

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

Raphael Calel^{1,2} and Antoine Dechezleprêtre^{2,3}

¹Department of Agricultural and Resource Economics, UC Berkeley

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

³Centre for Economic Performance, London School of Economics

First version: March 2012

This version: November 2013

Abstract

This paper investigates the impact of the European Union Emissions Trading System (EU ETS) on technological change, exploiting installations-level inclusion criteria to estimate the System's causal impact on firms' patenting. We find that the EU ETS has increased low-carbon innovation among regulated firms by as much as 10%, while not crowding out patenting for other technologies. We also find evidence that the EU ETS has not impacted patenting beyond the set of regulated companies. These results imply that the EU ETS accounts for nearly a 1% increase in European low-carbon patenting compared to a counterfactual scenario.

JEL: O3, Q55, Q58, C14

Keywords: Directed technological change, EU Emissions Trading System, Policy evaluation.

*E-mail: r.calel@berkeley.edu and a.dechezlepretre@lse.ac.uk. For their insightful comments and generous advice, we owe a great thanks to Devin Caughey, Sam Fankhauser, Matthieu Glachant, Bronwyn Hall, Ivan Hascic, Nick Johnstone, Carmen Marchiori, David Popp, and Paul Rosenbaum. Participants of seminars and conferences in Asheville, Cambridge, London, Madrid, Mannheim, Milan, Paris, Rome, Toulouse, Toxa and Venice have all improved the paper. 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. Lastly, but by no means least, we wish to thank the Editor and two anonymous reviewers for their constructive suggestions and comments.

1 Introduction

Emissions trading programs have assumed an ever more prominent role in environmental policy over the last few decades. In the US, the Acid Rain Program, the Regional Greenhouse Gas Initiative (RGGI), and California's cap-and-trade program are all examples of this trend. Australia, New Zealand, and the Canadian province of Quebec have all recently created their own cap-and-trade programs to regulate greenhouse gas emissions. China has initiated several pilot programs in anticipation of a national market that will be launched after 2015. Japan, South Korea, Brazil, Mexico, and Chile are individually making moves toward launching their own. Global carbon markets are worth over \$175 billion a year according to recent figures (Kossoy and Guigon, 2012), and cover over 20% of global greenhouse gas emissions (Kossoy et al., 2013). With so many new initiatives in the works, these numbers will likely grow much larger in years to come.

At present, most of the \$175 billion a year is accounted for by the European Union Emissions Trading Scheme (EU ETS), today's largest cap-and-trade program in the world. The EU ETS was launched in 2005, allocating tradable emissions permits to over 12'000 power stations and industrial plants in 24 countries, accounting for over 40% of the EU's total greenhouse gas emissions. Like all of the new emissions trading initiatives around the globe, the EU ETS was expected not only to reduce carbon emissions in a cost-effective manner, but also to spur the development of new low-carbon technologies. 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) and Acemoglu et al. (2012), suggests that part of this new investment will be directed toward developing and commercializing new emissions-reducing technologies. From an economic perspective, it is crucial that carbon market programs provide incentives for technological change, since new technologies may substantially reduce the long-run cost of abatement (Jaffe et al., 2003; Stavins, 2007). From a political perspective, induced innovation may improve the acceptability of these policies. Indeed, EU policy makers have often articulated their vision that the EU ETS would be a driving force of low-carbon innovation and economic growth (see, for instance, European Commission, 2005, and European Commission, 2012).

In this paper we conduct the first comprehensive investigation of the impact of the EU ETS on low-carbon technological change in the first 5 years of the System's existence. The EU ETS offers a unique opportunity to investigate the impact of environmental

policy on technological change. It is the first and largest environmental policy initiative of its kind anywhere in the world, which by itself would make it an interesting case to study. But more important is the fact that, 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, although the firms themselves might be just as large as those affected by the regulations.¹ Because innovation takes place at the level of the firm, we can exploit these installation-level inclusion criteria to compare firms with similar resources available for research and similar patenting histories, but which have fallen under different regulatory regimes since 2005. This provides an opportunity to apply the sort of quasi-experimental techniques most suited to assessing the causal impacts of environmental policies (List et al., 2003; Greenstone and Gayer, 2009). Studies employing these methods have found that environmental regulations inhibit new-plant formation (List et al., 2003), but stimulate capital investment in existing plants (Fowlie, 2010). To the authors' knowledge, though, this is the first time these methods have been employed to study the impact of environmental policy on directed technological change.

We use a newly constructed data set that records patenting activities, key firm characteristics, and regulatory status with respect to the EU ETS. Our data set includes information on over 30 million firms across 23 countries, 18 of which took part in the 2005 launch of the EU ETS. We identify over 5'500 firms operating more than 9'000 installations regulated under the EU ETS, accounting for over 80% of EU ETS-wide emissions. Using this data set, we are able to compare unregulated and would-be regulated firms both before and after the EU ETS launched. The low-carbon patent classification recently developed by the European Patent Office (EPO) allows us to identify emissions reduction technologies. A matched difference-in-differences study design enables us to control for confounding factors that affect both regulated and unregulated firms (input prices, sector- and country-specific policies, etc.), as well as firm-level heterogeneity (Heckman et al., 1998b,a; Smith and Todd, 2005; Abadie, 2005). Our estimates provide the first comprehensive empirical assessment of the impact of the EU ETS on directed technological change.

A casual look at aggregate patent data reveals a surge in low-carbon patenting since 2005. The increase appears larger among EU ETS regulated companies, and our matched difference-in-differences estimate of the treatment effect implies that the EU ETS is responsible for a 36.2% increase in low-carbon patenting among our matched sample

¹Although the EU ETS regulations are applied at the level of the installation, we will often use 'EU ETS firms' or 'regulated firms' as shorthand for firms operating at least one EU ETS regulated installation.

of 3'428 EU ETS firms, or an increase of 8.1% across all of the 5'500 EU ETS firms. Because these firms only account for a small portion of all patents, however, this would account for less than a 1% increase of low-carbon patenting at the EPO. Put another way, only 2% of the post-2005 surge in low-carbon patenting can be attributed to the EU ETS.

With respect to concerns that low-carbon innovation would crowd out development of other technologies (Popp and Newell, 2012), we find evidence that the EU ETS has in fact encouraged patenting for other technologies, but by a very small amount. We investigate several challenges to the internal and external validity of our results—e.g. omitted variable bias and a failure of ‘selection on observables’—but our conclusions appear to be robust.

For fear that a focus on EU ETS firms would have blinkered us to a broader indirect impact of the EU ETS, we identify 12'000 of likely third-party technology providers and purchasers and test whether these firms have also responded to the EU ETS. The estimates are only indicative, but we find no compelling evidence that the EU ETS has had either a net positive or net negative impact on the patenting activities of third parties. Taken together, our findings suggest that, while EU ETS regulated firms have responded strongly, the System so far has had at best a very limited impact on the overall pace and direction of 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), but in its current form the EU ETS may not be providing incentives for low-carbon technological change on a large scale.

Technological change may be the single most important determinant of the long-run cost of emissions abatement. Consequently, the ability of an environmental policy to influence technological change is perhaps one of the most important criteria on which to judge its success (Kneese and Schultze, 1975; Pizer and Popp, 2008). In light of this, it is not surprising that there are ongoing efforts from both theoretical and empirical economists to better understand the capacity of environmental policies to induce clean innovation. On the theoretical side, the past few decades have seen the emergence of a considerable literature further developing the induced innovation hypothesis, especially in the context of climate change mitigation (Goulder and Schneider, 1999; van der Zwaan et al., 2002; Popp, 2004; Gerlagh, 2008; Acemoglu et al., 2012).

On the empirical side, a large and growing research enterprise is trying to understand and quantify the link between environmental policies and directed technological change, often with innovation measured at the level of economic sectors or countries (Jaffe and Palmer, 1997; Newell et al., 1999; Brunnermeier and Cohen, 2003; Popp, 2002; Aghion

et al., 2012, and many others. See Popp et al., 2009, Popp, 2010, and Ambec et al., 2010, for recent surveys). Our study contributes to this literature, and analyzes the policy impacts at the firm-level. The handful of studies that have begun to investigate the innovation impact of the EU ETS rely on interview-based methodologies and most analyze small unrepresentative samples (Hoffmann, 2007; Tomás et al., 2010; Anderson et al., 2011). Martin et al. (2011) take extra precautions to ensure consistency across interviews with different firms, and they conduct the largest study to date covering 450 EU ETS firms in 6 countries. We use patent portfolios as an objective proxy of technological change, and our study considers over 5'500 EU ETS firms in 18 countries, accounting for roughly 80% of the program as a whole. With this, we provide the first comprehensive empirical estimates of the System's impact on directed technological change.

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 inform us about how the EU ETS is likely to have impacted technological change. In section 3 we familiarize ourselves with our newly constructed data set, and use it 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 on regulated firms, and in section 5 we examine its indirect impact on third-party technology providers and purchasers. Section 6 summarizes and discusses the evidence in light of the broader empirical literature. We conclude by considering some of the potential policy implications of our findings, and directions for future research.

2 Emissions trading and directed technological change

2.1 Empirical background

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, 2006; Arimura et al., 2007; Lanoie et al., 2007; Johnstone et al., 2010). But while there appears to be a general link between environmental policy and directed technological change, a more careful reading of the literature yields two cautionary observations that seem particularly relevant for the EU ETS.

Firstly, the impact of emissions trading programs specifically, rather than environmental policies more broadly construed, appear to be more modest. Most studies con-

sider the Acid Rain Program, which in 1995 replaced the traditional regulatory regime for sulphur dioxide emissions from US power plants. Although patenting for sulphur dioxide control technologies began a precipitous decline after 1995 (Taylor, 2012), there was an increase in patents that improve the efficiency of sulphur scrubbers (Popp, 2003). This effect was confined to early years of the new regime though, and the Program has not provided ongoing incentives for technological advancement (Lange and Bellas, 2005). Early estimates suggested that 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), but the use of scrubber technology as an abatement strategy has declined over time (Burtraw and Szambelan, 2009).² To put it simply, past emissions trading programs like the Acid Rain Program do not provide a precedent for the kind of induced technological change EU policy makers are hoping the EU ETS will provide.

Secondly, if we expected the incentives for technological development to be mediated primarily by augmenting energy prices, historical estimates of the energy price elasticity of energy-saving technology patents might provide very rough idea of the effect the EU ETS might be having. Popp (2002) suggests that, even at the height of the energy crisis of the late 1970s, the hike in energy prices only eventually boosted the share of energy-saving patents by 3.14%. The carbon price in the EU ETS, having ranged from a peak of near €30 to a low of near €0 (and spending more time in the lower part of that range), does not imply anything close to the patenting response seen after the oil shock.³ One might therefore expect the patenting response, if any, to be small. This back-of-the-envelope comparison comes with serious health warnings, of course, not the least of which is that innovation may be driven more by expectations than currently prevailing prices (Martin et al., 2011). Nevertheless, it may aid our expectations about the likely impact of the EU ETS.

²It is worth noting, also, that Title IV of the Clean Air Act, which establishes the Acid Rain Program, also included special provisions that rewarded firms specifically for the use of scrubbers. It is not entirely clear, therefore, how much of the initial investment in scrubbers was the market's doing.

³Popp (2002) estimates that the energy price hike of nearly 10% increased the share of energy-saving patenting by 3.14%. European energy production emitted roughly 355 grams of carbon dioxide per kWh in 2005, and industrial energy users paid about €0.07 per kWh that year. If the average carbon price of €10 was entirely passed onto users, that would imply a circa 5% increase of industrial energy prices, and an eventual boost to patenting of 1.87%. The number is likely to be substantially lower in practice, however, if we account for lower rates of cost-pass-through, and the fact that most low-carbon innovation in Europe takes place in the countries that already have relatively higher energy prices and that are less carbon-intensive to begin with. In France, for instance, even with 100% of regulatory costs passed on to users, one would expect the share of patenting to rise by less than 0.5%.

2.2 The EU ETS and directed technological change

In 2005, the EU ETS launched in 24 countries across Europe, covering roughly 40% of the EU's total greenhouse gas emissions. Power stations and industrial plants across Europe were classified according to their main activity: “combustion”, “cement”, “paper and pulp”, etc. Activity-specific size criteria then determine which installations would be included in the EU ETS. For instance, only combustion installations with a yearly thermal input exceeding 20 MWh were covered. Each year a smaller and smaller number of tradable emissions permits are then allocated to the more than 12'000 qualifying installations, which are each legally required to surrender enough permits every year to cover its emissions. Prior to the compliance date, however, installation operators can freely trade permits with each other (as well as with financial intermediaries and private citizens).⁴ Since 2005, the spot price has varied between €0 and €30. The average price between 2005–2009 was around €10, although the actual price spent more time closer to €0. The price of forward contracts has remained steadily above the spot price, though, suggesting firms are taking the progressive stringency of the cap into account. Installations, or rather the firms that operate them, can then make abatement and investment decisions according to the carbon price revealed in the market.

