Raphael Calel and Antoine Dechezlepretre

Environmental policy and directed technological change: evidence from the European carbon market

Article (Published version) (Refereed)

Original citation:
Calel, Raphael and Dechezlepretre, Antoine (2016) Environmental policy and directed technological change: evidence from the European carbon market. Review of Economics and Statistics, 98 (1). pp. 173-191. ISSN 0034-6535
DOI: 10.1162/REST_a_00470

ENVIRONMENTAL POLICY AND DIRECTED TECHNOLOGICAL CHANGE: EVIDENCE FROM THE EUROPEAN CARBON MARKET

Raphael Calel and Antoine Dechezleprêtre*

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 affected 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.

I. Introduction

EMISSIONS trading programs have assumed an ever more prominent role in environmental policy over the past few decades. In the United States, 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 in 2017. 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 & 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 System (EU ETS), today the 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 other 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. The primary objective of carbon market programs is of course to reduce emissions, but from an economic perspective, it is crucial that they also provide incentives for technological change, since new technologies may substantially reduce the long-run cost of abatement (Jaffe, Newell, & Stavins, 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 European Commission, 2005, 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 five 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 firm level, 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 & 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 our

Received for publication February 19, 2013. Revision accepted for publication March 27, 2014. Editor: Philippe Aghion.

* Calel: Georgetown University, Berkeley, and London School of Economics; Dechezleprêtre: Grantham Research Institute and Centre for Economic Performance, London School of Economics.

We thank Philippe Aghion and two anonymous reviewers for their constructive suggestions. For their insightful comments and generous advice, we also owe 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. R.C. is grateful for funding provided by the ESRC, the Jan Wallander and Tom Hedelius Foundation, and the Ciriacy-Wantrup Foundation. A.D. gratefully acknowledges the support of ADEME, the Global Green Growth Institute, and the ESRC under the ESRC Postdoctoral Fellowship Scheme (award PTA-026-27-2756). We also acknowledge financial support from the Grantham Foundation for the Protection of the Environment. We thank Xavier Vollenweider for excellent research assistance.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00470.

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

The Review of Economics and Statistics, March 2016, 98(1): 173–191

doi:10.1162/REST_a_00470

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 (e.g., input prices, sector- and country-specific policies), as well as firm-level heterogeneity (Heckman, Ichimura, Smith et al. 1998; Heckman, Ichimura, & Todd, 1998; 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 of 3,428 EU ETS firms, or an increase of 9.1% across all of the 5,500 EU ETS firms. Because these firms account for only 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 the development of other technologies (Popp & 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 likely third-party technology providers and purchasers and test whether these firms 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 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 & Schultze, 1975; Pizer & Popp, 2008). In light of this, it is not surprising that both theoretical and empirical economists are engaged in ongoing efforts 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 & 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 & Palmer, 1997; Newell, Jaffe, & Stavins, 1999; Brunnermeier & Cohen, 2003; Popp, 2002; Aghion et al., 2012; and many others; see Popp, Newell, & Jaffe, 2010, Popp, 2010, and Ambec et al., 2013, for recent surveys). Our study contributes to this literature and analyzes 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, Convery, & Maria, 2011). Martin, Muûls, and Wagner (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 II surveys the evidence on environmental policy and directed technological change, especially in the context of emissions trading. Evidence from the U.S. Acid Rain Program and early studies of the EU ETS inform us about how the EU ETS is likely to have affected technological change. In section III, we familiarize ourselves with our newly constructed data set and use it to begin unpacking the characteristics of low-carbon technological change. In section IV, we turn our eye to estimating the impact of the EU ETS on regulated firms, and in section V we examine its indirect impact on third-party technology providers and purchasers. Section VI 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.

II. Emissions Trading and Directed Technological Change

A. Empirical Background

Several studies have found evidence that environmental policy does have an impact on the direction of technological change (Lanjouw & Mody, 1996; Brunnermeier & Cohen, 2003; Popp, 2002, 2003, 2006; Arimura, Hibiki, & Johnstone, 2007; Lanoie, et al., 2007; Johnstone, Haščič, & Popp, 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.

