Acid Rain Program appears to have been focused on operational rather than technological change, therefore (Burtraw, 2000). There is nevertheless some evidence of very narrowly directed technological change. Popp (2003) detects an increase in patents that improved the efficiency of scrubbers.[2] This effect was confined to early years under the new regime though, and the Program has not provided ongoing incentives for

[2]It is worth noting that Title IV of the Clean Air Act, which establishes the Acid Rain Program, also includes special provisions that reward firms specifically for the use of scrubbers. It is not entirely clear, therefore, how much was "the market's doing".

technological advancement (Lange and Bellas, 2005). This squares with findings that the use of scrubber technology as an emissions abatement strategy has actually declined over time (Burtraw and Szambelan, 2009).

2.2 The EU ETS and technological change

Since its launch in 2005, there has been vigorous debate about whether the EU ETS would induce firms to develop new emissions-reducing technologies, many arguing that it would not because of an overly generous allocation of emissions permits (Schleich and Betz, 2005; Gagelmann and Frondel, 2005; Grubb et al., 2005). Indeed, a few early case studies summarized by Petsonk and Cozijnsen (2007) indicate that rather than developing new technologies, firms have been introducing well-known technological solutions that had simply not been economically viable before the EU ETS imposed a carbon price on regulated firms.

A growing literature of interviews and case-studies provides support for this conclusion. Tomás et al. (2010) study four large EU ETS regulated Portuguese chemical companies, suggesting that the EU ETS may have encouraged some emissions reducing innovation, but largely in the form of energy efficiency improvements. Hoffmann (2007), reporting on case studies in the German electricity sector, find that the EU ETS has an effect only on innovation decisions with a short time horizon. Development of new low-carbon technologies generally does not fall into this category. Martin et al. (2011) conduct semi-structured interviews with nearly 800 European manufacturing firms, of which almost 450 operated EU ETS regulated installations, finding a positive effect of the EU ETS on process innovation, but not on product innovation. The former involves operational innovations to a greater degree, while the latter is more closely associated with technological advancement. Similarly, a survey of Irish EU ETS firms tentatively suggests that almost no resources were available for low-carbon R&D in Phase 1 of the trading program (2005–2007), while many of the firms had pursued operational innovations like installing new machinery or equipment, making process or behavioral changes, and/or employing fuel switching to some degree (Anderson et al., 2011).

The practice of fuel switching in particular appears to have been very important so far. Fuel switching requires neither capital investment nor R&D, only that power providers bring less polluting gas-fired plants online before coal-fired ones as demand ramps up. This changes the average fuel mix in favor of natural gas, and therefore reduces the carbon intensity of output. Fuel switching is a purely organizational innovation. Macroeconomic estimates suggest that the EU ETS reduced total emissions by roughly

50–100 million tons of CO_2 annually in Phase 1, or roughly 3–6%, compared with a "business-as-usual" scenario (Ellerman and Buchner, 2008; Anderson and Di Maria, 2011). Meanwhile, model-based estimates of power sector emissions abatement from fuel switching alone range from 26–88 million tons (Delarue et al., 2008, 2010). These estimates suggest that fuel switching very likely accounts for the lion share of emissions reductions in the EU ETS so far. While this is not a problem in and of itself, one should be conscious that the capacity for emissions reductions through fuel switching is very limited compared to the EU's longer term targets. Delarue et al. (2008) estimate that fuel switching has the potential to reduce emission by up to 300 million tons per year, which is no more than a tenth of what is needed to meet the EU target to cut emissions by 80% by 2050 against 1990-levels.[3]

These observations motivate a special interest in whether the EU ETS is encouraging firms to develop new technologies that can help achieve the ambitious long term targets. Though most of the evidence so far suggests companies are employing fuel switching and other short-term strategies to reduce emissions, at least during Phase 1 of the scheme (2005-2007) for which studies are available, it is conceivable that they simultaneously start to develop new technologies that will facilitate future emissions reductions. Yet, there is little systematic evidence so far as to what impact the EU ETS is having on low-carbon technological change. This paper uses a newly constructed patent data set to help answer this question.

2.3 Measuring technological change with patents

We use patent applications to focus on that part of "innovation" concerned with technological change. Patents have been used extensively as a measure of technological change in the recent induced innovation literature (Popp, 2002, 2006b; Johnstone et al., 2010), and the advantages and drawbacks of patents are now well-understood (see OECD, 2009, for a recent survey). Patents provide a useful measure of the output of innovative activity and are available at a highly disaggregated technological level. Moreover, recent advances in linking patent data with company data makes it possible to construct firm-level patent portfolios.

Our patent data are drawn from the World Patent Statistical Database (PATSTAT) maintained by the European Patent Office (EPO). PATSTAT is the world's most comprehensive patent data set, including over 60 million patent documents from 80 patent offices. The PATSTAT database reports the name of patent applicants. We use this

[3]The EU target amounts to reducing annual emissions by roughly 4'500 million tons compared to 1990, or roughly 3'500 million tons compared to current emission levels.

information to link patents with the companies that hold them. Note that while the EU ETS regulations apply at the level of the installation, patents measure innovation at the level of the firm.

As our main measure of innovation activity we use patent filings with the EPO. EPO patents provide a common measure of innovation for all of Europe, unlike self-reported innovation measures or patents filed with national patent offices, for which the standards vary from firm-to-firm or country-to-country. In addition, EPO patents provide a useful quality threshold as only high value inventions typically get patented at the EPO.[4] However, in some of our robustness tests we also look at patents filed with national patent offices to gauge whether our findings depend on how narrowly we define the patents of interest.

All patents filed at the EPO are categorized using 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. These low-carbon technologies include, to name a few, efficient combustion technologies (e.g. combined heat and power generation) and energy storage (e.g. fuel cells). This class helps us measure the direction of technological change. Because the EPO low-carbon classification is not comprehensive, we also test the robustness of our results to the inclusion of additional patents that other authors have considered low-carbon, in particular patents pertaining to energy-efficient industrial processes.[5] The precise description of the various sub-classes for low-carbon patents used in the paper can be found in appendix C.

Finally, because the value of patents is known to be heterogeneous, we also check the robustness of our analysis to using quality-weighted patent counts. We use two ways to account for the quality of patents: forward citations and family size. Citation data have been widely used in the literature to control for the quality of patents. With this method, patents are weighted by the number of times each of them is cited in subsequent patents (see Trajtenberg, 1990; Harhoff et al., 1999; Hall et al., 2005). The family of a patent is the set of patents protecting the same invention in various countries.[6] Counting the number of countries in which a patent is filed is another common measure of patent quality (Harhoff et al., 2003; van Zeebroeck, 2011). Family data also presents the

[4] Evidence shows that the highest value technologies are patented in several countries (Harhoff et al., 2003), and indeed, one of the methods used to measure the value of patents is to count the number of countries is which they are filed (van Zeebroeck, 2011). Patents filed at the EPO get patented in 6 EPO member countries on average.

[5] An updated list of environment-related patent classification codes is available from the OECD's Environmental Policy and Technological Innovation (EPTI) website: www.oecd.org/environment/innovation.

[6] Patent family information comes from the DOCDB family table in PATSTAT.

advantage of being more rapidly available than citations [7], which is especially valuable when dealing with very recent patents as we do.