Since it launched in 2005, there has been vigorous debate about whether the EU ETS would induce firms to develop new emissions-reducing technologies, many arguing that an overly generous allocation of emissions permits would largely undermine the incentives to innovate (Schleich and Betz, 2005; Gagelmann and Frondel, 2005; Grubb et al., 2005). So far, fuel switching appears to have been very important. Fuel switching is a purely organizational innovation, and 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 fuel-mix in favor of natural gas, and therefore reduces the carbon intensity of output. Macroeconomic estimates suggest that the EU ETS reduced total emissions by roughly 50–100 million tonnes of carbon dioxide 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 range from 26–88 million tonnes per year (Delarue et al., 2008, 2010), which suggests that fuel switching

⁴The System has been implemented in 3 trading phases, with successively more stringent emissions caps for each phase. Phase 1, which ran from 2005–2007, was insulated from later phases by prohibiting banking and borrowing of permits across the phase boundary. Phase 2 (2008–2012) and Phase 3 (2013–2020) allow firms to bank unused permits for later use, as well as a limited form of borrowing against future emissions reductions. See Ellerman et al. (2010) for a more comprehensive review of the design and implementation of the EU ETS.

likely accounts for the lion's share of emissions reductions in the EU ETS so far.

This is not a problem in and of itself, of course. As mentioned earlier, the US Acid Rain Program achieved its emissions targets in large part by analogous fuel switching strategies, and with little technological change. However, one should be conscious that in the case of the EU ETS, the capacity for emissions reductions through fuel switching is far more limited relative to the EU's long-term targets. Delarue et al. (2008) estimate that fuel switching has the potential to reduce emission by up to 300 million tonnes annually, 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.⁵

In addition to the evidence on fuel switching, a growing literature of case-studies and expert interviews indicates that, rather than developing new technologies, firms have been introducing well-known technological solutions that had simply not been economically viable without the EU ETS carbon price (Petsonk and Cozijnsen, 2007; Tomás et al., 2010). Martin et al. (2011) conducted interviews with nearly 800 European manufacturing firms, of which almost 450 fell under EU ETS regulations. Using their interview-based measure of innovation, they find a positive effect of the expected future stringency of EU ETS.

Few studies have inquired about more objective proxies of innovation, like R&D or patenting. A survey of Irish EU ETS firms tentatively suggested that almost no resources were made available for low-carbon R&D in the first trading phase (2005–2007), while many of the firms had pursued more operational innovations like installing new machinery or equipment, making process or behavioral changes, and employing fuel switching to some degree (Anderson et al., 2011). Hoffmann (2007), reporting on the German electricity sector, find that the EU ETS has had an effect on decisions about small-scale investments with short amortization times, but not on R&D efforts. Neither study, however, provides a sufficiently large or representative sample of EU ETS firms to provide a reliable picture of the innovation response to the EU ETS. Moreover, neither study offers for comparison a group of non-EU ETS firms.

All of this provides only fragmentary or indirect evidence on directed technological change, however, and it is difficult to summarize our expectations of the EU ETS' impact in terms of a clear quantitative hypothesis. The general literature on induced innovation would lead us to expect the EU ETS to have a positive impact on low-carbon innovation. Studies of earlier emissions trading programs, however, indicate a weaker impact, and studies of the EU ETS so far generally offer an even dimmer view than that. Our

⁵The EU target amounts to reducing annual emissions by roughly 4'500 million tonnes compared to 1990, or roughly 3'500 million tonnes compared to current emission levels.

purpose next, therefore, is to obtain more direct empirical evidence on whether, and to what extent, the EU ETS is encouraging firms to develop new low-carbon technologies.

3 Unpacking low-carbon technological change

While EU ETS regulations apply at the level of the installation, innovation takes place at the level of the firm, and recent advances in linking patent data with company data make it possible to construct firm-level patent portfolios. This paper exploits a newly constructed data set, joining patent portfolios with key firm characteristics, including whether or not the firm operates any installations covered by EU ETS regulations.

Patents have been used extensively as a measure of technological change in the recent induced innovation literature (Popp, 2002, 2006; Johnstone et al., 2010; Aghion et al., 2012), and the advantages and drawbacks of patents are well understood (see OECD, 2009, for a survey). Patents do not weigh or capture all aspects of innovations equally, but are generally considered to provide a useful proxy measure of the output of innovative activity and are available at a highly disaggregated technological level. Having said that, it is also worth noting that patent counts (output) and R&D expenditures (input) have been found to be highly correlated in cross-section (Griliches, 1984), and shift concurrently over time and in response to shocks (Kaufer, 1989).

Our main measure of technological change uses patents filed with the European Patent Office (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.⁶ Nevertheless, as a robustness test we also repeat our analysis using quality-weighted patent counts.⁷

⁶Evidence 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 in which they are filed (van Zeebroeck, 2011). Patents filed at the EPO get patented in 6 EPO member countries on average.

⁷Although the EPO provides a common measure of minimum patent quality, the value of patents is still known to be heterogeneous. We account for the quality of patents in two ways: 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 (patent family information comes from the DOCDB family table in PATSTAT). 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 has the advantage of being more rapidly available than citations (patents are typically mostly cited two years after their publication, hence four years after they are first filed), which is especially valuable

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. This new category (the “Y02” class) is the result of an unprecedented effort by the European Patent Office, whereby patent examiners specialized in each technology, with the help of external experts, developed a tagging system for all patents ever filed at the EPO that are related to climate change mitigation technologies. 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.⁸ It includes, to name a few examples, efficient combustion technologies (e.g. combined heat and power generation), carbon capture and storage, efficient electricity distribution (e.g., smart grids) and energy storage (e.g. fuel cells), which helps us measure the direction of technological change.⁹ A complete list of the sub-classes of low-carbon patents used in the paper can be found in appendix C.

The EPO was set up 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. Our newly constructed data set includes the patent portfolios of over 30 million firms located in 23 countries (22 EU countries, plus the US). 18 of these countries launched the EU ETS in 2005. The other 5 (Norway, Switzerland, Romania, Bulgaria, and the US) have either joined later or have remained outside of the EU ETS altogether. While our data is somewhat more geographically restricted than the EPO, the firms in our data set account for just over 95% of all patents filed at the EPO, so we are confident that we have managed to include the patent history of the vast majority of companies.¹⁰

The share of patents protecting low-carbon technologies shows a distinct pattern over time (figure 1). There was a surge in patenting for these technologies in the early 1980s, often attributed to the second oil price shock in the late 1970s (Dechezleprêtre et al., 2011). The share of low-carbon patents filed each year then stayed roughly level until the mid-1990s, after which it began to rise again. The share of low-carbon patents has

when dealing with recent patents, as we do here.


⁸Importantly, the Y02 class is consistently applied to patents filed both before and after the EU ETS was introduced. See Veeffkind et al. (2012) for more details on how this class was constructed.

⁹We 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. 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.

¹⁰We have also conducted extensive manual double-checking, so we can reasonably assume that companies for which we were unable to locate patent records have not filed any patents 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 EPO filings, which involve high administrative costs.

increased rapidly in recent years, as 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 is not present for any set of environmentally friendly technologies. To see this, figure 1 also plots the share of patents protecting non-greenhouse gas “pollution control technologies”, as defined by Popp (2006),¹¹ 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 has been attributed to the oil price shock, the recent surge might be due to rising oil prices. When comparing the share of low-carbon patenting with the evolution of oil prices (see figure 2), one notices that the present upsurge in patenting follows immediately on the heels of rapid oil price increases 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, clearly, 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, then, we can try comparing the experience of firms regulated under the EU ETS with those not covered by the regulation.

¹¹These technologies pertain to reduction of local pollutants including SO_2 and NO_x .

Both groups will have faced the same oil prices and other macroeconomic conditions, but starting in 2005 they were subject to different regulatory regimes.

Due to a technology supplier’s imperfect ability to appropriate the gains from her invention, economic theory predicts that environmental regulations would produce greater incentives to develop new technologies for regulated firms than for unregulated firms (Milliman and Prince, 1989; Fischer et al., 2003). Even if the System increases the incentive for low-carbon innovation for everyone by creating demand for low-carbon technologies among EU ETS firms, regulated firms receive an additional benefit because they can fully appropriate the gains from reducing their own compliance costs. To this, one may add whatever effects may result from the EU ETS increasing the salience of carbon management issues within regulated companies. It is of course an empirical question whether the EU ETS has encouraged low-carbon innovation for unregulated firms as well, one that we return to in sections 4 and 5, but for now it is enough to realize that the EU ETS is likely to encourage innovation for regulated and unregulated firms to different extents.

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



Our data set also records the regulatory status of 30 million firms—5’568 firms in our data set operate at least one installation regulated under the EU ETS. Together they operate 9’358 EU ETS regulated installations, accounting for over 90% of regulated installations and emissions in Phase 1 in the 18 EU ETS countries we are studying, and roughly 80% of installations and emissions EU ETS-wide (see table 1).¹²

Having identified the subset of firms directly affected by the EU ETS, we can now

¹²See appendix A for more details on how the link between firm data and regulatory data was constructed.

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	Mtonnes of emissions	Percent of installations covered	Percent of emissions covered
Austria	217	97.8	92.2	100.0
Belgium	345	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	77.8	99.9
Finland	637	133.9	84.6	100.0
France	1100	450.2	97.5	99.6
Germany	1944	1486.3	98.6	99.6
Ireland	121	57.7	76.9	94.7
Lithuania	113	34.4	87.6	91.4
Luxembourg	15	9.7	100.0	100.0
Netherlands	418	259.3	87.1	95.6
Poland	869	712.7	90.0	98.6
Portugal	265	110.7	99.2	99.9
Slovakia	191	91.4	90.6	99.9
Spain	1072	498.1	98.5	99.9
Sweden	774	67.6	93.9	98.8
UK	1107	628.0	83.3	97.0
Total	10056	5256.6	93.1	98.7
Total EU ETS	12122	6321.3	77.2	82.0

look separately at the EU ETS and non-EU ETS trends in low-carbon patenting. Figure 3 shows that 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. After 2005, the share of low-carbon patents among EU ETS firms looks to have risen faster than among non-EU ETS firms.¹³ The difference does not become apparent until the start of the second trading phase in 2008, which was widely expected to constrain emissions more tightly than Phase 1 had done. Could the post-2005 surge in low-carbon patenting be a consequence of the EU ETS after all?

Let us naively suppose for a moment that the differences visible in figure 3 are entirely due to the EU ETS. EU ETS firms filed 2’189 low-carbon patents in 2005–2009, compared to 972 patents in the 5 preceding years (an increase of 125%), while non-EU ETS firms filed 19’841 and 12’037 low-carbon patents in the corresponding periods (an

¹³One 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 patented 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.

Figure 3: Comparing the share of low-carbon patents (1978–2009)



increase of 65%). If we then were to assume that the number of low-carbon patents filed by EU ETS firms, had they not been regulated, would have grown at the same rate experienced by non-EU ETS firms, we can naively estimate how many low-carbon patents the EU ETS has added so far: $2'189 - 1.65 \times 972 = 585.2$. This amounts to a 2.6% increase in the number of low-carbon patents at the EPO compared to what it would have been without the EU ETS.

This is clearly a very naive estimate. It assumes that the patenting of non-EU ETS firms provides 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 of these two groups reveals that while only 1 in about 5'500 firms is EU ETS regulated, they account for roughly 1 in 12 low-carbon patents filed in the 5 years before the EU ETS launched. Clearly, EU ETS companies are not representative. One could quite easily imagine, then, that some unobserved change or shock (other than the EU ETS) would have had systematically different impacts on these two sets of firms. The naive calculation above cannot isolate the impact of EU ETS in this case.

To address this shortcoming, we need to restrict our view to a subset of companies that are more similar on pre-2005 characteristics. For such a group of firms, it is more difficult to imagine post-2005 changes (apart from the EU ETS) that would have systematically different impacts on the patenting activities of EU ETS and non-EU ETS firms. Rather than comparing all EU ETS firms with all unregulated firms, this more

restricted comparison is likely to yield a better estimate of the impact of the EU ETS. Let us now turn, therefore, to the task of constructing such a comparison.