First, the impact of emissions trading programs specifically, rather than environmental policies more broadly construed, appears to be more modest. Most studies consider the Acid Rain Program, which in 1995 replaced the traditional regulatory regime for sulfur dioxide emissions from U.S. power plants. Patenting for sulfur dioxide control technologies began a precipitous decline after 1995 (Taylor, 2012), although there was an increase in patents that improve the efficiency of sulfur scrubbers (Popp, 2003). The latter effect was confined to early years of the new regime, though, and the program has not provided ongoing incentives for technological advancement (Lange & 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 sulfur content (Schmalensee et al., 1998), but the use of scrubber technology as an abatement strategy has declined over time (Burtraw & 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 that EU policymakers are hoping the EU ETS will provide.

Second, 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 a 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 eventually boosted the share of energy-saving patents by only 3.14%. The carbon price in the EU ETS, having ranged from a peak of near 30 euros to a low of near 0 euros (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.

B. 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, and so on. Activity-specific size criteria then determined 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, fewer and fewer tradable emissions permits are allocated to the more than 12,000 qualifying installations, which are each legally required to surrender enough permits every year to cover their 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 euros and 30 euros. The average price between 2005 and 2009 was around 10 euros, although the actual price spent more time closer to 0 euros. 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 the EU ETS launched in 2005, there has been vigorous debate about whether it 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 & Betz, 2005; Gagelmann & Frondel, 2005; Grubb, Azar, & Persson, 2005). So far, fuel switching appears to have been very important. Fuel switching is a purely organizational innovation and requires neither capital investment nor

[2] It is worth noting also that Title IV of the Clean Air Act, which established the Acid Rain Program, 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.

[3] Popp (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 kilowatt-hour in 2005, and industrial energy users paid about 0.07 euros per kilowatt-hour that year. If the average carbon price of 10 euros was entirely passed onto users, that would imply about a 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 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%.

[4] The system has been implemented in three trading phases, with successively more stringent emissions caps for each phase. Phase 1, which ran from 2005 to 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, Convery, and de Perthuis (2010) for a more comprehensive review of the design and implementation of the EU ETS.

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.[5] Macroeconomic estimates suggest that the EU ETS reduced total emissions by roughly 50 million to 100 million tonnes of carbon dioxide annually in phase 1, or roughly 3% to 6%, compared with a business-as-usual scenario (Ellerman & Buchner, 2008; Anderson & Di Maria, 2011). Meanwhile, model-based estimates of power sector emissions abatement from fuel switching range from 26 million to 88 million tonnes per year (Delarue, Ellerman, & D'haeseleer, 2010; Delarue, Voorspools, & D'haeseleer, 2008), which suggests that fuel switching 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 U.S. 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 emissions by up to 300 million tonnes annually, which is no more than one-tenth of what is needed to meet the EU target to cut emissions by 80% by 2050 against 1990 levels.[6]

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 & 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, finds 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 offer 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, and it is difficult to summarize our expectations of the EU ETS's 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 weak or absent impact, and studies of the EU ETS generally have been unable to detect an effect thus far. Our 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.

III. 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 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). For instance, not all innovations are patentable, and even when one is, patenting is only one of several ways to protect it. The propensity to file patents and the economic value of patents consequently differ between sectors. However, there are very few examples of economically significant inventions that have not been patented (Dernis, Guellec, & Pottelsberghe, 2001), and the production of patented knowledge and tacit knowledge has been found to be positively correlated (Cohen, Nelson, & Walsh, 2000; Arora, Ceccagnoli, & Cohen, 2008). Moreover, it is possible to mitigate the deficiencies in patent-based measures by comparing companies active in the same sector and focusing on higher-value patents. In sum, patent-based measures 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. 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,

[5] In other contexts, "fuel switching" may refer to structural and technological changes over long time horizons, such as the global shift from biomass to fossil fuels as the dominant energy carrier over the past two centuries. Throughout, we use the term more narrowly to refer to the short-run operational shift between coal and gas.

[6] The 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.

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.[7] Nevertheless, as a robustness test we also repeat our analysis using quality-weighted patent counts.[8]

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" ("low-carbon technologies," for short). This new category (the Y02 class) is the result of an unprecedented effort by the EPO 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.[9] 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.[10] (A complete list of the subclasses of low-carbon patents used in the paper can be found in appendix C in the online appendix.)