3 Unpacking low-carbon technological change

The EPO was created in 1978. Since then, over 2.5 million patents have been filed with the EPO, of which just over 50'000 (or 2%) have been classified as low-carbon inventions. The number of patents for these low-carbon technologies shows a distinct pattern over time. There was a surge in patenting for these technologies in the early 1980s, widely attributed to the second oil price shock in the late 1970s (Dechezleprêtre et al., 2011). The number of low-carbon patents filed each year then stayed roughly level until the mid-1990s, after which it began to rise again. The number has increased rapidly in recent years, with over 32'000 of these 50'000 patents filed since the year 2000. Patents for low-carbon technologies have also been rising rapidly as a share of all patents (see figure 1). This trend is particularly evident after 2005, with the share doubling from 2% to 4% in just a few years. A simple Chow test strongly rejects the hypothesis that there is no structural break in 2005 ($P < 0.001$).

Figure 1: Share of low-carbon patents (1978–2009)

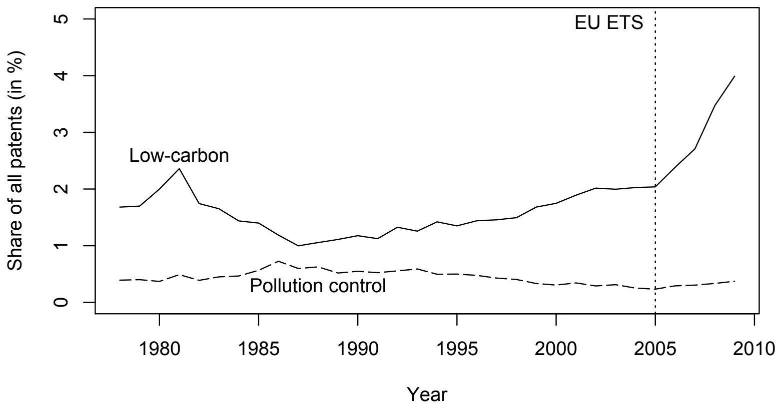

While this pattern is robust to using an expanded definition of "low-carbon technologies", it does not apply to environmentally friendly technologies in general. To see this, figure 1 also plots the share of patents for non-greenhouse gas "pollution control

[7] Patents are typically mostly cited two years after their publication, hence four years after they are first filed.

technologies", as defined by Popp (2006b),[8] which does not display the same structural break (one cannot reject the hypothesis of no structural break in 2005 at conventional significance levels). The sudden surge in patenting activity, therefore, appears to be specific to low-carbon technologies and to coincide with the launch of the EU ETS. Could the structural break in low-carbon patenting, then, be a consequence of the EU ETS?

Just like the increase in low-carbon patenting in the early 1980s was due to the second oil price shock, the recent upsurge in patenting could be due to increasing oil prices. When comparing the share of low-carbon patenting with the evolution of oil prices (see figure 2), one notices that the current surge in patenting follows immediately on the heels of the rapid oil price increases staring in the early 2000s. Patenting for pollution control, on the other hand, was not responsive to the oil price in the 1980s, and so it is not surprising it has stayed flat recently. Looking at the aggregate trends over time, therefore, is not enough to determine whether the increase in low-carbon patenting since 2005 is the result of the EU ETS, oil prices, or some other factor. In order to isolate the impact of the EU ETS we must compare the experience of firms regulated by the EU ETS with those not covered by the regulation. Both groups will have faced the same oil prices and other macroeconomic conditions, but starting in 2005 they were subject to different regulatory regimes.

Figure 2: Share of low-carbon patents and Crude oil price(1978–2009)

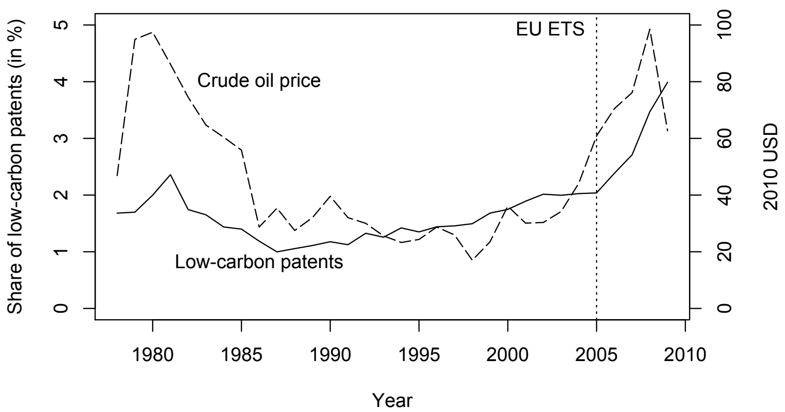

Our newly constructed data set includes roughly 8.5 million firms in 22 EU countries, operating in sectors of the economy covered to some degree by EU ETS regulations, including the power sector, iron and steel, cement, glass, pulp and paper, etc. 18 of

[8]These technologies pertain to reduction of local pollutants including SO_2 and NO_X.

these countries launched the EU ETS in 2005, and the other 4 (Norway, Switzerland, Romania, and Bulgaria) either have either joined later or have stayed outside of the EU ETS altogether. Our data set is therefore more restricted in terms of geographical coverage than the EPO and also covers fewer sectors than the EPO does (see appendix A for details). Nevertheless the firms in our data set have filed nearly 760'000 patents altogether, which accounts for just under 30% of all patents filed at the EPO, and 63% of all patents filed at the EPO by firms located in one of the 22 countries in our dataset. Since EU ETS regulated sectors only represent a share of the economy in these countries, this very high percentage gives us great confidence that we managed to include the patent history of a vast majority of the companies operating in EU ETS regulated sectors.[9] Of these, just over 14'000 (or 1.9%) were climate related (compared to 2% for the EPO as a whole), which again represents around 30% of all low-carbon patents filed at the EPO.[10]

Our data set also records the regulatory status of these 8.5 million firms. 5'521 firms in our data set are EU ETS regulated. They account for over 90% of EU ETS regulated installations and emissions in Phase 1 in the 18 EU ETS countries, and nearly 80% of regulated installations and emissions EU ETS-wide (see table 1). Because our data set records the firms' regulatory status, we can compare the patenting activities of EU ETS and non-EU ETS firms.

Figure 3 decomposes the overall trend seen in figure 1 into its EU ETS and non-EU ETS components. First, it is reassuring that the pattern does not alter notably. In fact, the correlation coefficient of the low-carbon patent share in EPO and in our data set is 0.94. Second, at least visually it appears as if the share of low-carbon patents was roughly the same among EU ETS and non-EU ETS firms in the 5 years before the EU ETS launched, and also during Phase 1 of the EU ETS (2005–2007). As mentioned in section 2, it was widely argued that too many emissions permits were issued for Phase 1, and that this meant firms would have little incentive to develop new abatement technologies. The second trading phase began in 2008 and was widely expected to constrain emissions more tightly. Coinciding with the start of Phase 2, the share of low-carbon patents among EU ETS firms rose more rapidly than among non-EU ETS firms.

[9] We have also conducted extensive manual double-checking, so we can reasonably assume that companies for which we were unable to find any patent data have actually not filed any patent at the EPO. It is well documented that only a fraction of companies ever file patents, and this is likely to be especially true of the EPO that has high administrative costs.

[10] The companies in our data set account for 32% of all patents filed at the EPO since 2005 (29% since 1978), and 31% of all climate related patents (27% since 1978).

Table 1: Coverage of the EU ETS – The first two columns of this table show the number of Phase 1 installations in each of the 18 countries in our sample, and their allocated emissions (source: CITL). The following two columns show the percentages of installations and emissions for which the operating firm has been identified. The two rows at the foot of the table summarise our data set's EU ETS coverage for our 18 countries as well as as a proportion of the EU ETS as a whole.