4 The direct impact of the EU ETS

4.1 Matching

Comparing two groups of firms that are more similar prior to 2005 makes it more difficult to explain away any difference in outcomes by factors other than the EU ETS. Ideally one would like to match each EU ETS firm with one or more non-EU ETS firms with similar resources available and facing similar demand conditions, regulations (other than the EU ETS), input prices, etc. Because of how the EU ETS was designed and implemented, this is at least theoretically possible. Regulatory status is determined by applying inclusion criteria to installations, not firms. For instance, installations for which the main activity is “combustion of fuels” are included only if their annual thermal input exceeds a threshold of 20 MWh. For steel plants, the relevant inclusion criterion is instead that installations have a production capacity exceeding 2.5 tonnes per hour. Installations manufacturing glass and glass fibre are included only if their melting capacity exceeds 20 tonnes per day. These three examples, taken from a longer list, make clear that regulated installations are bound to systematically differ from unregulated installations. Meanwhile, however, this configuration also means that what we refer to as EU ETS and non-EU ETS firms can in principle be identical in all respects relevant to their patenting behaviour, except for the size of a single installation. This allows us, in theory at least, to form groups of similar EU ETS and non-EU ETS firms, although in practice, as we restrict ourselves to more closely matched firms, there will inevitably be a 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 (see, for instance, Dehejia and Wahba, 1999).

Along with patent portfolios, our data set contains information on the country and economic sector in which firms operate,¹⁴ as well as other firm-level information such as turnover and employment. Using this data, we have tried to assign to each of the 5’568 EU ETS firms a group of similar but unregulated firms (setting aside all companies with ownership ties to EU ETS firms, see appendix A). Though, this has not always been

¹⁴Economic sectors are defined at the 3-digit level for the NACE Rev. 2 industry classification. A 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”.

possible, for two main reasons. Firstly, the records of turnover become less and less complete further back in time. In fact, we only have pre-2005 records on the turnover for 3'564 out of the 5'568 EU ETS firms. Secondly, though EU ETS regulations were applied at the installation level rather than directly to the firm, one might expect two very similar firms to receive the same regulatory treatment more than occasionally. Different regulatory fates are possible if, say, an EU ETS firm operates an installation just large enough to be covered by EU ETS regulations, while the matched control operates one or more installations just below the threshold. But even though we have a very large pool of firms to start with, sometimes there will be no such comparators available within the same country and sector. Due to lack of suitable comparators, the sample of EU ETS firms is further reduced to 3'428. We return to the omitted firms below in section 4.3, to consider the possible consequences of dropping them from our sample.


For each of the 3'428 matched 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 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 response to the EU ETS.¹⁵ The resulting matched sample consists of 3'428 EU ETS firms and 4'373 non-EU ETS firms.

Figure 4 compares the empirical distributions of EU ETS and non-EU ETS firms in our matched sample on a few key 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 reject the hypotheses that the empirical distributions differ between the EU ETS and non-EU ETS firms.

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 (though we subject this claim to further scrutiny below).

¹⁵See appendix B or technical details about how the matching was implemented.

Figure 4: Comparison of matched EU ETS and non-EU ETS firms



Panel (a) displays the empirical quantile-quantile (e-QQ) plot for average turnover in the 4 years before the EU ETS (2001–2004). Each dot gives the value for one EU ETS firm and the average for a group of matched non-EU ETS firms, shown on logarithmic scales. 2001 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, once again shown on logarithmic scales.

4.2 Results

Perhaps the most transparent and intuitive way to view the results is with the aid of a simple graph plotting the patenting of matched EU ETS and non-EU ETS firms, side by side, both before and after the EU ETS came into effect (see figure 5). There are several noteworthy features of this graph. Firstly, matching appears to have produced a set of EU ETS and non-EU ETS firms roughly comparable prior to 2005, both in their general level of low-carbon patenting and in that they do not appear to exhibit different trends. Secondly, the two groups begin to diverge after 2005, coinciding with the introduction of the new policy.

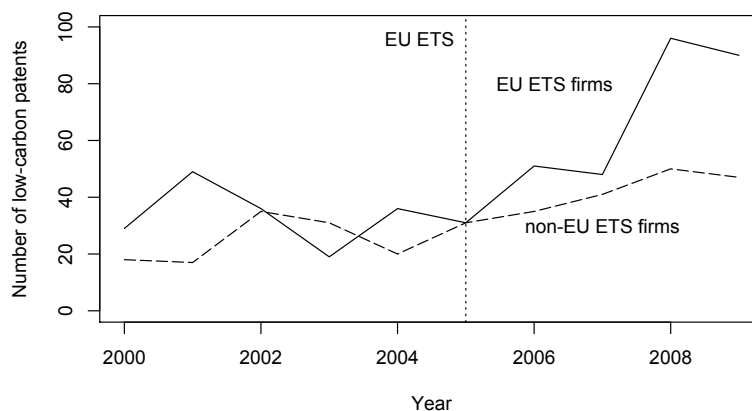
To examine this pattern more precisely, we measure the change in the number of low-carbon patents from 2000–2004 to 2005–2009 for each firm. This means that, even after matching, we take account of any additional time invariant firm-level heterogeneity. The outcomes of the matched control firms are then subtracted from the outcomes of the EU ETS firms to obtain the difference-in-differences. A striking feature of the patent counts used to calculate these difference-in-differences is the large number of zeros. It is a very common feature of patent data that most firms do not file any patents at all, and this arises from a similar censoring problem that usually motivates the use of the

Table 2: Equivalence tests for matched EU ETS and non-EU ETS firms

	Median difference between EU ETS and non-EU ETS firms	Equivalence range	Critical equivalence range (5% sign. lev.)
Turnover (in Mil. Euro)	1.60	± 523.39	± 13.25
Patents	0	± 9.30	± 1.99
Low-carbon patents	0	± 0.25	± 1.99
Year of incorporation	0	± 5.97	± 0.49
Any pre-2005 patents (binary)	Exactly matched	–	–
Economic sector	Exactly matched	–	–
Country	Exactly matched	–	–

The first column from the left reports the median difference between EU ETS firms and non-EU ETS firms in our sample for the key matching variables. Apart from those variables shown in figure 4, matched on the year of incorporation interacted with other variables, since turnover and cumulative patent filings mean different things for an old and new firms. We have also matched exactly for whether (1) or not (0) a firm filed any patents before 2005, for country of operation, and for economic sector (defined at the 3-digit level for NACE Rev. 2). The empirical distributions of EU ETS and non-EU ETS characteristics are judged to be substantively equivalent if the location shift parameter (as defined for Wilcoxon’s signed-rank test) lies within the ‘equivalence range’ reported in the second column. We follow the convention of letting this range be ± 0.2 standard deviations of the distribution of the pooled sample (Cochran and Rubin, 1973; Ho and Imai, 2006). Using Wilcoxon’s signed-rank test, we are just unable reject at the 5% significance level the hypothesis that the location shift parameter lies within the the ‘critical equivalence range’ reported in the final column. (The signed-rank test has been adjusted to account for the fact that our variables are censored at zero, using a method outlined by Rosenbaum (2009, ch. 2). More details in section 4 below.) As can be seen by the fact that the range in the third column is contained within that in the second column, we can reject the hypotheses of substantive differences for all variables, except for low-carbon patents. This last failure to reject is because of the small number of firms that filed any low-carbon patents prior to 2005, as is evidenced by the fact that the same test also fails to reject the hypothesis that the difference is zero. Standard *t*-tests for differences in means reject the hypotheses of substantive differences for all variables (not reported).

Figure 5: Low-carbon patents by matched EU ETS and non-EU ETS firms



Tobit estimator. We can imagine there being a latent variable that can take any value, but we can only observe numbers of zero or greater.

To implement a Tobit estimator in our case, though, we would have to explicitly model the propensity of firms to file at least one patent. This is by no means a straight-

forward exercise, and getting the model wrong carries with it the risk of introducing new biases. The analogous maximum likelihood estimator will likewise generally be inconsistent, especially when applied to panel data (Chay and Powell, 2001). Instead, we can account for the censoring at zero using a Tobit-modified empirical-likelihood estimator, as outlined by Rosenbaum (2009, ch. 2). The idea is as follows. We observe the low-carbon patents filed by EU ETS firms and non-EU ETS firms. In estimating a treatment effect, we would normally search for a number that, if subtracted from each of the observations in one of our two samples, would as nearly as possible equate the distributions of the two samples (using some metric of similarity). The problem, of course, is that this assumes a constant treatment effect that applies even to firms with zero patents. Instead, we can adjust our observed difference-in-differences in a way that takes the censoring into account, and then re-calculate our similarity measure. Each of the difference-in-differences, Δ , is adjusted according to the formula:

$$\Delta = \begin{cases} \max((T_t - T_{t-1}) - \tau, -T_{t-1}) - (C_t - C_{t-1}) & \text{if } \tau \geq 0 \\ (T_t - T_{t-1}) - \max((C_t - C_{t-1}) + \tau, -C_{t-1}) & \text{otherwise} \end{cases}$$

where T_t and T_{t-1} are the numbers of low-carbon patents filed by an EU ETS firm in the treatment period t (2005–2009) and the pre-treatment period $t - 1$ (2000–2004), respectively. C_t and C_{t-1} , are the corresponding numbers for the matched non-EU ETS firms, and τ is the treatment effect. The point estimate of the treatment effect is then the value of τ for which the similarity measure is maximized, and the $(1 - \alpha)\%$ confidence interval is the set of values of τ for which we cannot reject the alternative of difference at the $\alpha\%$ level of significance. We implement this estimator using as our similarity measure the p -value calculated from the Wilcoxon signed-rank test. This provides a non-parametric alternative to the Tobit estimator.

We estimate a treatment effect of $\tau = 2$ additional low-carbon patents for our EU ETS firms, with a 95% confidence interval of (1, 5). The matched EU ETS firms filed a total of 316 low-carbon patents in the period 2005-2009. Subtracting 2 low-carbon patents from each of our matched EU ETS firms (and accounting for censoring at zero) tells us that these firms together would have filed 232 low-carbon patents in the absence of EU ETS regulations. Our estimated treatment effect therefore implies that EU ETS has prompted 84 (53, 129) additional low-carbon patents amongst our sample of EU ETS firms, or an increase of 36.2% (20.2%, 69.0%) compared to what we expect would have happened in the absence of the EU ETS. Because these firms only account for a small portion of all patents, however, this remarkable impact translates into an increase of low-carbon patenting at the EPO of only 0.38% (0.24%, 0.58%) compared to what we expect

it would have been in the absence of the EU ETS. If we think our estimate applies to all of the 5'568 EU ETS firms, we can use their patenting records to calculate that, once we account for censoring at zero, the EU ETS is responsible for 188 (114, 319) additional low-carbon patents. This amounts to a 8.1% (4.7%, 14.5%) increase in their low-carbon patenting, or a 0.85% (0.51%, 1.45%) increase in the total number of low-carbon patents filed at the EPO in 2005–2009 compared to the counterfactual. The first thing to note about these numbers is that they are substantially smaller than what was suggested by our naive calculations above (585.2 additional low-carbon patents, or a 2.6% increase in low-carbon patents at the EPO, see table 3). Second, because these numbers are so small relative to the totals, it is likely we would not have recognized the impact to be anything different from zero, had we been studying patent counts at a more aggregated level.

To address the issue of the *direction* of technological change, we must compare this with the impact on patenting for other technologies. Environmental regulations like the EU ETS could in principle increase patenting for other technologies as well. For instance, even if they are not classified as low-carbon technologies, they may be complementary to low-carbon technologies. More generally, environmental regulations that increase the cost of production can in principle encourage patenting for any technology that reduces it, be it a low-carbon technology or not. The induced innovation hypothesis holds that a policy like the EU ETS would have a disproportionate impact on low-carbon technologies, but this is an essentially empirical matter. There is a related concern, also, that the increase in low-carbon innovation will actually displace, or crowd out, development of other technologies (Popp and Newell, 2012). We can address these questions using the same matched sample and estimator described above. We estimate that the EU ETS has added on average 1 other patent (1, 1.99). This translates into 305 (305, 512.9) additional patents for other technologies, which represents an increase of 1.9% (1.9%, 3.2%) in their patent filings for non-low-carbon technologies, or a 0.041% (0.041%, 0.068%) increase in patenting for other technologies at the EPO. Comparing these numbers with the estimates from the previous paragraph, we see that the EU ETS has had a disproportionate impact on patenting for low-carbon technologies: 36.2% vs. 1.9% (difference is significant at 5% level). Put another way, the System has nearly had a 20 times greater impact on low-carbon patenting, but it has not crowded out patenting for other technologies. If we think our estimate applies to all of the 5'568 EU ETS firms, the EU ETS would be responsible for 554 (554, 963.86) additional other patents, which amounts to a 0.77% (0.77%, 1.34%) increase in their other patenting, or a 0.074% (0.074%, 0.13%) increase in the total number of other patents filed at the EPO

in 2005–2009.