The EPO was set up in 1978. Since then, over 2.5 million patents have been filed with it, 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 (21 EU countries, Switzerland, and the United States). Eighteen of these countries launched the EU ETS in 2005. The other five (Norway, Switzerland, Romania, Bulgaria, and the United States) either joined later or have remained outside the EU ETS altogether. While our data are 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.[11]

Figure 1.—Share of Low-Carbon Patents, 1978–2009

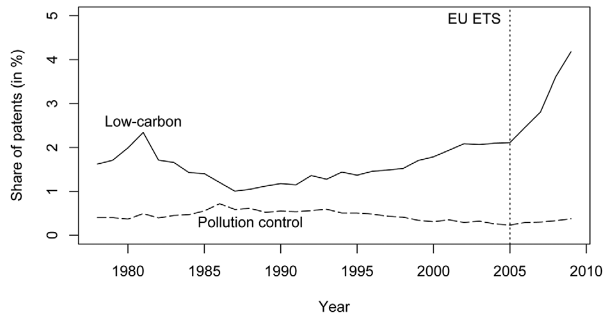

The share of patents protecting low-carbon technologies shows a distinct pattern over time (figure 1). There was a surge in patenting for these technologies in the early 1980s, often attributed to the second oil price shock in the late 1970s (Dechezleprêtre et al., 2011). The share of low-carbon patents 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 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 ($p < 0.001$).

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 nongreenhouse gas "pollution control technologies," as defined by Popp (2006), which does not display the same structural break (one cannot reject the hypothesis of no structural break in 2005 at conventional significance levels).[12] 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 as 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

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

[8] Although 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, Jaffe, & Trajtenberg, 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 have 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 when dealing with recent patents, as we do here.

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

[10] We also test the robustness of our results to the inclusion of additional patents that other authors have considered low carbon, in particular patents pertaining to energy-efficient industrial processes. 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.

[11] We 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 have high administrative costs.

[12] These technologies pertain to reduction of local pollutants including sulfur dioxide and nitrogen oxide.

FIGURE 2.—SHARE OF LOW-CARBON PATENTS AND CRUDE OIL PRICES, 1978–2009

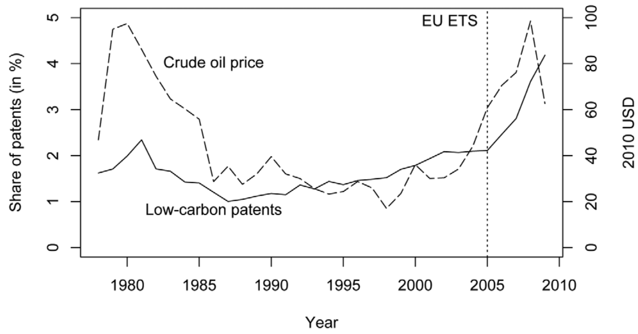

(see figure 2), one notices that the recent surge in patenting follows immediately on the heels of rapid oil price increases in the early 2000s. Patenting for pollution control, however, was not responsive to the oil price in the 1980s, and so it is not surprising it has stayed flat recently. Clearly, looking at the aggregate trends over time is not enough to determine whether the increase in low-carbon patenting since 2005 is the result of the EU ETS, oil prices, or some other factor.

In order to isolate the impact of the EU ETS, we can try comparing the experience of firms regulated under the EU ETS with those not covered by the regulation. Both groups 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 & Prince, 1989; Fischer, Parry, & Pizer, 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 IV and V, 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.

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 eighteen EU ETS countries we are studying, and roughly 80% of installations and emissions EU ETS–wide (see table 1).[13]

[13] See appendix A in the online appendix for more details on how the link between firm data and regulatory data was constructed.

TABLE 1.—COVERAGE OF THE EU ETS

	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 Republic	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	1,100	450.2	97.5	99.6
Germany	1,944	1,486.3	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	1,072	498.1	98.5	99.9
Sweden	774	67.6	93.9	98.8
United Kingdom	1,107	628.0	83.3	97.0
Total	10,056	5,256.6	93.1	98.7
Total EU ETS	12,122	6,321.3	77.2	82.0

The first two columns of this table show the number of phase 1 installations in each of the eighteen 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 summarize our data set's EU ETS coverage for our eighteen countries, as well as as a proportion of the EU ETS as a whole.