	Number of installations	Mton of emissions	Percent of installations covered	Percent of emissions covered
Austria	217	97.8	91.7	100.0
Belgium	347	178.7	98.6	100.0
Czech Rep.	415	290.8	92.5	96.9
Denmark	399	93.1	92.7	95.2
Estonia	54	56.3	66.7	15.3
Finland	637	133.9	80.7	95.4
France	1100	450.2	97.5	99.6
Germany	1944	1486.3	97.2	99.4
Ireland	121	57.7	76.0	94.6
Lithuania	113	34.4	85.8	91.4
Luxembourg	15	9.7	66.7	78.1
Netherlands	419	259.3	58.0	54.1
Poland	869	712.7	90.0	98.6
Portugal	265	110.7	97.0	98.9
Slovakia	191	91.4	86.9	99.3
Spain	1072	498.1	97.9	98.3
Sweden	774	67.6	93.9	98.8
UK	1107	628.0	86.3	97.7
Total	10059	5256.6	91.3	95.4
Total EU ETS	12125	6321.3	73.8	79.3

Figure 3: Comparative shares of low-carbon patents (1978–2009)

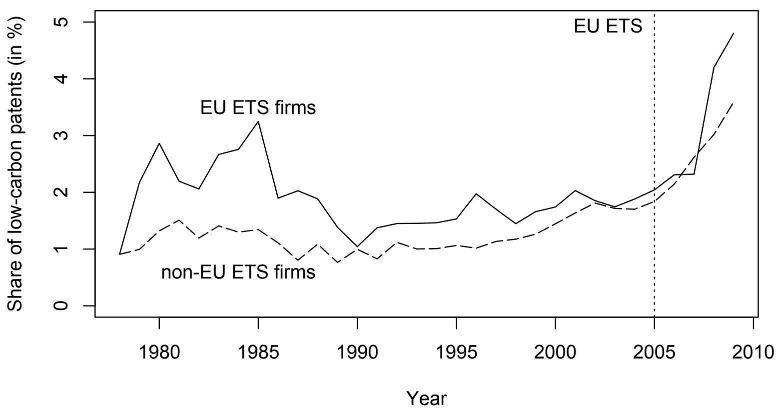

One might be concerned that the surge in patenting activity by EU ETS firms compared to non-EU ETS companies might have been accompanied by a concurrent drop in the relative average quality of inventions filed by EU ETS companies. However, the average number of citations received by low-carbon patents filed by EU ETS companies since 2005 does not significantly differ from those filed by non-EU ETS companies. Similarly, the size of low-carbon patent families is the same for EU ETS and non-EU ETS companies.

The pattern across the two groups is even starker when viewed in terms of year-on-year changes (see figure 4). Up until the mid-1990s there is no clear pattern at all. Then in 1997 the Kyoto Protocol was agreed, which among other things established emissions reduction targets for EU countries. Over the next 5 years the number of low-carbon patents grows for both EU ETS and non-EU ETS firms alike, EU ETS firms being slightly outpaced, but the growth gradually declines until 2003. Then in 2003 the European Parliament adopted the EU ETS directive, and low-carbon patent growth began to pick up again, slightly earlier for EU ETS companies than non-EU ETS companies. After the 2005 launch of the EU ETS, the price of a permit initially varied between €10 and €30. It soon became clear, however, that the regulations were not as stringent as had been anticipated. The number of low-carbon patents filed by EU ETS firms quickly leveled off. In late 2006 it became evident that way too many permits had been issued for Phase 1, and the market price of emission permits collapsed, seemingly taking the growth of low-carbon patents in 2007 with it. It was not until 2008, when Phase 2 of the EU ETS began, that we can once again see an increase in the growth of low-carbon patenting, and a clear divergence between EU ETS and non-EU ETS firms.

Figure 4: Change in number of low-carbon patents (1980–2009)

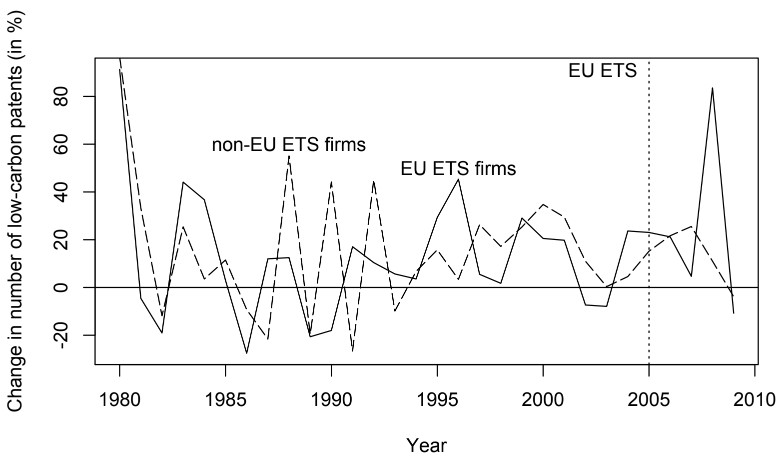

4 Estimating the impact of the EU ETS on low-carbon innovation

4.1 Naive estimates of the impact of the EU ETS

Since 2005 EU ETS firms have filed 2'232 climate related patents, compared to 1'018 patents in the 5 preceding years (an increase of 119%). Non-EU ETS firms filed 4'666 and 2'539 patents protecting low-carbon technologies in the corresponding periods (an increase of 84%). Low-carbon patenting grew at similar rates among EU ETS and non-EU ETS firms in the pre-EU ETS period. If we then were to assume that the number of low-carbon patents filed by EU ETS firms would have grown at the same rate experienced by non-EU ETS firms, had they not been regulated, we can calculate a naive estimate of how many low-carbon patents the EU ETS has added so far: 2'232 - 1.84 $\times$ 1'018 = 361.2, or 16% of their 2'232 low-carbon patents. Put another way, the EU ETS would account for 29.7% of the 5-year increase in low-carbon patenting by EU ETS covered firms.

This is clearly a very naive estimate. The underlying assumption is that EU ETS firms and non-EU ETS firms are comparable in all relevant respects apart from regulatory status. More precisely, we are assuming that the patenting activities of unregulated firms provide an accurate counterfactual estimate of how EU ETS companies would have behaved had they not become regulated. This assumption may be problematic in case non-EU ETS firms are also responding to the new regulations. A more pressing concern, though, is that the two groups of firms appear to be very different even before the EU ETS. Just looking at the patenting activities of these two groups reveals that while only 1 in roughly 1'700 firms is EU ETS regulated they account for nearly 1 in 3 low-carbon patents in the 5 years before the EU ETS launched. Clearly, EU ETS companies do not appear to be representative of the population of firms as a whole.

One simple way to make the two groups more comparable with respect to patenting activities is to restrict our view to the most prolific patenters. Excluding the 10 most prolific, as they are clear outliers, we then look at the next top 100 low-carbon patent holders. 25 of these 100 firms were directly regulated by the EU ETS, and a further 17 either owned or were owned by EU ETS companies. Putting to one side the 17 companies with unclear regulatory status, EU ETS companies in this subsample filed 678 low-carbon patents over the period 2005-2009, and 408 in the 5 preceding years (an increase of 66%). Non-EU ETS firms filed 1'347 and 897 low-carbon patents in the corresponding periods (an increase of 50%). This very crude selection rule appears to produce a set of much more comparable firms, with 1 in 3 firms in this sample being EU ETS regulated and the EU ETS companies accounting for nearly 1 in 3 low-carbon patents in 2000–2004. Repeating the same calculation as before for this much more balanced sample, we attribute 678 - 1.5 × 408 = 66 out of their 678 low-carbon patents to the EU ETS, or 10%. Put another way, the EU ETS would account for 24.4% of the 5-year increase in low-carbon patenting by EU ETS covered firms. Non-EU ETS firms that are more similar on pre-2005 characteristics are likely to provide a better counterfactual estimate of what the EU ETS firms would have done had they not been regulated. When we use this improved counterfactual the estimated impact of the EU ETS on low-carbon technological change, it turns out, is smaller.