The EU ETS may also have affected the direction of change within the class of low-carbon technologies itself, encouraging more patenting for certain types of low-carbon technologies. Unfortunately, our firm-level identification strategy is ill-suited to look at patenting at such a disaggregated level. Due to the large number of zeros typically present in patent datasets, the small number of regulated companies active in each sector, and the even smaller number of patents each firm holds in a particular technology class, this method does not yield informative technology-level estimates. However, once we have estimated that each EU ETS firm filed 2 additional low-carbon patents, it is a small step to consider what types of technologies those patents protect (i.e. conditional on the estimated treatment effect). Since firms often hold several patents protecting different technologies, there is no definite way of identifying which 2 low-carbon patents were additional. If we look at the average across all possible permutations, however, we find that most of the additional low-carbon patents appear to protect alternative energy and energy storage. The focus of the remaining ones is on energy efficiency and carbon sequestration. Alternative energy technologies appear to account for a greater number of additional low-carbon patents than do improvements of conventional combustion technologies. Broken down by economic sector, most of the additional low-carbon patents belong to chemicals manufacturers, energy companies, and automobile manufacturers (see appendix D for further explanation of the methodology as well as for all the technology- and sector-level estimates). These stylized conclusions should be read more as indicative than final, though, and since they are conditional on our estimated treatment effect, their soundness ultimately depends on the robustness of our earlier estimates.

Our main results are summarised for convenience in table 3, along with comparable naive estimates for the full sample of EU ETS firms (calculated as in section 3). The naive estimates substantially overestimate the impact of the EU ETS, yet they display the same general pattern as our matching estimates, showing increases in patenting for both low-carbon and other technologies, but with a pronounced direction. The matching estimates suggest the EU ETS has had a positive and notable impact on low-carbon patenting among EU ETS firms, though the impact appears much smaller relative to the overall pace of low-carbon technological development, boosting low-carbon patenting by only a fraction of a percent. On the one hand, our findings contradict early prognostications that over-allocation of emissions permits in the EU ETS would completely undermine the incentives for low-carbon innovation. On the other hand, even a quite remarkable response among EU ETS firms—whether 36.2% among matched EU

Table 3: Summary of results

	Matching estimates		Naive estimates
	Matched sample	Full sample	Full sample
Additional low-carbon patents	84 (53, 129)	188 (114, 319)	585.2
As % increase	36.2 (20.2, 69.0)	8.1 (4.7, 14.5)	36.5
As % increase of EPO	0.38 (0.24, 0.58)	0.85 (0.51, 1.45)	2.6
Additional other patents	305 (305, 512.9)	554 (554, 963.86)	9072.8
As % increase	1.9 (1.9, 3.2)	0.77 (0.77, 1.34)	16.0
As % increase of EPO	0.041 (0.041, 0.068)	0.074 (0.074, 0.13)	1.2

Point estimates, along with 95% confidence intervals in brackets where applicable. The matched sample estimates consider the impact only for the 3'426 matched EU ETS firms, while full sample estimates consider the impact for all 5'568 EU ETS firms in our data set. The matching estimates are calculated using our point estimates of τ obtained for the matched sample of 3'426 EU ETS firms and 4'373 non-EU ETS firms. Naive estimates are included for comparison. They have been calculated using the full set of 30 million non-EU ETS firms to construct a counterfactual, as in section 3.

ETS firms or 8.1% among the full sample—translates into rather small impact from an economy-wide perspective, less than a 1% increase at the EPO. Putting it another way, of the post-2005 surge in low-carbon patenting seen in figure 1, roughly 2% can be attributed to the EU ETS.¹⁶ It is worth noting that this apparently small impact relative to the overall pace of technological change is not simply an arithmetical artifact of the small number of EU ETS firms, however, as is demonstrated by the fact that the naive estimator is more than three times higher.

Before settling on an interpretation of these estimates, though, we must ask whether they are really best explained by the EU ETS having had a very small impact. Perhaps these small numbers should instead caution us that we may have underestimated the impact? Let us therefore investigate challenges to the internal and external validity of our results.

4.3 Robustness tests

Is our conclusion driven by an omitted variable? The primary challenge for any matching study is to justify the assumption that firms that appear similar are similar in

¹⁶The number of low-carbon patents filed at the EPO increased by 9054 from the period 2000-2004 to 2005-2009. The 188 additional low-carbon patents we have attributed to the EU ETS correspond to 2% of this increase. Even under the more generous framing that the upward trend from 2000-2004 would have continued unabated in 2005-2009, the post-2005 'surge' was only 4725.5 low-carbon patents, of which the 188 additional low-carbon patents would amount to barely 4%.

unmeasured dimensions as well—often called ‘selection on observables’. 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 read to as evidence 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. As figure 6 and table 4 show, the empirical distributions of number of employees of the EU ETS and non-EU ETS 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.

Figure 6: Comparison of matched EU ETS and non-EU ETS firms on ‘unobserved’ variable




Table 4: Equivalence test for matched EU ETS and non-EU ETS firms on ‘unobserved’ variable

	Median difference between EU ETS and non-EU ETS firms	Equivalence range	Critical equivalence range (5% sign. lev.)
Employees	25	± 904.07	± 106.75

See caption of table 2 for details on how to read this table.

This test, though reassuring, is perhaps too simplistic. Other unobserved differences between regulated and unregulated firms might still bias our findings. Such differences might arise, for instance, if firms could influence to some degree whether or not they would be regulated by the EU ETS. In general, there is very little evidence to suggest that firms had such influence; most of the installation-level inclusion criteria already appeared in draft legislation in 2002, and have remained unchanged to this day. One small exception, though, is the debate over whether or not to regulate installations that produce chemicals or aluminum. These types of installations were to be regulated according to the 2002 proposal, but were omitted from a later draft, before a final compromise that allowed chemicals and aluminum installations to opt-in to the EU ETS (Markussen and Svendsen, 2005). Ultimately, 575 such installations opted in, slightly less than 5% of all EU ETS installations. Our estimates might be biased, then, if the firms with chemicals or aluminum installations that are opting into the EU ETS are systematically different from their non-EU ETS counterparts in some unobserved dimension that is predictive of patenting responses. To see whether our estimates are biased by the possibility of self-selection in this subset of firms, we re-estimate the treatment effect after dropping any matched pairs where the EU ETS firm has opted in at least one of its installations (this reduces our sample size by nearly one hundred matched pairs). This returns an estimate of 2 (1, 5.99) additional low-carbon patents, and of 1 (1, 1.99) other additional patents. These estimates are identical to our original specification (although one of the confidence intervals is slightly wider), offering no indication that our estimates are biased by self-selection.

The two above tests look at specific sources of omitted variable bias. Neither tests finds evidence to suggest that our estimates are biased by variable omissions, but the possibility remains that our estimates are confounded by bias from some unknown source. Let us therefore ask the more general question: what kind of an omitted variable could in principle undermine confidence in our estimate?

Imagine that we have an omitted binary variable that is negatively correlated with EU ETS regulations and positively correlated with increases in low-carbon patenting (or vice versa). This could be, for instance, a variable that tells us whether a firm would be covered by a complementary carbon policy that targets the types of firms unlikely to be regulated by the EU ETS. Omitting such a variable would cause us to underestimate the impact of the EU ETS. Using the model for sensitivity analysis developed by Rosenbaum (1987) and Rosenbaum and Silber (2009), we can infer precisely how large the omitted variable bias would have to be in order to undermine confidence in our estimate relative to some larger alternative.

In order for our 3'428 matched EU ETS firms to have boosted the number of low-carbon patents filed at the EPO by 5%, say, they would have to have filed 1062 additional low-carbon patents. Since they did not file this many low-carbon patents in 2005–2009 in total, we can comfortably rule out that the EU ETS would have had such a large treatment effect even if all of the patents were additional. To have boosted low-carbon patents by just 1%, 223 of their low-carbon patents would have to have been additional. This translates back into a treatment effect of $\tau = 20.4$ —more than 10 times higher than our original estimate. In order to increase our point estimate beyond this level, we would have to postulate an omitted variable that, if observed before 2005, would successfully predict more than 83 times out of a 100 (a) which firm in our matched pairs escapes EU ETS regulations *and* (b) which firm in our matched pairs would most increase their low-carbon patenting. Even if the omitted variable predicted (a) almost perfectly, it would still have to predict (b) 73 times out of 100. For the milder threshold of just being unable to reject the hypothesis that the the treatment effect is 20.4, we would still have to postulate an omitted variable that makes these prediction successfully more than 70 times out of 100.¹⁷ We have estimated above that our sample of matched EU ETS firms account for only a 0.38% increase in low-carbon patenting at the EPO. If one finds an example of a complementary policy that was implemented in such a systematic fashion across the EU and caused such a predictable boost in the low-carbon patenting, we would have to concede that they may have boosted low-carbon patenting by as much as 1%. Even then, it is not obvious that this would seriously challenge the conclusion that the EU ETS has had but a limited direct impact on low-carbon patenting overall.

Another category of potential omitted variables are those generally expected to be positively correlated with both a firm's chances of becoming regulated and with their chances of increasing their low-carbon patenting. Examples of this include, for instance, whether a firm had high or low carbon emissions prior to 2005, or a complementary carbon policy that targets the same types of firms regulated under the EU ETS. The omission of a variable with these properties would imply we have overestimated the impact of the EU ETS above. To reduce our point estimate to zero, we would need to postulate an omitted variable that predicts more than 81 times out of 100 (a) which firm in our matched pairs became EU ETS regulated and (b) which firm in our matched pairs would most increase their low-carbon patenting. It would need to make these predictions successfully more than 71 times out of 100 to make us just unable to reject at the 5%

¹⁷In Rosenbaum's notation, it is just possible that the estimated treatment effect is 20.4 for a sensitivity parameter of $\Gamma = 2.65$, and we are just unable to reject this treatment effect at the 5% significance level for $\Gamma = 1.4$. This can be decomposed into the biases present in treatment assignment and outcomes using propositions in Rosenbaum and Silber (2009).

level the hypothesis that the treatment effect is really zero.¹⁸ In appendix E, we examine two suggested omitted variables—company growth rates and the number of innovation locations—but neither predicts a firm’s EU ETS status well enough to challenge our conclusions.¹⁹

In sum, matching has achieved balance on at least one ‘unobserved’ characteristic, which might suggest it has balanced other unobserved variables as well, like a truly randomized experiment would have. Even if this is not the case, though, it appears our estimate of the low-carbon treatment effect is reasonably robust to both negative and positive omitted variable biases.

Are the estimates valid beyond our sample? A more serious challenge to our conclusion, perhaps, is to justify extrapolating from our sample of 3’428 EU ETS firms to all EU ETS firms. This type of calculation might lead us to underestimate the impact of the EU ETS if the firms omitted from estimation have had a systematically stronger reaction compared to those firms in our sample. This is a question of selection bias.

The first thing to look at is whether the EU ETS firms we have matched successfully exhibited substantially different patenting behaviour prior to 2005 from the EU ETS firms dropped from our matched sample. An unmatched EU ETS firms would have been dropped either because it was an outlier or because crucial data was missing that prevented matching. In practice, most were dropped because financial data was missing. This has two consequences. Firstly, we can only reliably compare the patenting behaviour of matched and unmatched EU ETS firms. Secondly, there may be substantial overlap in the levels of patenting of matched and unmatched EU ETS. Keeping in mind that some proportion of the unmatched EU ETS firms are probably outliers, though, matched EU ETS firms are likely to have slightly lower patenting levels on average.

Let us apply the same procedure used in table 2 to compare matched EU ETS and non-EU ETS firms.²⁰ For low-carbon patenting, we cannot reject the hypothesis that the empirical distributions of matched and unmatched EU ETS firms are different, although we can reject at the 5% level of significance the hypothesis that the two distributions differ by a shift-parameter greater than ± 1.99 (equivalence range: ± 0.72 ; critical equivalence range for 5% sign. level: ± 1.99). This mirrors our findings in table 2, and follows in

¹⁸In Rosenbaum’s notation, it is just possible that the estimated treatment effect is 0 for a sensitivity parameter of $\Gamma = 2.34$, and we are just unable to reject this treatment effect at the 5% significance level for $\Gamma = 1.45$.