FIGURE 3.—COMPARING THE SHARE OF LOW-CARBON PATENTS, 1978–2009

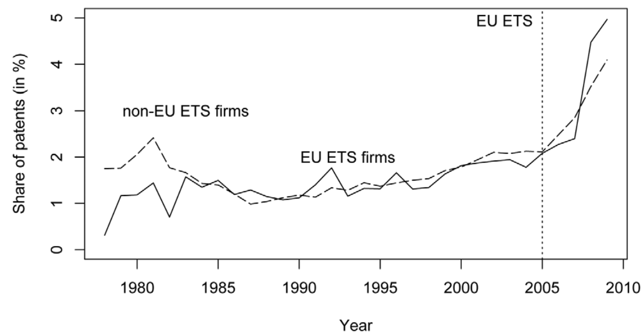

Having identified the subset of firms directly affected by the EU ETS, we can now look separately at the EU ETS and non–EU ETS trends in low-carbon patenting. Figure 3 shows that the share of low-carbon patents was roughly the same among EU ETS and non–EU ETS firms in the five years before the EU ETS launched. After 2005, the share of low-carbon patents among EU ETS firms looks to have risen faster than among non–EU ETS firms.[14] 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

[14] One might be concerned that the surge in patenting activity by EU ETS firms compared to non–EU ETS companies might have been accompanied by a concurrent drop in the relative average quality of inventions 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.

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 from 2005 to 2009 compared to 972 patents in the five 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 increase of 65%). If we 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 one in twelve low-carbon patents filed in the five 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. We now turn to the task of constructing such a comparison.

IV. The Direct Impact of the EU ETS

A. *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, and so on. 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 that installations have a production capacity exceeding 2.5 tonnes per hour. Installations manufacturing glass and glass fiber 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, 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 behavior 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, e.g., Dehejia & Wahba, 1999).

Along with patent portfolios, our data set contains information on the country and economic sector in which firms operate, as well as other firm-level information such as turnover and employment.[15] Using these 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). This has not always been possible, for two main reasons. First, the records of turnover become less and less complete further back in time. In fact, we have pre-2005 records on the turnover for only 3,564 out of the 5,568 EU ETS firms. Second, 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 although 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 a lack of suitable comparators, the sample of EU ETS firms is further reduced to 3,428. We return to the omitted firms in section IV.C, 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

[15] Economic sectors are defined at the three-digit level for the NACE Rev. 2 industry classification. A few examples of these sector definitions 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."