4.2 Matching for observed characteristics

We face a difficult identification problem. Looking at changes over time is not sufficient, because it is not possible to adequately control for things like oil price fluctuations and changes in macroeconomic conditions. Comparing EU ETS firms with non-EU ETS firms at a given time allows us to better control for these time-variant factors. On the other hand, as we have discovered, the typical EU ETS firm appear to be very different from the typical unregulated firm even before the EU ETS launched in 2005. This comparison may therefore wrongly attribute some low-carbon patents to the EU ETS that are really the result of other systematic differences between EU ETS and non-EU ETS firms.

Selecting a group of firms that are more similar prior to 2005 would make it more difficult to explain away any difference in outcomes by factors other than the EU ETS. Our estimates so far have been constructed with lenient requirements for similarity, and our attempts to create a more balanced sample have been crude. Ideally one would like to match each EU ETS firm with a group of non-EU ETS firms with similar resources available and facing similar demand conditions, regulations, input prices, etc. In this section we perform just such a matching exercise in order to better estimate the impact of the EU ETS on low-carbon technological change. As we restrict ourselves to more closely matched firms there will inevitably be a greater number of EU ETS companies for which no good match can be found. What is lost in sample size, however, is regained in terms of accuracy and robustness.

Along with patent data, our newly constructed data set contains information on the country and economic sector in which firms operate,[11] as well as other firm-level

[11] Economic sectors are defined at the 3-digit level for the NACE Rev. 2 industry classification. A

information such as turnover and employment.

Using this data, we have tried to assign similar but unregulated firms to each of the 5'521 EU ETS firms. Though, this has not always been possible. EU ETS regulations were not haphazardly applied, so one would generally expect two very similar firms to receive the same regulatory treatment. However, in some cases two apparently similar firms could come to meet different regulatory fates. It is possible that the EU ETS firm operates one installation just large enough to be covered by EU ETS regulations, while the matched control operates an installation just below the threshold. More generally, it may be that the matched control operates several smaller installations, none of which exceed the threshold size individually, but collectively account for similar levels of emissions as the regulated firm's installations. The EU ETS regulations are applied at the installation-level, while our analysis is conducted at the level of the firm. This difference can explain why one firm is regulated and the other is not.

These examples illustrate how unmeasured differences can result in different regulatory treatments for two apparently similar firms. Because we do not have regulatory information on unregulated installations, however, we are unable to give specific examples.[12] Yet, one would not expect these unmeasured differences to be systematically related to the probability that a firm files low-carbon patents.[13] Rather, these differences introduce some quasi-randomness in the assignment of firms to the EU ETS.

As expected, we have unfortunately had to exclude a large number of EU ETS firms at this point because we cannot find suitable unregulated matches, notably for many of the largest firms and some of the most prolific patenters. As implied by the examples above, a very large company, for instance, is very likely to operate at least one EU ETS regulated installation, which would make it impossible to find suitable comparators. In many other cases we do not have enough data for the EU ETS firm to be confident that it is similar to any particular unregulated firm. For these reasons our final sample consists of 743 EU ETS firms and 1'019 non-EU ETS firms. Combined, they have filed 25'104 patents since 2000, of which nearly 2% protected low-carbon technologies.

For each of the 743 EU ETS firms we have found at least one unregulated firm that operates in the same country and economic sector. This means that they are likely exposed to much the same business and regulatory environment, input prices, country

few examples of these sector definitions will illustrate how narrowly sectors are defined: "electric power generation, transmission, and distribution", "steam and air conditioning supply", "manufacture of glass and glass products", "manufacture of plastic products", "manufacture of rubber products".

[12] This also precludes the direct application of a regression discontinuity design to identify the impact of the EU ETS.

[13] Below we investigate whether an omitted variable could reasonably create a spurious result by being correlated with both regulatory status and low-carbon patenting.

and sector specific shocks and trends. The firms are also matched to have similar pre-2005 turnover, patenting records, and age, since their available resources and capacity for R&D and patenting are likely important determinants of a firm's reaction to the EU ETS. See appendix B or technical details about how the matching was implemented.

Figure 5: Similarity of matched EU ETS and non-EU ETS firms

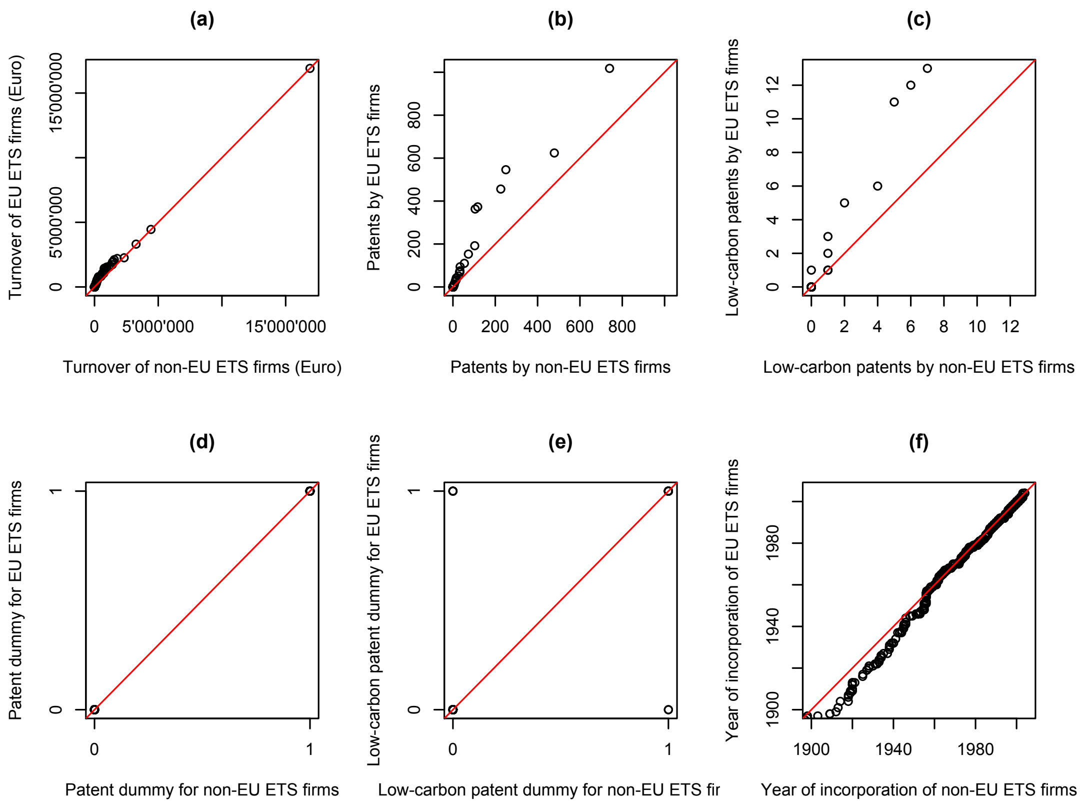

Panel (a) displays the empirical quantile-quantile (e-QQ) plot for average turnover in the 3 years before the EU ETS (2002–2004). Each dot gives the value for one EU ETS firm and the weighted average for a group of non-EU ETS firms matched to a single EU ETS firms. 2002 is the first year for which turnover is recorded in our data set for any firm. Panels (b) and (c) show the e-QQ plots for the total number of patents and the number low-carbon patents filed 2000–2004, respectively. Because the difference between zero and one patent is different from other unit increases, the firms have also been matched for two dummy variables: whether (1) or not (0) the firm filed any patents or low-carbon patents before 2005. Panels (d) and (e) plots the results. Finally, turnover and cumulative patent filings mean different things for an old and new firms, so we have also matched on the year of incorporation interacted with these other variables. Panel (f) displays the the e-QQ plot for year of incorporation on its own. In addition, all pairs are matched exactly for country of operation and for economic sector (defined at the 3-digit level for NACE Rev. 2).