¹⁹We would like to thank the anonymous referee for this suggestion.

²⁰Since matched and unmatched EU ETS firms are not paired, we here substitute Wilcoxon’s rank-sum test for the signed-rank test.

part from the relative rarity of low-carbon patents. For other patents, we can reject the hypothesis that the empirical distributions are substantively different (equivalence range: ± 34.52 ; critical equivalence range for 5% sign. level: ± 1.99). The sectoral composition is somewhat different for matched and unmatched EU ETS firms, but all economic sectors with at least a handful of unmatched EU ETS firms are also well-represented among our matched firms, including in the electric power generation, transmission, and distribution sector. Naturally, matched and unmatched EU ETS firms are not identical—if they were, we would have been able to match them all in the first place (apart from where data was missing). Nevertheless, our tests here suggest that unmatched EU ETS firms do not appear to be substantially different from the EU ETS firms in our matched sample, which is perhaps reassuring for our earlier attempt to extrapolate.

This may not entirely allay one's concern that matched and unmatched EU ETS firms have had systematically different reactions to the EU ETS. Maybe there was selection on some other relevant variable that we are unable to check. We can address this concern in three ways: (1) increasing the sample size by matching some of those unmatched EU ETS firms, (2) calculating an upper bound for our estimates, and (3) calculating a lower bound for the out-of-sample response necessary to qualitatively affect our conclusions. Firstly, because turnover figures become more widely available in 2005, we are able to increase sample size if we 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 greater number of EU ETS and non-EU ETS firms, it may still provide a reasonable test of whether our findings apply to the EU ETS more broadly.

Matching using 2005 turnover figures allows us to successfully match an additional 427 EU ETS firms, producing 3'855 matched groups in total. The point estimates for this sample are 2.75 (1, 5.99) for low-carbon patents and 1 (1, 1.99) for other patents, which is almost identical to our original estimates. The typical matched firm still looks much the same, which is what one would expect if we were simply finding more firms around the same EU ETS thresholds. The EU ETS firms in our original matched sample therefore appear to be representative of a larger portion 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 atypical companies for which suitable unregulated matches could not be found the first time around.

It is, nevertheless, possible to bound the effect that these atypical firms can have on

the impact estimates. Suppose we were able to perfectly match every one of the 2'140 EU ETS firms we were forced to omit. Suppose further that the hypothetically matched non-EU ETS firms have not filed any patents since 2005, a strict lower bound. Because we observe the low-carbon patenting of the EU ETS firms, these two assumptions allow us to calculate the upper bound difference-in-differences for each of these 2'140 EU ETS firms. Pooling them with the 3'428 previously difference-in-differences, we can then estimate the upper bound of the treatment effect.²¹ This procedure produces point estimates of 13 (4, 43.99) for low-carbon and 6 (4, 10.99) for other patents. These high point estimates are driven in large part by a small number of prolific patenters that were previously omitted, but are now matched to hypothetical non-EU ETS firms with zero patents after 2005. Subtracting a large number of patents from each firm and accounting for censoring at zero, therefore, does not add as many patents as the higher point estimates perhaps might suggest. The new estimates translate into 524 (275, 952.9) additional low-carbon patents and 2093 (1582, 3176.95) additional other patents, or increases of 26.7% (12.4%, 62.2%) and 3% (2.3%, 4.7%) respectively. While there is still a clear direction to induced technological change, it is less pronounced than for our original estimates. In comparison with the total numbers of patents that would otherwise have been filed at the EPO in each category in this period, the additional patents represent a 2.4% (1.2%, 4.5%) increase in low-carbon patenting and a 0.28% (0.21%, 0.42%) increase in patenting for other technologies. In economic terms, the upper bounds are perhaps slightly more noteworthy than our original estimates, though we are now very aware of the kind of extremely favorable and unrealistic assumptions needed to generate such results.

Our third strategy to address concerns about external validity is to calculate what out-of-sample response would be necessary in order to qualitatively affect our conclusion. Our sample covers 9'358 out of the 12'122 installations that fell under EU ETS regulation in 2005 (see table 1). In order for the EU ETS to have boosted low-carbon patenting by 5%, say, EU ETS firms would together have to have filed 1062 additional low-carbon patents in 2005–2009. Subtracting our best estimate of 188 additional low-carbon patents for the 5'568 firms operating 9'358 EU ETS installations, this leaves the operators of the remaining 2'764 installations to have filed 874 additional low-carbon patents. To put it another way, we estimate that the average EU ETS firm in our sample filed roughly 0.03 extra low-carbon patents, but even if the remaining 2'764 were operated by as many firms (another charitable assumption), the EU ETS firms outside our sample would have

²¹This bound is analogous to the sharp bounds derived by Manski (2007, ch. 2) for situations with missing data. The bound is sharp in the sense that it does not impose any restrictions on the process that leads to 'missingness'.

to have filed 0.32 additional low-carbon patents in the same period. The out-of-sample response would have to be 10 times greater than the in-sample response. Even if we use the upper bound estimate (in-sample firms filed 524 additional low-carbon patents), the out-of-sample firms would have to have filed 538 extra low-carbon patents, or at least 0.19 per firm, which is still more than twice the upper bound for our in-sample firms (0.09). These strong responses appear especially unlikely in light of the fact that most of the out-of-sample firms operate in countries with lower patenting propensities (Cyprus, Greece, Hungary, Italy, Latvia, and Slovenia).

It seems, therefore, that none of the strategies to address concerns about external validity—increasing sample size, computing upper bounds, and calculating necessary out-of-sample responses—seriously challenge our earlier conclusion. The EU ETS appears to have had a positive and notable impact on low-carbon patenting among EU ETS firms, but partly because these firms account for a small proportion of low-carbon patents, the direct impact on low-carbon technological change has been much more limited on a European scale.

Other robustness tests. Above we have tried to address the most pertinent challenges to our interpretation of the results, but one can imagine still other explanations for why the direct impact of the EU ETS appears to have been so small. We have tried to test several of these:

- Are matched non-EU ETS firms also responding to EU ETS? If so, firms less exposed to the EU ETS and to direct competition with EU ETS firms would perhaps be expected to respond less. We re-matched our EU ETS firms to similar firms in Norway, Switzerland, Bulgaria, and Romania (4 countries that did not launch the EU ETS in 2005, and two of which have remained outside). We also re-matched our EU ETS firms to similar US firms. Neither comparison returns an estimate of the treatment effect significantly different from that reported above (see appendix E for further details).
- Did the main patenting response occur after the Directive was adopted in 2003, but before the EU ETS launched in 2005? Some authors have highlighted the possibility that firms patent in anticipation of new regulations (Dekker et al., 2012). To address this concern, we re-matched our EU ETS firms using 2003 as the treatment year instead of 2005. The treatment effect for the period 2003–2004 actually indicates that prospective EU ETS firms would actually have filed 1.75 additional low-carbon patents *if not* for the EU ETS, though the number is not

significantly different from zero. In other words, there is no significant difference in the low-carbon patenting activities of EU ETS and non-EU ETS firms in this period.

- Is the result an artifact of how we measure low-carbon patents? To address this, we looked at using an expanded definition of low-carbon patents. This does not materially affect our conclusions, however. Nemet (2009) and Hoppmann et al. (2013) raise a related concern, that a policy like the EU ETS might discourage non-incremental innovation (more likely to be counted as high value patents). However, we do not find evidence that the quality of patents held by EU ETS firms (measured by citations and family size) has changed relative to non-EU ETS firms (see appendix E for more details).
- Is there some other hidden bias? Perhaps we are only picking up the low-carbon technology component of a broader trend toward environmental technologies going on among our EU ETS firms. We look at the number of patents filed by matched EU ETS and non-EU ETS firms protecting other ‘pollution control technologies’, as defined by Popp (2006). Since these technologies do not help mitigate emissions covered under the EU ETS, we would not expect the EU ETS to have had any impact. A hidden bias in our study design, perhaps some unknown omitted variable, would manifest itself as finding a treatment effect here that is significantly different from zero. Our estimated treatment effect is $\tau = 0.75$, but it is not significantly different from zero.²²

For convenience, table 5 summarizes the results from robustness tests that involved re-estimating the treatment effect under alternative assumptions. More information, and a few additional robustness tests, can be found in E.

It appears, then, that EU ETS has had a positive and notable impact on low-carbon patenting among EU ETS firms. It has spurred development of low-carbon technologies without crowding out innovation for other technologies. Since EU ETS firms account for only a small proportion of low-carbon patents, however, the impact on EU ETS regulated firms is negligible on a European scale. None of the above challenges seems to offer a compelling alternative explanation to this interpretation of the results.²³

²²Roughly 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.

²³One must be careful also because some 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). The point here, however, is that to replicate our results each time,

Table 5: Summary of treatment effect estimates

	Low-carbon	Other
Original estimate	2 (1, 5)	1 (1,1.99)
<i>Alternative specifications</i>		
Excluding opt-ins	2 (1, 5.99)	1 (1, 1.99)
Matching with 2005 turnover	2.75 (1, 5.99)	1 (1, 1.99)
Expanded low-carbon definition	1.75 (1, 3.99)	1 (1, 1.99)
Non-EU ETS firms from Norway, Switzerland, Romania, and Bulgaria	1 (0, 1.99)	2 (1, 3)
Non-EU ETS firms from USA	-1 (-1.99, 0.99)	0 (-0.99, 0.99)
Treatment years 2003–2004	-1.75 (-∞, 1.99)	-1 (-4, -0.01)
<i>Upper bounds</i>		
Assuming 1% boost to EPO low-carbon patenting	20.4 –	–
Assuming all patents of unmatched EU ETS firms are additional	13 (4, 43.99)	6 (4, 10.99)

If we accept that the impact of the EU ETS on regulated firms does not account for the post-2005 surge in low-carbon patenting seen in figure 1, might the EU ETS still be indirectly responsible? Has it encouraged third parties to develop low-carbon technologies in the hope of selling or licensing them to newly regulated EU ETS firms? We investigate this question next.

5 The indirect impact of the EU ETS

The preceding analysis strongly suggests that the direct impact of the EU ETS has not been sufficient to account for the apparent surge in low-carbon patenting since 2005. Could the impact of the EU ETS instead have been largely indirect, spurring third parties to develop new low-carbon technologies?

There are three major reasons why we would expect the indirect impact to be comparatively small. Firstly, since technology providers cannot perfectly appropriate the gains from their technologies, economic theory predicts that environmental regulations would produce greater incentives to develop new technologies for directly regulated firms than for third parties (Milliman and Prince, 1989; Fischer et al., 2003). The asymmetry

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.

arises because the latter group is not discharging costly emissions themselves and receive no additional benefit reducing its own compliance cost. To the extent that the EU ETS is encouraging low-carbon technological change, therefore, economic theory predicts this response to be strongest among EU ETS firms.

Secondly, EU ETS firms have filed over 120'000 patents with the EPO since 2000, circa 2.5% of which protect low-carbon technologies. These are clearly firms with above average innovation capabilities. To argue that the bulk of the response to the EU ETS comes from third-party technology providers amounts to saying that these EU ETS firms with well-developed low-carbon innovation capabilities are responding mostly by purchasing technologies from others, rather than developing the technologies in-house to suit their own specific needs.

Thirdly, the EU ETS firms in our sample are very likely technology providers themselves. As highlighted in the previous paragraph, EU ETS firms do develop new technologies themselves, including low-carbon technologies. While some firms may innovate in the hope of meeting new demand from EU ETS firms, others might expect greater opportunities to purchase the technologies developed by EU ETS firms. The indirect impact of the EU ETS is the net of these two responses.