FIGURE 4.—COMPARISON OF MATCHED EU ETS AND NON–EU ETS FIRMS

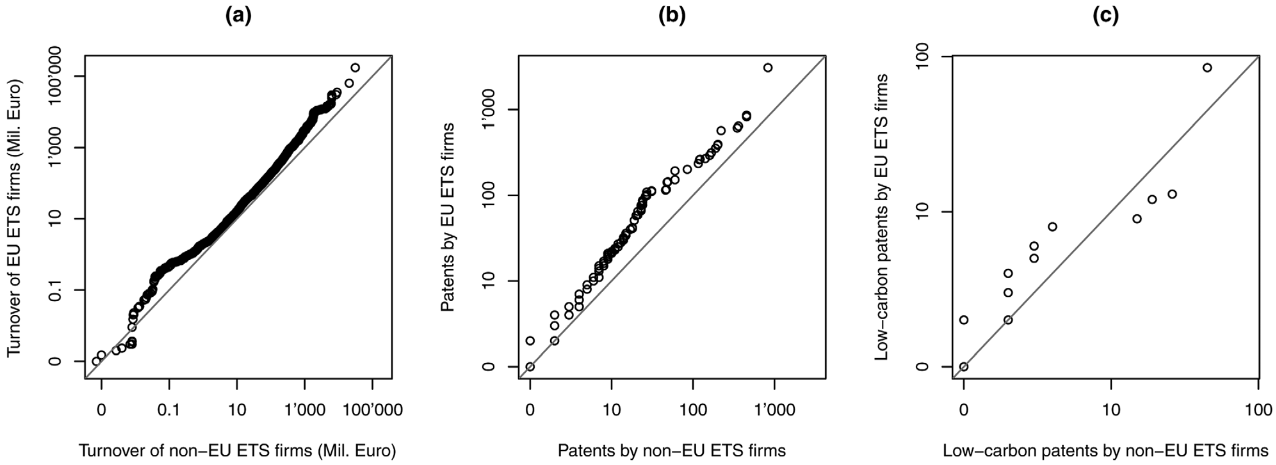

Panel (a) displays the empirical quantile-quantile (e-QQ) plot for average turnover in the four years before the EU ETS (2001–2004). Each dot gives the value for one EU ETS firm and the average for a group of matched non–EU ETS firms, shown on logarithmic scales. The year 2001 is the first one for which 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 from filed 2000 to 2004, respectively, once again shown on logarithmic scales.

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% significance level)
Turnover (in million euros)	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 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 three-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 & Rubin, 1973; Ho et al., 2007). Using Wilcoxon's signed-rank test, we are just unable to reject at the 5% significance level the hypothesis that the location shift parameter lies within 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 0, using a method outlined by Rosenbaum, 2009. More details are in section IV). 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 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 0. Standard t-tests for differences in means reject the hypotheses of substantive differences for all variables (not reported).

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.[16] 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).

B. *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. First, matching appears to have produced a set of EU ETS and non–EU ETS firms roughly comparable prior to 2005 in both their general level of low-carbon patenting and in that they do not appear to exhibit different trends. Second, the two groups begin to diverge after 2005, coinciding with the introduction of the new policy.

[16] See appendix B for technical details about how the matching was implemented.

FIGURE 5.—LOW-CARBON PATENTS BY MATCHED EU ETS AND NON–EU ETS FIRMS

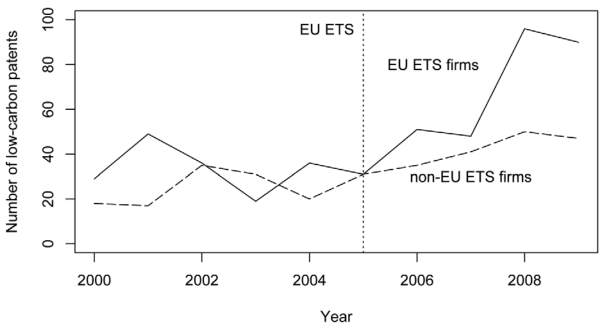

To examine this pattern more precisely, we measure the change in the number of low-carbon patents from 2000–2004 to 2005–2009 for each firm. This means that even after matching, we take account of any additional time-invariant firm-level heterogeneity. The outcomes of the matched control firms are then 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 Tobit estimator. We can imagine there being a latent variable that can take any value, but we can observe only numbers of 0 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 straightforward exercise, and getting the model wrong carries with it the risk of introducing new biases. The analogous maximum likelihood estimator will also generally be inconsistent, especially when applied to panel data (Chay & Powell, 2001). Instead, we can account for the censoring at zero using a Tobit-modified empirical-likelihood estimator, as Rosenbaum (2009) outlined. 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 recalculate our similarity measure. Each of the difference-in-differences, Δ, is adjusted according to the formula:

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

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 pretreatment 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 100 × (100 "times" (1-alpha) gives the correct confidence level) $(1-\alpha)\%$ 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 with Wilcoxon's signed-rank test. This provides a nonparametric 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 316 low-carbon patents in the period 2005 to 2009. Subtracting two low-carbon patents from each of our matched EU ETS firms (and accounting for censoring at 0) 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 among our sample of EU ETS firms, or an increase of 36.2% (20.2%, 69.0%) compared to what we expect would have happened had they not been regulated under the EU ETS. Because these firms account for only 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 0, the EU ETS is responsible for 183 $(111, 299)$ additional low-carbon patents. This amounts to a 9.1% (5.3%, 15.8%) increase in their low-carbon patenting, or a 0.83% (0.50%, 1.36%) increase in the total number of low-carbon patents filed at the EPO over 2005 to 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 0, 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,