Figure 5 compares the empirical distributions of EU ETS and non-EU ETS firms in our sample on the variables used to construct the match. EU ETS regulated firms have slightly greater pre-EU ETS turnover on average, and filed slightly more patents. However, as can be seen in table 2, we strongly reject the hypotheses that the empirical distributions differ between the EU ETS and non-EU ETS firms.

Table 2: Similarity of regulated and unregulated firms

	Mean EU ETS firms	Mean non-EU ETS firms	Equivalence range	P-value	Critical equivalence range (5% sign. lev.)
Turnover (in Euro)	195'263.00	121'036.40	± 140'498.50	<0.001	± 18'470.00
Patents	5.98	3.28	± 9.30	<0.001	± 0.00
Low-carbon patents	0.08	0.04	± 0.13	<0.001	± 0.00
Year of incorporation	1979.17	1980.69	± 5.51	<0.001	± 1.50
Number of employees	943.32	612.06	± 716.51	<0.001	± 149.00

The first two columns from the left report the mean values for the key matching variables for the EU ETS firms and weighted non-EU ETS firms in our sample. The third column reports a range for within which two firms are judged to have 'equivalent' values for a given variable, following the convention of letting this range be ± 0.2 standard deviations of the pooled sample (Cochran and Rubin, 1973; Ho and Imai, 2006). The fourth column reports the P-values from testing the hypothesis that the two distributions differ by more than the equivalence range. We can reject this hypothesis in favour of equivalence for all variables. The tests were implemented using Wilcoxon's rank sign test, which is sensitive to differences along the whole distribution, not just at the mean. The final column gives the equivalence range for which we would just reject the hypothesis of equivalence in favour of difference at the 5% significance level. Where the range is ±0, we were not able to reject the null hypothesis of any non-negligible difference at the 5% level. The numbers in this table have been rounded.

Because firms look similar within each match, the firms' pre-2005 observable characteristics do not help us predict (better than chance) which firm in each matched group would become regulated after 2005 and which firm in each group would file more low-carbon patents. Conditional on pre-EU ETS observable characteristics, the assignment of firms to the EU ETS appears random. In a naive sense, we have recovered the identifying conditions present in a randomized experiment.

Of course, in a randomized experiment one can rely on the law of large numbers to achieve similarity between a treated and control group on both observed and unobserved characteristics. Matching, on the other hand, achieves an observed similarity by construction, so similarity on matched characteristics cannot be taken to imply that the treated and control firms are also similar on unobserved characteristics. A simple test of whether matching has achieved balance on unobserved variables is to look at a variable that was not used to construct the matches. We have one such variable in our data set: the number of employees of a firm. As figure 6 and the final row of table 2 show, the empirical distributions of number of employees of the EU ETS firms and the unregulated firms are very similar, and we can reject the hypothesis that they are materially different. We can therefore have some confidence that matching has indeed recovered the central identifying condition of a randomized experiment. From a statistical perspective, therefore, the behavior of the control firms reflect how one would have expected the EU

ETS regulated firms to behave had they not become regulated.

Figure 6: Similarity of matched EU ETS and non-EU ETS firms on 'unobserved' variable

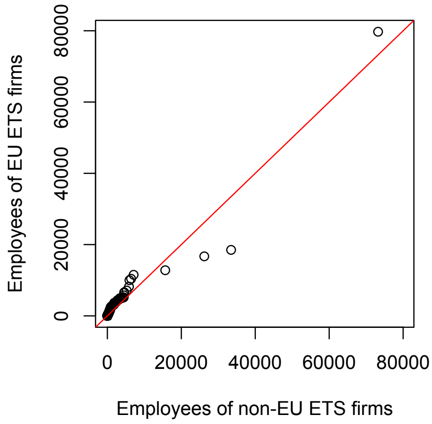

Being reasonably reassured that our regulated and unregulated firms are comparable, we next turn to estimating the impact of the EU ETS on low-carbon technological change.

4.3 Matching estimates

For each firm we measure the change in the number of low-carbon patents from 2000–2004 to 2005–2009. The outcomes of the EU ETS and non-EU ETS firms in each group are then compared. The mean difference-in-differences is 0.04, which rises to 1.5 for the subset of groups in which at least one member had filed at least one low-carbon patent before 2005. The Hodges-Lehmann point estimate for the average treatment effect on the treated is 1.27×10^{-5}. We can reject both the hypothesis that the EU ETS has increased average low-carbon patenting by 0.001 patents or more at the 5% level of significance, and the hypothesis that the EU ETS has reduced average low-carbon patenting by the same amount. We are 95% confident, therefore, that the impact of the EU ETS on the average change in low-carbon patent filings of EU ETS firms lies in the interval $[-0.001, 0.001]$.

The EU ETS appears to have had virtually no impact at all on low-carbon technological change, contrary to what our naive estimates suggested. The closer we get to a sample of EU ETS and non-EU ETS firms that were comparable before 2005, the smaller the estimated impact of the EU ETS appears to be. The naive estimates calculated earlier, therefore, are more likely explained by systematic differences between EU ETS and non-EU ETS companies. The matching results are in line the expectations expressed by Schleich and Betz (2005); Gagelmann and Frondel (2005); Grubb et al. (2005) and others, who early on said that the EU ETS would not encourage low-carbon innovation due to overallocation of emissions permits.

But is our finding best explained by the EU ETS having no impact? Before we can say this it is necessary to examine possible competing explanations.

Competing explanation 1: Artifact of outcome definition It is possible that our finding is an artifact of our particular measure of low-carbon technological change. However, re-matching for an expanded definition of "low-carbon technologies", for citation- and family size-weighted patent measures, and for a measure that uses patents filed with national patent offices instead of the EPO, produces the same result (see table 3). It appears, therefore, that our finding is robust to how the outcome is defined.

Table 3: Difference-in-differences estimates with different outcome definitions

	HL Point estimate	95% C.I.
Using extended "low-carbon" definition	4.30×10^{-5}	$\pm$ 0.001
Using citation-weighted patents	1.99×10^{-5}	$\pm$ 0.001
Using family size-weighted patents	-2.77×10^{-5}	$\pm$ 0.001
Using patents filed with national patent offices	7.76×10^{-6}	$\pm$ 0.001

Competing explanation 2: Biased study design The study design may be biased in some unknown way, which might coincidentally lead us to underestimate the impact of the EU ETS. One way to investigate such hidden biases is to look at a secondary outcome for which you have a stronger expectation. If this secondary outcome behaves as expected, this provides reassurance that study design itself does not harbor any hidden bias.

In the present context, we would expect the EU ETS to have no impact on patents filed to protect non-greenhouse gas "pollution control" technologies, as defined by Popp (2006b). If there is a hidden built-in bias in our study design, we would expect to observe either a positive or negative impact of the EU ETS on patenting for these technologies. The Hodges-Lehmann point estimate of the difference-in-differences in "pollution control" patents is 1.27×10^{-5}, and we are 95% confident that the impact of the EU ETS lies in the region $[-0.001, 0.001]$.[14] This suggests our finding is not due to a hidden bias in the study design.