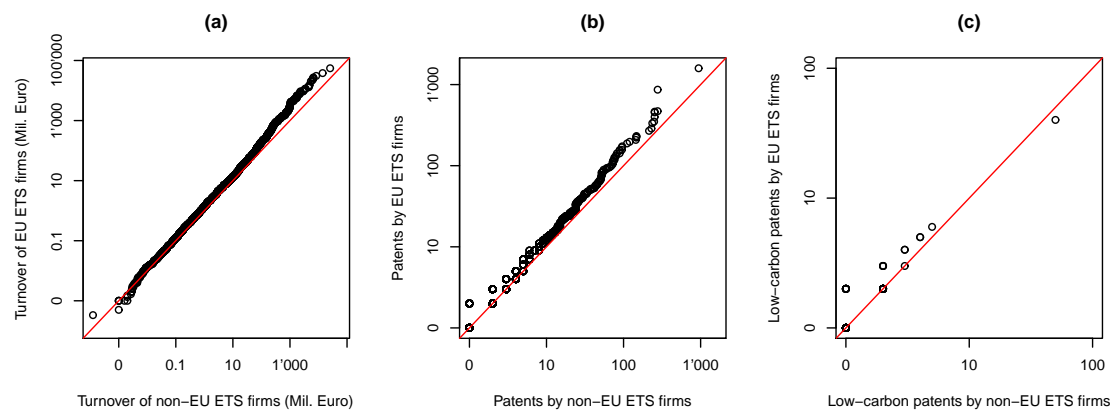
These three reasons suggest that the indirect impact of the EU ETS would be comparatively small, but all claims about the indirect effect need to be met with the same level of skepticism as any other empirical hypothesis. It is a very difficult task to cleanly estimate the indirect impact of the EU ETS, not least because of the difficulty involved in identifying firms more likely to either provide new technologies to EU ETS firms or to which EU ETS firms are more likely to provide new technologies. We can, nevertheless, make a start.

Consider the set of firms that had filed at least one patent jointly with an EU ETS firm prior to 2005. A joint patent filing records a technological partnership with an EU ETS firm. One might then expect these firms to be more likely than an average non-EU ETS to either provide technologies to EU ETS firms once the regulations came into force, or to demand new technologies from EU ETS firms. They are likely to be good candidates for studying the indirect impact of the EU ETS. By comparing this set of firms with other non-EU ETS firms, therefore, we might hope to gain at least some partial insight as to the net indirect impact of the EU ETS. It is worth noting, though, that while technology provision is an asymmetric relationship, co-patenting is of course symmetric. Hence, we cannot separate co-patenters into technology providers and demanders even if each co-patenter could in principle be classified as one or the other. Nevertheless, we can provide an indicative estimate of the *net* indirect impact of

the EU ETS.

From patent records we can identify 11'603 non-EU ETS firms that each filed at least one patent jointly with an EU ETS firm in 1978–2004. Many of these firms are no longer active or operate in countries not in our data set, which prevents us from matching them. Additionally, as before there are many firms for which historical data are missing, and a few for which we simply cannot find suitable comparators. Our matched sample therefore contains 2'784 co-patenters and 19'361 similar firms that had not filed a joint patent with an EU ETS firm prior to 2005.²⁴ Figure 7 and table 6 show the properties of our matched sample.²⁵

Figure 7: Comparison of matched co-patenters and non-co-patenting firms



We estimate a treatment effect of $\tau = 0.99$ additional low-carbon patents among our co-patenters, with a 95% confidence interval of $(-0.99, 1.99)$. We cannot say with confidence, therefore, that the EU ETS has had any net impact on the low-carbon patenting of co-patenters. Even taking the point estimate at face value, it translates into a mere 47.52 additional low-carbon patents. Although it would represent a quite dramatic response, on the order of a 32.4% increase compared to what it would have been without the EU ETS, it would still translate into a negligible increase relative to

²⁴Compared to when EU ETS firms were matched earlier, finding a single good comparator here was a good indicator that there were many good comparators available. We have kept all of these comparators in our matched sample to reduce the variance of our estimates.

²⁵On average, co-patenters have historically filed more patents than EU ETS firms. It is no mystery why—to be a co-patenters a firm must have filed at least one patent prior to 2005, while EU ETS firms had no such requirement to meet.

Table 6: Equivalence tests for matched co-patenters and non-co-patenting firms

	Median difference between EU ETS and non-EU ETS firms	Equivalence range	Critical equivalence range (5% sign. lev.)
Turnover (in th. Euro)	14.90	$\pm 304'382.80$	$\pm 1'421.00$
Patents	0	± 7.07	$\pm < 0.01$
Low-carbon patents	0	± 0.17	± 0.99
Year of incorporation	0	± 5.48	± 0.50
Any pre-2005 patents (binary)	Exactly matched	–	–
Economic sector	Exactly matched	–	–
Country	Exactly matched	–	–
Employees	1.66	$\pm 1'613.82$	± 20.66

See caption of table 2 for details on how to read this table. Again, the failure to reject the hypothesis of difference for low-carbon patents is a consequence of the small number of firms that filed any low-carbon patents prior to 2005. The same test also fails to reject the hypothesis that the difference is zero. Standard t -tests for differences in means reject the hypotheses of substantive differences for all variables (not reported). For completeness, the results from the robustness test of checking balance on employees is also included at the bottom of this table.

the number of low-carbon patents filed at the EPO (0.2%). Extrapolating the number to all 11'603 co-patenters would naturally make it look as if the EU ETS has had a more impressive indirect impact, but since the estimate does not even stand up to a conventional significance test, such an exercise is not likely to be informative.

The picture is not much different for other technologies either. We estimate that the EU ETS has on average *subtracted* 0.745 other patents ($-0.99, -0.01$) for co-patenters. We are just barely able to reject the hypothesis that the effect is actually zero, but this rejection does not withstand even the slightest challenge to robustness. Moreover, even if the point estimate were true, it would suggest that the EU ETS has crowded out patenting for non-low-carbon technologies among co-patenters.

These numbers offer no compelling evidence that the EU ETS has had an indirect impact on patenting. A patent filed jointly with an EU ETS firm is a record of a technological partnership, be it the case that the co-patenter has provided technologies to EU ETS firms or vice versa. In either case, one would expect that co-patenters are more likely than an average non-EU ETS firm to supply new technologies to EU ETS firms once the EU ETS launched, or to demand new technologies from EU ETS firms. Yet, taken together, co-patenters appear to behave no different to other non-EU ETS firms. It is of course incredibly difficult to identify potential technology providers and demanders for the purposes of estimation, so our results should not be over-interpreted. Nevertheless, our findings can perhaps be read as a reasonable indication that the EU ETS has had no net indirect impact on directed technological change. At the very least, it poses an empirical challenge for those wishing to argue otherwise.

6 Discussion

The EU ETS launched in 2005, amid both promises and pessimism. An important objective of carbon market programs like the EU ETS is encourage the development of low-carbon technologies (Stavins, 2007; European Commission, 2005, 2012). In this paper we have investigated the System’s success in this regard during the 5 years subsequent to its launch.

A casual look at aggregate patenting suggests there has been an increase in low-carbon patenting since 2005, but there are several obstacles to isolating the impact of the EU ETS. Comparing patenting behaviour prior to and after 2005 risks conflating the impact of the EU ETS with other changes, like rapidly rising oil prices. Yet, looking only at the period after 2005 and comparing EU ETS regulated firms with those that escaped regulation risks conflating the impact of the EU ETS with other systematic differences in firm characteristics that might also drive patenting. Employing a matched difference-in-differences study design has permitted us to account for firm-level time-invariant heterogeneity, and to isolate that part of the change that does not depend on systematic differences in firm characteristics.

We find evidence that the EU ETS has had a strong impact on the patenting behaviour of EU ETS regulated firms. Our best estimate for a sample of 3’428 EU ETS firms implies that the System has increased their low-carbon patenting by 36.2% compared to what we expect would have happened in the absence of the EU ETS. What is more, our estimates suggest that the System has also encouraged EU ETS firms to increased their patent filings for non-low-carbon technologies by 1.9%. The EU ETS thus appears to have had a disproportionate impact on patenting for low-carbon technologies, but it has not crowded out patenting for other technologies.

Extrapolating our point estimates to 5’568 EU ETS firms across 18 countries (and accounting for censoring at zero), the EU ETS would account for an 8.1% increase in low-carbon patenting and a 0.77% increase in patenting for other technologies. Because of the targeted nature of EU ETS regulations, however, these responses translate into a quite unremarkable nudge on the pace and direction of technological change—a 0.38% boost to low-carbon patenting at the EPO (0.85% for the full sample), and a meagre 0.041% boost to patenting for other technologies (0.074% for the full sample). We should remain cognizant of the fact that patent counts will tend to emphasize technological changes, and that other measures may provide a better understanding of the System’s impact on other aspects of innovation, such as development of new operational strategies.

To test whether our focus on EU ETS firms would have blinkered us to the System’s

broader effects, we have also attempted to estimate the indirect impact of the EU ETS. To this end, we have compared non-EU ETS firms with at least one patent jointly filed with an EU ETS firm, with otherwise similar non-EU ETS firms. Although we can only provide indicative estimates, we find no compelling evidence that the EU ETS has had either a net positive or net negative impact on the patent filings of potential technology providers and demanders. If data on patent licensing agreements could be obtained, researcher may in the future be able to study questions like this in greater detail.

Our findings suggest a way to reconcile the findings of the broader empirical literature on environmental policy and directed technological change. Several studies of the impacts of inclusive standards and energy or pollution taxes find evidence that environmental policy does indeed encourage directed technological change (Lanjouw and Mody, 1996; Brunnermeier and Cohen, 2003; Popp, 2002, 2003, 2006; Arimura et al., 2007; Lanoie et al., 2007). In contrast, studies of previous emissions trading programs, like the US Acid Rain Program, at best unearth evidence of very small impacts on directed technological change (Popp, 2003; Lange and Bellas, 2005). Our results indicate that the discrepancy between the findings of cap-and-trade studies and studies of other instruments may be a consequence, not of weaker innovation incentives provided by emissions trading instruments, but of the fact that they tend to concern a comparatively small number of firms. The impact on these firms may in fact be quite large, even in the EU ETS where permits in the initial trading phases were very likely over-allocated. When their response is compared to the overall pace of technological change, however, the effect appears negligible. Our estimates at the aggregate level are consistent with the weak effects found in the empirical literature on cap-and-trade programs, but our firm-level estimates provide additional detail. The weak aggregate effect is an average of the non-reaction of a large number of firms that are more or less unaffected by the program, and the strong reaction of a small group of regulated firms. Someone studying the impact of an emissions trading program by looking only at patenting records at a more aggregated level is effectively pooling together these two groups of firms, and is therefore likely to overlook the System's strong but targeted effect. Conversely, the impact of more inclusive environmental policies, like energy and pollution taxes, may be more easily detected because these policies affect so many firms, even if the change in behaviour for each firm is quite small. Debates about the relative costs and benefits of different environmental policy instruments already consider the impacts on pace and direction technological change of central importance (Kneese and Schultze, 1975; Pizer and Popp, 2008). Our results, read in combination with the findings of the broader literature, suggest that environmental policy instruments may differ also in the distribu-

tion of impacts on directed technological change. This could be potentially significant because of the positive spill-overs usually associated with innovation. It is an interesting question for future research, therefore, whether this could change the economic, or indeed the political calculus of instrument choice for environmental policy.

Our aim has been to estimate the overall impact of the EU ETS on directed technological change. However, once we have estimated that each EU ETS firm filed 2 additional low-carbon patents, it is a small step to consider what types of technologies those patents protect, and which economic sectors the patent holders are active in (i.e. conditional on the estimated treatment effect). When we look at how these additional patents break down across technologies, we find that most of them appear to protect alternative energy and energy storage, with the remaining ones focusing on energy efficiency. Alternative energy technologies appear to account for a greater number of additional low-carbon patents than do improvements of conventional combustion technologies. Broken down by economic sector, most of the additional low-carbon patents belong to chemicals manufacturers, energy companies, and automobile manufacturers (see appendix D for an explanation of the methodology as well as for the technology-level and sector-level estimates). These preliminary conclusions are of course based on conditional estimates, and future research may give us a more granular picture of the impact of the EU ETS.

There are many questions, too, that we have not answered in this paper. For instance, would we have observed a greater innovation impact if the price of permits had been higher? Or if the permits had been auctioned instead of allocated for free? Or if there had been less uncertainty about the policy? Given the lack of variation in EU ETS rules so far, it has not been feasible to construct the counterfactual scenarios needed to test these hypotheses—an EU ETS with different prices, with different allocation rules, etc. The impact observed until now of the *de facto* EU ETS on low-carbon technological change is consistent with a number of alternative hypotheses about the impacts of specific future reforms. Future changes to the rules may provide opportunities to study the impacts of such reforms.

In focusing on the EU ETS, moreover, we have not identified what has caused the post-2005 surge in low-carbon patenting in Europe. The number of low-carbon patents filed in Europe has risen rapidly in recent years, and it would be an interesting exploratory exercise to search for the other factors that have contributed to this development (e.g. renewable energy policies). At present, we can establish only that the EU ETS seems to have played no more than a very limited part.