Competing explanation 3: Non-EU ETS firms also responding to EU ETS Firms not regulated by the EU ETS may nevertheless respond to it, either directly or indirectly because they engage in competition with EU ETS firms. In this case, the EU ETS firm might behave the same as the matched unregulated firms, not because neither is innovating, but because they are both innovating. To examine this possibility, we have re-matched our EU ETS companies to firms operating in countries that did not participate in the 2005 launch of the EU ETS (Norway, Switzerland, Romania, and Bulgaria). This comparison is less likely to suffer from this kind of bias, because the firms are less exposed to the market created by the EU ETS and less directly engaged in competition with EU ETS companies.[15]

Using this set of matches, the estimated average treatment effect is -3.73×10^{-5}, and we can still be 95% confident that the effect lies within $[-0.001, 0.001]$.

Competing explanation 4: No impact in subsample only A criticism of the external validity of the result is that it is estimated for an unrepresentative sample of EU ETS firms. The discrepancy between the naive estimates and the matching estimates might suggest to some that the EU ETS is indeed encouraging low-carbon technological change, but that the matching estimate fails to detect it because it looks at a smaller sample of EU ETS firms.

In many cases we were forced to exclude EU ETS firms because no suitable matches could be found. In a large number of cases, though, missing data, in particular turnover figures from before 2005, prevents us from adequately assessing similarity. Turnover figures become much more widely available starting in 2005. One way to address the external validity concerns, therefore, is to allow ourselves to use 2005 turnover figures to construct the matches. This is not generally desirable, because the EU ETS might have affected 2005 turnover, which in turn had some effect on low-carbon patenting. If this is the case, the matching estimate using 2005 turnover would be biased because it omits this channel. However, because using 2005 turnover gives us access to a far greater number of EU ETS and non-EU ETS firms, it may provide a reasonable test of whether our findings apply to the EU ETS more broadly.

[14] Roughly 20% of EPO patents classified as one of Popp's pollution control technologies also fall into the low-carbon category. Excluding these, however, does not substantively affect the outcome.

[15] While this comparison helps address a potential bias introduced by non-EU ETS firms responding to the EU ETS, it is not able to control for between-country differences.

Matching using 2005 turnover figures produces 3'177 matched pairs. The point estimate for this sample is -8.51×10^{-5}, and we can be 95% confident the effect lies within $[-0.001, 0.001]$. The typical matched firm still looks much the same, though, which is what one would expect if we were simply finding more firms around the same EU ETS thresholds. This illustrates that the 743 EU ETS firms in our original matched sample appear to be representative of a much larger part of the EU ETS. On the other hand, it also means that this re-match does not so much help address concerns that the EU ETS is affecting low-carbon patenting among the types of companies for which suitable unregulated matches still could not be found.

Competing explanation 5: Omitted variable bias Our matching estimate relies on the assumption that firms that appear similar are similar in unmeasured dimensions as well—often called "selection on observables". Conditional on the observed pre-treatment characteristics, we assume the odds of receiving treatment is the same for the EU ETS firm and the non-EU ETS match. We have some justification for making this assumption, since the matching appears to have done a decent job of recovering the identifying condition of a randomized experiment. However, unobserved systematic differences between regulated and unregulated firms might still bias our findings. It is therefore not entirely satisfactory to let our results stand without examining how sensitive they are to violations of the 'even odds' assumption. What kind of an omitted variable could in principle skew the odds so far as to undermine confidence in our estimate?

In order to argue that we have underestimated the effect of the EU ETS, one would have to postulate an omitted variable that at once *increases* the prior odds of the EU ETS firm having become regulated and *reduces* the odds of the EU ETS firm filing more low-carbon patents, or vice versa. The prime candidate for an omitted variable—firm-level emissions—is generally thought to be positively correlated with both the probability of becoming regulated and of filing more low-carbon patents, so it would not explain why we find that the EU ETS has had no impact. If anything, this omission would imply we have overestimated the impact of the EU ETS.

The omission of complementary regulations, however, could result in an underestimate of the impact of the EU ETS. If before 2005 countries began implementing emissions regulations specifically targeting those firms exempted from the EU ETS, our estimate would be more accurately interpreted as the difference between the impact of the EU ETS and these complementary policies. To think more systematically about the potential problem caused by omitting such a variable, the model for sensitivity analysis developed by Rosenbaum (1987) and Rosenbaum and Silber (2009) allows us to inspect just how much a hypothetical omitted variable would have to skew the prior odds of regulation and low-carbon innovation in order to make our estimate just insignificant, so that we no longer can say with 95% confidence it lies within $[-0.001, 0.001]$. To make the matching estimate just insignificant at the 5% level, one would need to postulate an omitted variable that, if measured before 2005, would have correctly predicted more than 99 out of 100 times which firm would become EU ETS regulated and which firm would file fewer low-carbon patents.[16] The omitted variable criticism, therefore, is only valid if one can propose a near-deterministic omitted variable, and it appears unlikely that complementary policies have been implemented in such a systematic fashion across the EU.

Competing explanation 6: Innovation only further up the technology supply chain Both the EU ETS firms and non-EU ETS firms we study are technology-users, but they are not necessarily technology-suppliers. To the extent that third-party technology-supplier are an important source of new technologies, our estimate may underestimate the impact of the EU ETS.

Two points are worthy of attention. First, economic theory predicts that environmental regulations would produce greater incentives to develop new technologies for directly regulated firm than for third-party technology-suppliers (Milliman and Prince, 1989; Fischer et al., 2003). The asymmetry arises because the latter are not discharging emissions themselves and receive no private benefit from the new technology. To the extent that the EU ETS is encouraging low-carbon technological change, therefore, economic theory predicts this response to be strongest among regulated firms.

Second, empirically we see that the firms in our sample are in fact technology-suppliers. As mentioned, the firms in our sample have filed over 25'000 patents since 2000, circa 2% of which to protect low-carbon technologies. These are firms with above average innovation capabilities. This competing explanation would imply that these firms with well-developed low-carbon innovation capabilities are responding solely by purchasing technologies off-the-shelf from third-party technology-suppliers in other sectors, and not at all by innovating themselves. Though we cannot rule out this possibility it appears somewhat doubtful, especially in light of the aforementioned economic theory.

Ultimately, the claim that the EU ETS is having an impact *only* further up the technology supply chain needs to be met with the same level of skepticism as any other

[16] The sensitivity parameter at which the result becomes insignificant is $\Gamma = 52$, in Rosenbaum's notation. This can be decomposed into the biases present in treatment assignment and outcomes using propositions in Rosenbaum and Silber (2009).

empirical hypothesis. To offer it as a credible competing explanation, one would need to implement an identification strategy that could persuasively attribute to the EU ETS the patents filed by these third-party technology-suppliers. In the absence of such evidence, one cannot conclusively rule in or out the 'technology-supplier'-hypothesis, but we have both theoretical and empirical reasons to think it unlikely.

None of these competing explanations seem to offer a compelling alternative to the finding that the EU ETS has had no impact on low-carbon technological change, though it is of course impossible to conclusively rule them in or out. One must be careful also because many of the tests we have used to investigate these alternative explanations, though addressing one potential source of bias, may introduce new biases of their own (e.g. using 2005 turnover figures, or matching to firms in other countries). The point here, however, is that to replicate our results each time, the new bias would have to be of the same sign and magnitude as the hypothesized bias in the original match. This explanation becomes increasingly unlikely with each new test, and the explanation that our estimate is unbiased appears more likely by comparison.

Finally, it is worth repeating that our matching sample was selected to allow estimation of a reliable counterfactual. Though the EU ETS firms in this sample do appear to represent a sizable portion of EU ETS firms, they are not necessarily representative of all EU ETS companies. It may be, therefore, that the naive estimates from earlier are detecting some EU ETS-induced low-carbon patents filed by regulated companies for which we simply cannot find appropriate matches.

5 Conclusion

Economic theory suggests that imposing a price on emissions would encourage firms to develop new technologies to reduce their emissions more cheaply. The EU ETS was launched in 2005 and there is now patent data for the 5 subsequent years. Previous studies have found that most of the induced innovation response is observed in the first 5 years following the introduction of a new environmental policy (Popp, 2002). This paper has used a newly constructed data set covering 8.5 million European firms, of which over 5'500 are regulated under the Scheme, to investigate whether the EU ETS has induced low-carbon technological change.

The EU ETS offers a unique opportunity to investigate the impact of environmental policy on technological change. The EU ETS is the first and largest policy initiative of its kind anywhere in the world, which makes it an interesting large-scale case study.

But more importantly, EU ETS rules are designed to cover only large installations, so firms operating smaller installations are not covered by the new regulations. Because EU ETS rules are applied at the level of the installation rather than the firm, we are able to compare regulated companies with otherwise similar but unregulated companies.

Exploratory data analysis reveals a rapid increase in low-carbon patenting activities since 2005. The rise appears disproportionately large among EU ETS regulated companies, especially during the more stringent second trading phase that started in 2008. Naive estimates obtained by comparing EU ETS and non-EU ETS firms suggest that the Scheme may be responsible for up to 30% of the increase in low-carbon patenting of regulated companies. However, it is difficult to invoke the EU ETS as a causal explanation here because of many possible confounding influences, in particular potentially important systematic differences between EU ETS and non-EU ETS firms.

In order to estimate the causal impact of the EU ETS on low-carbon technological change we compute the difference-in-differences estimates for a matched sample of firms. We find evidence that the EU ETS has not impacted the direction of technological change, and while it is impossible to conclusively rule out alternative explanations for our results, our investigations of these alternatives leads us to conclude our findings are robust. Furthermore, even if the EU ETS has impacted only large companies for which suitable unregulated comparators cannot be found, our findings imply that the EU ETS so far has had at best a very limited impact on low-carbon technological change.

Our results have important policy implications. The EU ETS forms an integral part of the European Union's roadmap to a low-carbon economy in 2050. Moreover, policy makers in the process of implementing new carbon market programs in New Zealand, the North-Eastern United States, Australia, and elsewhere, can learn from the EU ETS experience. It appears that emissions reductions in past emissions trading programs like the US Acid Rain Program have come largely from operational rather than technological changes. The same appears to have happened with the EU ETS. Emissions reductions in the EU have so far come largely from measures like fuel switching, but such abatement strategies will not be sufficient to reach the EU's ambitious longer term targets. New low-carbon technologies are needed. Our findings suggest, however, that the EU ETS in its current form might not be enough to incentivize low-carbon technological change.

Even before the EU ETS launched, many argued it would not impact firm innovation behaviour because of an overly generous allocation of emissions permits, and that permits were awarded to polluters free of charge (Schleich and Betz, 2005; Gagelmann and Frondel, 2005; Grubb et al., 2005). We have not attempted to test these explanations, nor would it be feasible to do so given the lack of variation in the EU ETS rules so far.

Future changes to the rules of the EU ETS may provide opportunities to study these specific questions. To the extent that these factors account for our findings, however, there are relatively clear policy implications—tighten the emissions cap and/or sell permits instead of giving them away free. The current move to set aside permits, as well as the increased reliance on auctions to distribute permits in the third trading phase, would in these cases appear to be moves in the right direction.

There is a further interpretation of our results. The EU ETS has been a tremendously ambitious policy intended in part to create a demand for new low-carbon technologies, but demand-pull policies like the EU ETS only address one part of the problem. Both pollution and innovation are characterized by market failures, which result in over-pollution and under-innovation—low-carbon innovation is doubly under-provided (Jaffe et al., 2005). Demand-pull policies like the EU ETS may therefore fail to bring about low-carbon technological change unless combined with complementary technology-push policies (Fischer and Newell, 2008; Acemoglu et al., 2012). Our findings are consistent with the conclusion that a price on carbon emissions alone is not enough to have a substantial impact on low-carbon technological change.

In spite of our findings, it is still possible that EU ETS firms have begun devoting more resources to low-carbon R&D, but that this is yet to translate into new patents. We cannot address this at present, but the analysis can be updated as more patent data becomes available.

Finally, it is important to emphasize that the matching estimates are calculated using a subsample of EU ETS firms for which suitable unregulated matches could be found. It is conceivable, therefore, that the EU ETS has spurred low-carbon innovation elsewhere in the economy. In particular, it is possible that the EU ETS has induced unregulated companies up and down the supply chain (i.e. third-party technology-suppliers) to develop new low-carbon technologies. However, it is a difficult task to confidently attribute those patents to the EU ETS, or rule out the EU ETS as an important source of encouragement. Our analysis is conservative in that we take great care not to misattribute low-carbon patents to the EU ETS. In order to convincingly attribute low-carbon patents elsewhere in the supply chain to the EU ETS, future research efforts will require much richer data on technology supply relationships, e.g. licensing data. We leave this as a future project.

References

Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012). The Environment and Directed Technical Change. *The American Economic Review*, 102(1):131–166.

Ambec, S., Cohen, M., Elgie, S., and Lanoie, P. (2010). The porter hypothesis at 20: Can environmental regulation enhance innovation and competitiveness? *CIRANO Working Papers*.

Anderson, B., Convery, F., and Maria, C. D. (2011). Technological change and the EU ETS: the case of Ireland. *IEFE Working Paper Series*, (43).

Anderson, B. and Di Maria, C. (2011). Abatement and Allocation in the Pilot Phase of the EU ETS. *Environmental and Resource Economics*, pages 1–21.

Arimura, T. H., Hibiki, A., and Johnstone, N. (2007). An empirical study of environmental R&D: what encourages facilities to be environmentally innovative? In Johnstone, N., editor, *Corporate Behaviour and environmental Policy*, chapter 4. Cheltenham, UK: Edward Elgar in association with OECD.

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.

Burtraw, D. (2000). *Innovation Under the Tradeable Sulphur Dioxide Emission Permits Programme in the US Electricity Sector*, chapter 4, pages 63–84.

Burtraw, D. and Szambelan, S. (2009). US emissions trading markets for SO2 and NOx. *Resources for the Future Discussion Paper*, pages 09–40.

Cochran, W. and Rubin, D. (1973). Controlling bias in observational studies: A review. *Sankhyā: The Indian Journal of StatisticsSeries A*, 35(4):417–446. Dedicated to the Memory of P. C. Mahalanobis.

Dechezleprêtre, A., Glachant, M., Haščič, I., Johnstone, N., and Ménière, Y. (2011). Invention and transfer of climate change–mitigation technologies: a global analysis. *Review of Environmental Economics and Policy*, 5(1):109.

Delarue, E., Ellerman, D., 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.

Ellerman, A. and Buchner, B. (2008). Over-allocation or abatement? A preliminary analysis of the EU ETS based on the 2005–06 emissions data. *Environmental and Resource Economics*, 41(2):267–287.

European Commission (2011). A Roadmap for moving to a competitive low carbon economy in 2050. Technical Report COM(2011) 112, European Union.

Fischer, C. and Newell, R. G. (2008). Environmental and technology policies for climate mitigation. *Journal of Environmental Economics and Management*, 55(2):142–162.

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.

Gagelmann, F. and Frondel, M. (2005). The impact of emission trading on innovation-science fiction or reality? *European Environment*, 15(4):203–211.

Gerlagh, R. (2008). A climate-change policy induced shift from innovations in carbon-energy production to carbon-energy savings. *Energy Economics*, 30(2):425–448.

Goulder, L. and Schneider, S. (1999). Induced technological change and the attractiveness of CO2 abatement policies. *Resource and Energy Economics*, 21(3-4):211–253.