In focusing on the EU ETS, moreover, we have not identified what has caused the

post-2005 surge in low-carbon patenting in Europe. The number of low-carbon patents filed in Europe has risen rapidly in recent years. Our estimates imply that the EU ETS accounts for only about 2% of the post-2005 surge. It would be an interesting exploratory exercise to search for the other factors that have contributed to this development (e.g. renewable energy policies), but at present, we can only establish that the EU ETS seems to have played no more than a very limited part.

Our results also have broader policy implications. The EU ETS forms an integral part of the European Union's roadmap to a low-carbon economy in 2050 (European Commission, 2011). Policy makers in New Zealand, the United States, Australia, China, Japan, South Korea, and elsewhere, can also learn from the EU ETS experience. So far, it appears that emissions reductions in the EU ETS have come largely from operational changes like fuel switching rather than technological changes, much like in past emissions trading programs. Such abatement strategies will not be enough to reach the EU's ambitious longer term targets, however. New low-carbon technologies are needed. Indeed, our results indicate that EU ETS regulated firms are cognizant of this fact, and are responding accordingly. Even so, because the impact of emissions trading appears to be concentrated to a relatively small group of firms, their response looks to nearly vanish when considered in relation to the overall pace and direction of technological change. For this reason, the System in its current form might not be providing the economy-wide incentives necessary to bring about low-carbon technological change on a larger scale.

References

- Abadie, A. (2005). Semiparametric difference-in-differences estimators. *The Review of Economic Studies*, 72(1):1–19.
- 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. (2012). Carbon Taxes, Path Dependency and Directed Technical Change: Evidence from the Auto Industry. *NBER Working Paper Series*, (18596).
- 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. and Szambelan, S. (2009). US emissions trading markets for SO₂ and NO_x. *Resources for the Future Discussion Paper*, (09-40).
- Chay, K. Y. and Powell, J. L. (2001). Semiparametric censored regression models. *Journal of Economic Perspectives*, 15(4):29–42.
- 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.
- Dechezleprêtre, A., Glachant, M., Hašič, 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.
- Dehejia, R. and Wahba, S. (1999). Causal effects in nonexperimental studies: Reevaluating the evaluation of training programs. *Journal of the American Statistical Association*, 94(448):1053–1062.
- 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.

- 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.
- Ellerman, A. D., Convery, F. J., and de Perthuis, C. (2010). *Pricing Carbon: The European Union Emissions Trading Scheme*. Cambridge University Press.
- 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 (2011). A Roadmap for moving to a competitive low carbon economy in 2050. Technical Report COM(2011) 112, European Union.
- European Commission (2012). Emissions trading: annual compliance round-up shows declining emissions in 2011. Press Release IP/12/477.
- 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. (2010). Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement. *American Economic Review*, 100(3):837–869.
- 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.
- Greenstone, M. and Gayer, T. (2009). Quasi-experimental and experimental approaches to environmental economics. *Journal of Environmental Economics and Management*, 57(1):21–44.
- Griliches, Z. (1984). *R & D, patents, and productivity*. University Of Chicago Press.
- 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.
- Heckman, J., Ichimura, H., Smith, J., and Todd, P. (1998a). Characterizing Selection Bias Using Experimental Data. *Econometrica*, 66(5):1017–1098.
- Heckman, J. J., Ichimura, H., and Todd, P. (1998b). Matching as an econometric evaluation estimator. *The Review of Economic Studies*, 65(2):261–294.
- 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.
- Hoppmann, J., Peters, M., Schneider, M., and Hoffmann, V. H. (2013). The two faces of market support—How deployment policies affect technological exploration and exploitation in the solar photovoltaic industry. *Research Policy*, 42(4):989–1003.
- 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.
- 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.
- Kaufer, E. (1989). *The economics of the patent system*. Routledge.
- Kneese, A. V. and Schultze, C. (1975). Pollution, Prices, and Public Policy. (*Brookings Institution, Washington, DC*).

- Kossoy, A. and Guigon, P. (2012). State and trends of the carbon market 2012. Annual report, World Bank.
- Kossoy, A., Oppermann, K., Reddy, R. C., Bosi, M., Boukerche, S., Höhne, N., Klein, N., Gilbert, A., Jung, M., Borkent, B., Lam, L., Röser, F., Braun, N., Hänsel, G., and Warnecke, C. (2013). Mapping Carbon Pricing Initiatives: Developments and Prospects. Technical report, World Bank and Ecofys, 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*.
- 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.
- Manski, C. F. (2007). *Identification for Prediction and Decision*. Harvard University Press.
- Markussen, P. and Svendsen, G. (2005). Industry lobbying and the political economy of ghg trade in the european union. *Energy Policy*.
- 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*.
- 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.
- Nemet, G. (2009). Demand-pull, technology-push, and government-led incentives for non-incremental technical change. *Research Policy*, 38(5):700–709.
- 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.
- 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. *Environmental Defense*, pages 1–23.
- Pizer, W. A. and Popp, D. (2008). Endogenizing technological change: Matching empirical evidence to modeling needs. *Energy Economics*, 30(6):2754–2770.
- 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: Endogenous Technological Change in the DICE Model of Global Warming. *Journal of Environmental Economics and Management*, 24(1):742–768.
- 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. (2010). Innovation and climate policy. *NBER Working Paper*.
- 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.
- 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.
- Rosenbaum, P. R. (2009). *Design of Observational Studies*. Springer.

- 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.
- 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.
- Smith, J. and Todd, P. (2005). Does matching overcome LaLonde’s critique of nonexperimental estimators? *Journal of econometrics*, 125(1):305–353.
- Stavins, R. (2007). A US cap-and-trade system to address global climate change. *Regulatory Policy Program Working Paper*, (RPP-2007-04).
- Taylor, M. R. (2012). Innovation under cap-and-trade programs. *Proceedings of the National Academy of Sciences*, 109(13):4804–4809. PMID: 22411797.
- 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.
- 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.

A Data

For 8 of the countries in our sample, the company registration numbers of the installation operators were obtained directly, either from national emissions trading registries or from the Community Independent Transactions Log (CITL) (the EU body to which national registries report). For the remaining 13 countries in our data set that participated in the 2005 launch of the EU ETS, a combination of exact and approximate text matching methods were used to establish a link between firm data and regulatory data. This was complemented by further manual searches, and extensive manual double-checking.

The firm data set allows us to identify majority ownership. Using this information, we excluded non-EU ETS firms that were owner, sister company, or subsidiary to an EU ETS firm. This reduces the chance of matching two potentially dependent observations.

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.

Firms have been matched so that each matched group operates in the same country and economic sector (defined at the 3-digit level of NACE Rev. 2 sector classification codes). The firms are also matched on the basis of average turnover in the period 2000–2004, the number of low-carbon patents and other patents filed that same period, and year of incorporation (to measure age). To improve covariate balance, the matches were also penalized for dissimilarities in the square of turnover, an indicator variable noting whether or not firms had filed any patents prior to 2005, another indicator variable noting whether or not firms had filed any low-carbon patents prior to 2005, and in the overall and the low-carbon patent counts interacted with the year of incorporation. Finally, calipers were applied to ensure that no matched groups were too dissimilar in terms of overall and low-carbon patent counts in the period 2000–2004.

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, detailed in Veefkind et al. (2012). Table 7, adapted from Veefkind et al. (2012), lists the main patent classes along with some examples of technologies for each class:

Table 7: Climate change mitigation patent categories (EPO's Y02 class)

Patent code	Description	Example technologies
Y02C 10/00	CO_2 capture or storage	Chemical or biological separation, ad- or absorption, membrane technology, condensation etc.; subterranean or submarine storage
Y02C 20/00	Capture or disposal of greenhouse gases other than CO_2	N_2O , methane, perfluorocarbons, hydrofluorocarbons or sulfur hexafluoride
Y02E 10/00	Energy generation through renewable energy sources	Geothermal, hydro, oceanic, solar (photovoltaic and thermal), wind
Y02E 20/00	Combustion technologies with mitigation potential	Combined Heat and Power (CHP), Combined Cycle Power Plant (CCPP), Integrated Gasification Combined Cycle (IGCC), synair, oxyfuel combustion, cold flame, etc.
Y02E 30/00	Energy generation of nuclear origin	Fusion and fission
Y02E 40/00	Technologies for efficient electrical power generation, transmission or distribution	Reactive power compensation, efficient operation of power networks, etc.
Y02E 50/00	Technologies for the production of fuel of non-fossil origin	Biofuels, from waste
Y02E 60/00	Technologies with potential or indirect contribution to greenhouse gas (GHG) emissions mitigation	Energy storage (batteries, ultracapacitors, flywheels.), hydrogen technology, fuel cells, etc.
Y02E 70/00	Other energy conversion or management systems reducing GHG emissions	Synergies among renewable energies, fuel cells and energy storage

The full list of low-carbon patent classes includes:

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

Additional patent classes for “extended” low-carbon patents definition include:

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

D Additional patents by technology and sector

The large number of zeros in the patent data unfortunately prevents us from obtaining informative estimates of the treatment effect at the level of individual technologies and economic sectors, without additional identifying assumptions. Perhaps the simplest assumption we can make at this point is that the EU ETS does in fact account for 2 additional low-carbon patents for each EU ETS firms (i.e. to condition on the estimated treatment effect). What would this imply for specific technologies and sectors?

Let us first look at the distribution of additional patents across technologies. An individual often holds several patents protecting different technologies, and even if we assume that 2 of the patents are additional, there is no definite way of identifying which. Instead, we adopt a probabilistic approach. Imagine randomly selecting 2 low-carbon patents from each firm (or as many as they have filed, if fewer), and then simply

counting the number of patents belonging to each patent class. If we repeat this exercise again and again, until we have selected every subset of 2 (or fewer) patents from every firm in combination with every subset of 2 (or fewer) patents from every other firm, we ultimately obtain the whole conditional distribution of additional patents for every technology.

Table 8 reports the minimum, mean, and maximum of these technology-specific distributions, organized by technology group, roughly in order of number of additional patents, from largest to smallest (and excluding all technologies with zero additional patents). Besides the means it is worth also keeping an eye on the minima, because if a firm filed some low-carbon patents protecting technologies seemingly unrelated to the EU ETS regulations, our method of estimating conditional distributions will assign a positive probability those patents being additional also. In most of these cases one would expect the minimum to be zero. Note also that some patents are tagged with multiple codes, which results in a small amount of double counting (so that the sum of means slightly exceeds the total number of additional patents). Most of this double counting appears in form of a handful of additional patents in categories seemingly unrelated to abatement of carbon dioxide emissions, or in categories outside of the Y02 class. Double counting also potentially raises the minimum for these technologies.

We provide these numbers without too much discussion or interpretation, only noting that most of the additional low-carbon patents appear to protect alternative energy and energy storage, followed by carbon sequestration and storage, with the remaining additional patents focusing primarily on energy efficiency. We may also notice that alternative energy technologies appear to account for more additional low-carbon patents than do conventional combustion technologies.

We can repeat the same thought experiment to obtain the number of additional low-carbon patents for each economic sector. Patents are assigned to the economic sector of the patent holder (i.e. the firm, not the installation), so this exercise should provide an indication of the main activities of the patenting firms. We should be aware, therefore, that the economic sector of a patent may in principle be different from the type of technology it protects, or from the type of activity for which the firm is regulated under the EU ETS (e.g. a generator operated by a chemicals manufacturer may be regulated as a ‘combustion’ installation, while the chemicals manufacturer may file patents to protect anything from its energy efficiency innovations to new alkaline or acid solution used in batteries). A firm’s patents are all assigned to a unique economic sector, which means the sector-level distributions will be degenerate. Table 9 presents the conditional estimates for all economic sectors with at least one additional patent represented among our EU

ETS firms (sectors are here defined in terms of NACE Rev. 2 codes, and aggregated to the 2 digit level). Since information about economic sector was missing for a few firms, the sum of additional patents across sectors is slightly smaller than the total number of additional patents.

Nearly half of all additional low-carbon patents were added by chemicals manufacturers, energy companies, and automobile manufacturers. Technologies relating to energy and transportation are relatively easy to identify in table 8 as well, but chemicals is perhaps a bigger surprise. Although the EU ETS regulates rather few chemicals installations, and although the firms operating those installations filed very few patents, in practice the EU ETS regulates several chemicals manufacturers for other activities, and these firms filed a fair number of low-carbon patents.