Grubb, M., Azar, C., and Persson, U. (2005). Allowance allocation in the European emissions trading system: a commentary. *Climate Policy*, 5(1):127–136.

Hall, B., Jaffe, A., and Trajtenberg, M. (2005). Market value and patent citations. *The RAND Journal of Economics*, 36(1):16–38.

Harhoff, D., Narin, F., Scherer, F., and Vopel, K. (1999). Citation frequency and the value of patented inventions. *Review of Economics and statistics*, 81(3):511–515.

Harhoff, D., Scherer, F., and Vopel, K. (2003). Citations, family size, opposition and the value of patent rights. *Research Policy*, 32(8):1343–1363.

Hicks, J. R. (1932). *The Theory of Wages*. MacMillan.

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.

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.

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.

Kneese, A. V. and Schultze, C. (1975). Pollution, Prices, and Public Policy. *(Brookings Institution, Washington, DC)*.

Lange, I. and Bellas, A. (2005). Technological change for sulfur dioxide scrubbers under market-based regulation. *Land Economics*, 81(4):546–556.

Lanjouw, J. and Mody, A. (1996). Innovation and the international diffusion of environmentally responsive technology. *Research Policy*, 25(4):549–571.

Lanoie, P., Laurent-Lucchetti, J., Johnstone, N., and Ambec, S. (2007). Environmental policy, innovation and performance: new insights on the Porter hypothesis. *CIRANO Working Papers*.

Martin, R., Muuls, M., and Wagner, U. (2011). Climate change, investment and carbon markets and prices – evidence from manager interviews. pages 1–51.

Milliman, S. and Prince, R. (1989). Firm incentives to promote technological change in pollution control. *Journal of Environmental Economics and ...*, 17(3):247–265.

OECD (2009). OECD Patent Statistics Manual. Technical report, OECD.

Petsonk, A. and Cozijnsen, J. (2007). Harvesting the low-carbon cornucopia: How the european union emissions trading system (eu-ets) is spurring innovation and scoring results. pages 1–23.

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. (2004). ENTICE: Endogenouse Technological Change in the DICE Model of Global Warming. *Journal of Environmental Economics and Management*, 24(1):742–768.

Popp, D. (2006a). Innovation in climate policy models: Implementing lessons from the economics of R&D. *Energy Economics*, 28(5-6):596–609.

Popp, D. (2006b). International innovation and diffusion of air pollution control technologies: the effects of NOX and SO2 regulation in the US, Japan, and Germany. *Journal of Environmental Economics and Management*, 51(1):46–71.

Popp, D. (2010). Innovation and climate policy. *NBER Working Paper*.

Popp, D., Newell, R., and Jaffe, A. (2009). Energy, the environment, and technological change. *NBER Working Paper*.

Porter, M. E. (1991). Essay: America's green strategy. *Scientific American*, 264(3).

Rosenbaum, P. (1987). Sensitivity analysis for certain permutation inferences in matched observational studies. *Biometrika*, 74(1):13–26.

Rosenbaum, P. and Silber, J. (2009). Amplification of sensitivity analysis in matched observational studies. *Journal of the American Statistical Association*, 104(488):1398–1405.

Schleich, J. and Betz, R. (2005). Incentives for energy efficiency and innovation in the european emission trading system. *Proceedings of the 2005 eceee Summer ...*.

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.

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.

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.

Trajtenberg, M. (1990). A penny for your quotes: patent citations and the value of innovations. *The Rand Journal of Economics*, pages 172–187.

van der Zwaan, B., Gerlagh, R., Schrattenholzer, L., et al. (2002). Endogenous technological change in climate change modelling. *Energy economics*, 24(1):1–19.

van Zeebroeck, N. (2011). The puzzle of patent value indicators. *Economics of Innovation and New Technology*, 20(1):33–62.

A Data

We constructed our data set in the following way. First, we gathered regulatory data from the CITL and national registries. We were then able to identify EU ETS regulated companies in the company database Orbis for 18 countries. Having identified the EU ETS regulated companies, we then identified the 3-digit NACE sectors associated with the EU ETS companies, and downloaded data from Orbis on all companies operating in the same countries and the same sectors. We then separately matched PATSTAT with Orbis, resulting in a data set that combines regulatory status, basic firm characteristics, and patenting activities.

B Matching

The matches were constructed using `GenMatch()` from the `R`-package `Matching`. It uses a genetic search algorithm to search the propensity score space for a specification that minimizes imbalances on the whole set of covariates (see Sekhon, 2007, for details). We used variable ratio matching with replacement, so that each EU ETS firm could be matched to one or more non-EU ETS firms depending on how many similar non-EU ETS firms could be found.

C Patents

We use the patent codes available at www.oecd.org/environment/innovation. For our main measure of low-carbon patents we use the EPO patent classes for low-carbon patents definition. These include the following classes:

B. ENERGY GENERATION FROM RENEWABLE AND NON-FOSSIL SOURCES

B.1. RENEWABLE ENERGY GENERATION

B.1.1. Wind energy: Y02E10/7

B.1.2. Solar thermal energy: Y02E10/4

B.1.3. Solar photovoltaic (PV) energy: Y02E10/5

B.1.4. Solar thermal-PV hybrids: Y02E10/6

B.1.5. Geothermal energy: Y02E10/1

B.1.6. Marine and hydro energy: Y02E10/3

B.2. ENERGY GENERATION FROM FUELS OF NON-FOSSIL ORIGIN

B.2.1. Biofuels: Y02E50/1

B.2.2. Fuel from waste: Y02E50/3

C. COMBUSTION TECHNOLOGIES WITH MITIGATION POTENTIAL (e.g. using fossil fuels, biomass, waste, etc.)

C.1. TECHNOLOGIES FOR IMPROVED OUTPUT EFFICIENCY (Combined combustion): Y02E20/1

C.2. TECHNOLOGIES FOR IMPROVED INPUT EFFICIENCY (Efficient combustion or heat usage): Y02E20/3

D. TECHNOLOGIES SPECIFIC TO CLIMATE CHANGE MITIGATION

D.1. CAPTURE, STORAGE, SEQUESTRATION OR DISPOSAL OF GREENHOUSE GASES

D.1.1. CO2 capture or storage (CCS): Y02C10

D.1.2. Capture or disposal of greenhouse gases other than CO2: Y02C20

E. TECHNOLOGIES WITH POTENTIAL OR INDIRECT CONTRIBUTION TO EMISSIONS MITIGATION

E.1. ENERGY STORAGE: Y02E60/1

E.2. HYDROGEN TECHNOLOGY: Y02E60/3

E.3. FUEL CELLS: Y02E60/5

We use additional patent classes for "extended" low-carbon patents definition:

Energy-efficient cement (see Dechezleprêtre et al., 2011, for list of codes)

Natural pozzuolana cements: C04B 7/1213

Cements containing slag: C04B 7/1421

Iron ore cements: C04B 7/22

Cements from oil shales, residues or waste other than slag: C04B 7/24-30

Calcium sulfate cements: C04B 11/00

HEATING (incl. water and space heating; air-conditioning)

Hot-water central heating systems - in combination with systems for domestic hot-water supply: F24D3/08

Hot-water central heating systems - using heat pumps: F24D3/18

Hot-air central heating systems - using heat pumps: F24D5/12

Central heating systems using heat accumulated in storage masses - using heat pumps: F24D11/02

Other domestic- or space-heating systems - using heat pumps: F24D15/04

Domestic hot-water supply systems - using heat pumps: F24D17/02

Use of energy recovery systems in air conditioning, ventilation or screening: F24F12

Combined heating and refrigeration systems, e.g. operating alternately or simultaneously: F25B29

Heat pumps: F25B30