E Details of other robustness tests

Are matched non-EU ETS firms also responding to EU ETS? The matched firms that are not regulated by the EU ETS may nevertheless respond to it, either directly, or indirectly because they engage in competition with EU ETS firms. This would bias our estimates. If very similar unregulated firms are responding by innovating more, a comparison of EU ETS firms and matched non-EU ETS firms will under-estimate the impact of the EU ETS. If very similar unregulated firms are responding by innovating less, this comparison will over-estimate the impact of the EU ETS. To examine these possibilities we have re-matched our EU ETS firms to companies operating in European countries that did not participate in the 2005 launch of the EU ETS (Norway, Switzerland, Romania, and Bulgaria), and then separately to US companies. These comparisons are less likely to suffer from this kind of bias, because the matched non-EU ETS firms are less exposed to the market created by the EU ETS and less likely to be directly engaged in competition with EU ETS companies.²⁶

Table 10 reports the estimated treatment effects for both the European and US re-matched samples, along with our original estimates for comparison. The re-matched point estimates are smaller than our original estimate (and both insignificantly different from zero), which would tend to indicate that very similar unregulated firms in EU ETS countries perhaps are innovating less than they would have without the EU ETS. Our original estimate, then, may if anything have overestimated the impact of the EU ETS. Due to between-country differences, however, which these re-matched estimates cannot

²⁶While 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.

control for, one should exercise caution in recommending such an interpretation. Neither of the re-matched estimates differ significantly from our original estimate, and as such do not seem to offer a substantive challenge to our findings.

Is the result an artifact of how we measure low-carbon patents? It is possible that our finding is an artifact of our particular measure of low-carbon technological change. If we compare our matched EU ETS and non-EU ETS firms using expanded definition of “low-carbon technologies”, the result does not appear to change materially (see table 11). Our original estimate was that the EU ETS accounts for a 36.2% increase in low-carbon patenting among matched EU ETS firms, a 8.1% increase across our full sample of EU ETS firms, and no more than a 1% increase across our study area. The new treatment effect estimates suggest the EU ETS may have increased low-carbon patenting among matched EU ETS firms by 32.4%, a 7.1% increase across our full sample, and no more than a 1% increase across our study area. The new numbers are well within our original confidence intervals, and do not appear to present a challenge for our interpretation of the results. Our findings therefore appear robust to how the outcome is defined.

A related concern is that patent counts would omit any EU ETS response that appears in the form of a change in the quality of patents. For instance, one might hypothesize that the EU ETS invest in increasing the quality of their patents, not just the number. Alternatively, Nemet (2009) and Hoppmann et al. (2013) raise the concern that a ‘demand-pull’ or ‘deployment’ policy, like the EU ETS, might hinder non-incremental innovation (which would likely be counted in the form of high value patents). If this were the case, we would expect the patent quality of EU ETS firms to deteriorate relative to their non-EU ETS counterparts. We test whether the EU ETS has systematically changed the quality of low-carbon patents filed by EU ETS relative to non-EU ETS firms, as measured by citations and family size. Our results are reported in table 12. Our estimates suggest that EU ETS firms typically have received 2.75 additional low-carbon patent citations in 2005–2009, which roughly means that each of their 2 additional patents received just over one citation. The estimate is insignificantly different from zero, however. The family size of the patent portfolios of EU ETS firms, which is expected to respond quicker to changes in regulation than citations, increased by 11.75 relative to non-EU ETS firms. Since EPO patents are filed in 6 countries on average, this estimate can be roughly interpreted as saying that the 2 additional low-carbon patents filed by EU ETS firms are of average quality. In sum, our estimates suggest that the EU ETS has not had an impact on the quality of patents.

Are there other relevant omitted variables? An omitted variable can bias our estimates if it is correlated with both the treatment (EU ETS or non-EU ETS) and the outcome (the change in patent filings from 2000-2004 to 2005-2009). Section 4.3 already considers a few specific instances of omitted variables, as well as investigates the sensitivity of our findings with respect to generic omitted variables. To better gauge whether there are other examples of omitted variables that might compromise our results, we look at a few more candidates here to see whether, even though they were not explicitly matched on, matching nevertheless achieved a reasonable balance on these variables.

First, consider the fact that firms which qualify must have at least one sufficiently large installation. For a given overall firm size, one might then expect that the activities of EU ETS firms were more concentrated among fewer installations, relative to non-EU ETS firms. We also know that something else is happening in the economy after 2005 that explains much of the surge in innovation—maybe simply a drop in the cost of low-carbon innovations. If it is easier for the more concentrated firms to adapt their research efforts—perhaps their R&D department is located at a single installation, say—this would create a systematic bias in our estimates. It is possible to address this concern by counting the number of locations of innovators for each firm, and then testing whether the distributions differ substantially between matched EU ETS and non-EU ETS firms.

Table 13 reports the result from an equivalence test, following the same procedure as before. Matching appears to have achieved a reasonably balanced set of firms in terms of the number of innovation locations. It is worth noting, though, that there is a lot of idiosyncratic spelling of innovator addresses in the patent database. This creates a great deal of measurement error in the location counts, which is one of the reasons why this variable was not used to match on in the first instance (matching on noise reduces the quality of matching estimates).

Second, suppose growing firms install extra capacity, to meet expected future demand, while shrinking firms get rid of their excess capacity. Growing firms might then be more likely to become regulated under the EU ETS. Suppose further that growing firms react more strongly to the EU ETS with low-carbon innovation. In combination, these two suppositions imply that company growth, even conditional on the *level* of turnover in the years prior to the EU ETS (which we have already matched on), might be correlated with both treatment assignment and outcomes. The first of these necessary conditions can be directly assessed by testing whether the distributions of company growth differ substantially for matched EU ETS and non-EU ETS firms.

While theoretically straightforward, this test presents a practical challenge. The

growth rate of turnover is likely to be measured with a greater degree of error than the *level* of turnover, since there is always one less observation to estimate the growth rate than to estimate the mean, and since a particularly low turnover in a given year—whether real, an accounting artifact, or a database error—is likely to give rise to a hugely inflated growth rate measured in the subsequent year. The amount of missing financial data before 2005, as discussed earlier, makes this a very real concern. The greater sensitivity of growth rates to measurement error will tend to produce a highly dispersed distribution. In principle, a single small value followed by a normal value for a single firm could vastly exaggerate both the mean and standard errors of the distribution. This is one of the reasons why this variable was not used to match on in the first place.

With these caveats in mind, table 13 reports the results an equivalence test on average annual growth in turnover in the pre-EU ETS period. As expected, the presence of a few inflated growth rates results in a pretty meaningless equivalence range, but the other two columns are perhaps more informative. The first column shows that the typical difference between matched EU ETS and non-EU ETS firms is less than $\pm 1\%$, and in fact, more non-EU ETS firms grew faster than their matched EU ETS firms than vice versa. In the last column, we see that we are able to reject (at the 5% significance level) hypotheses that the distributions of turnover growth for EU ETS and non-EU ETS firms differ by a shift of more than 1.81%. These findings suggest that our matched sample is fairly balanced with respect to pre-2005 company growth rates, and consequently, it is very unlikely that omission of pre-2005 company growth has produced substantial bias in our estimates.

Table 8: Additional low-carbon patents by technology (conditional estimates)

Technology	Patent code	Min	Mean	Max
Enabling technologies				
Energy storage	Y02E 60/10	19	33.31	60
Fuel cells	Y02E 60/50	3	18.96	45
Hydrogen technology	Y02 E60/30	3	6.29	25
Non-fossil fuel production				
Biofuels	Y02E 50/10	18	24.98	38
Fuel from waste	Y02E 50/30	10	12.73	22
Renewable energy				
Solar PV	Y02E 10/50	10	18.60	40
Solar thermal	Y02E 10/40	7	11.52	22
Wind	Y02E 10/70	4	9.03	19
Sea	Y02E 10/30	1	2.46	7
Hydro	Y02E 10/20	0	1.25	7
Geothermal	Y02E 10/10	1	1.00	1
Thermal-PV hybrid	Y02E 10/60	0	0.07	2
Combustion technologies with mitigation potential				
Combined combustion	Y02E 20/10	12	17.72	37
Efficient combustion or heat usage	Y02E 20/30	5	5.89	13
Carbon capture, storage, sequestration, disposal				
Carbon capture and storage	Y02C 10/00	7	9.95	24
Other CCS	Y02C 10/10	3	4.02	11
Other GHG capture and disposal				
Nitrous oxide	Y02C 20/10	3	4.70	14
Methane	Y02C 20/20	0	0.14	5
PFC, HFC, SF6	Y02C20:3	0	0.09	3
Efficient end-user electric power management and consumption				
Efficient power electronics conversion	Y02B 70/10	2	3.76	13
Transportation				
Fuel cell applications	Y02T 90/30	0	2.03	11
Energy harvesting for auxiliary power supply	Y02T 10/90	0	0.18	5
Nuclear energy				
Other fission	Y02E 30/40	2	2.00	3
Fission reactors	Y02E 30/30	0	1.20	2
Fusion reactors	Y02E 30/10	0	0.04	1
Efficient electrical power generation, transmission, and distribution				
Smart grids / System integration	Y02E 40/70	1	1.12	2
Reactive power compensation	Y02E 40/30	1	1.04	4
Superconductive systems	Y02E 40/60	0	0.21	3
Active power filtering	Y02E 40/20	0	0.03	2
Harmonics reduction	Y02E 40/40	0	0.02	3
Flexible AC transmission	Y02E 40/10	0	0.01	2
Polyphase network asymmetry reduction	Y02E 40/50	0	0.01	2
Efficient heating, ventilation, and air conditioning				
Control and regulation	Y02B 30/70	0	1.00	2
Boilers	Y02B 30/10	0	0.10	2
Other heating and cooling	Y02B 30/60	0	0.09	3
Home appliances efficiency				
Efficient batteries, ultracapacitors, supercapacitors or double-layer capacitors charging or discharging systems or methods specially adapted for portable applications	Y02B 40/90	0	0.68	3
Buildings				
Fuel cells applications	Y02B 90/10	0	0.67	2
Power network integration				
End-user control systems	Y04S 20/20	0	0.25	1
Power network elements and equipment	Y04S 10/10	0	0.02	2
Communication technology	Y04S 40/10	0	0.01	2
Electric or hybrid vehicle interoperability systems	Y04S 30/10	0	0.01	2
Energy conversion and management systems				
Combining non-fossil energy generation with energy storage	Y02E 70/30	0	0.02	2
Combining non-fossil energy generation with hydrogen electrolysis	Y02E 70/10	0	0.00	1

Table 9: Additional low-carbon patents by sector (conditional estimates)

Economic sector	NACE Rev. 2	Additional low-carbon patents
Chemicals	20	29
Electricity, gas, steam and air conditioning supply	35	25
Motor vehicles	29	18
Glass, ceramics, and cement	23	16
Computer, electronic, and optical products	26	9
Transport equipment (except motor vehicles)	30	9
Machinery and equipment (Engines, turbines, etc.)	28	9
Paper	17	7
Fabricated metal products (except machinery and equipment)	25	6
Iron and steel	24	4
Electrical equipment	27	4
Scientific research and development	72	4
Refined petroleum products	19	4
Pharmaceuticals	21	4
Food products	10	4
Wood products	16	3
Crude petroleum extraction	06	2
Engineering activities and related technical consultancy	71	2
Wholesale trade	46	2
Activities of holding companies	64	2
Land transport and transport via pipelines	49	1

Table 10: Treatment effect estimates using ‘distant’ matches

	Treatment effect
Norway, Switzerland, Romania, and Bulgaria	1 (0, 1.99)
USA	-1 (-1.99, 0.99)
Original estimate	2 (1, 5)

Table 11: Estimates with different definitions of “low-carbon technologies”

	Additional low-carbon patents			
	Matched sample		Full sample	
	As % increase	As % increase of EPO	As % increase	As % increase of EPO
Extended definition	32.4 (20.3, 62.5)	0.34 (0.24, 0.54)	7.1 (4.5, 12.3)	0.77 (0.50, 1.28)
Standard EPO definition	36.2 (20.2, 69.0)	0.38 (0.24, 0.58)	8.1 (4.7, 14.5)	0.85 (0.51, 1.45)

Table 12: Changes in quality of low-carbon patents

	Treatment effect
Citations	2.75 (0, 17.99)
Family size	11.5 (4, 35)

Table 13: Equivalence tests for matched EU ETS and non-EU ETS firms on omitted variables

	Median difference between EU ETS and non-EU ETS firms	Equivalence range	Critical equivalence range (5% sign. lev.)
Number of innovation locations	0	± 6.42	± 1.99
Turnover growth (% p.a.)	-0.69	± 21507.37	± 1.81