TABLE 3.—SUMMARY OF RESULTS

	Matching Estimates		Naive Estimates
	Matched Sample	Full Sample	Full Sample
Additional low-carbon patents	84 (53,129)	183 (111,299)	585.2
As % increase	36.2 (20.2, 69.0)	9.1 (5.3, 15.8)	36.5
As % increase of EPO	0.38 (0.24, 0.58)	0.83 (0.50, 1.36)	2.6
Additional other patents	305 (305, 512.9)	551 (551, 934)	9,072.8
As % increase	1.9 (1.9, 3.2)	0.83 (0.83, 1.44)	16.0
As % increase of EPO	0.041 (0.041, 0.068)	0.072 (0.072, 0.12)	1.2

Point estimates, along with 95% confidence intervals in parentheses 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 III.

be it a low-carbon technology or not.[17] 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. A related concern is that the increase in low-carbon innovation will displace, or crowd out, the development of other technologies (Popp & 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 one 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% versus 1.9% (the difference is significant at 5% level). Put another way, the system has nearly had a twenty 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 541 $(541, 934)$ additional other patents, which amounts to a 0.83% (0.83%, 1.44%) increase in their other patenting, or a 0.072% (0.072%, 0.12%) increase in the total number of other patents filed at the EPO over 2005 to 2009.

[17] Apart from technological complementarity and cost minimization, firms might fear that the EU ETS will make them less competitive, and hence innovate more across the board to maintain market share. Alternatively, the windfall profits that were earned from the free allowances may have eased pressure from shareholders, so it became easier for EU ETS firms to invest in previously sidelined research projects. One can, of course, imagine still other mechanisms whereby a price on carbon increases patenting for other technologies. The main point here is only that economic theory does not rule it out.

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 data sets, the small number of regulated companies active in each sectors 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 two 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 two 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 summarized for convenience in table 3, along with comparable naive estimates for the full sample of EU ETS firms (calculated as in section III). 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 overallocation 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 ETS firms or 9.1% among the full sample—translates into a 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

FIGURE 6.—COMPARISON OF MATCHED EU ETS AND NON–EU ETS FIRMS ON UNOBSERVED VARIABLE

be attributed to the EU ETS.[18] 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.

C. *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 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 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

[18] The number of low-carbon patents filed at the EPO increased by 9,054 from the period 2000–2004 to 2005–2009. The 183 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 4,725.5 low-carbon patents, of which the 183 additional low-carbon patents would amount to barely 4%.

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% significance level)
Employees	25	± 904.07	± 106.75

See the table 2 footnote for details on how to read this table.

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.

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 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 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 & Svendsen, 2005). Ultimately, 575 such installations—slightly less than 5% of all EU ETS installations—opted in. 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 100 matched pairs). This returns an estimate of 2 (1, 5.99) additional low-carbon patents, and of 1 (1, 1.99) other additional patent. 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 test 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 1,062 additional low-carbon patents. Since they did not file this many low-carbon patents over 2005 to 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 ten 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 made prediction (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.[19] We have estimated above that our sample of matched EU ETS firms accounts 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 their chances of increasing their low-carbon patenting. Examples 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 0, 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% level the hypothesis that the treatment effect is really 0.[20] 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.[21]

In sum, matching has achieved balance on at least one unobserved characteristic, which might suggest it has balanced other unobserved variables as well, as a truly randomized experiment would have. Even if this is not the case, though, it appears that 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 behavior prior to 2005 from the EU ETS firms dropped from our matched sample. An unmatched EU ETS firm would have been dropped either because it was an outlier or because crucial data were missing that prevented matching. In practice, most were dropped because financial data were missing. This has two consequences. First, we can reliably compare the patenting behavior of matched and unmatched EU ETS firms. Second, there may be substantial overlap in the levels of patenting of matched and unmatched EU ETS firms. 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.[22] 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% significance level:

[19] In 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).

[20] In 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 $\Gamma = 1.45$.

[21] We thank the anonymous referee for this suggestion.

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

± 1.99). This mirrors our findings in table 2 and follows in 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% significance 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 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: (a) increasing the sample size by matching some of those unmatched EU ETS firms, (b) calculating an upper bound for our estimates, and (c) calculating a lower bound for the out-of-sample response necessary to qualitatively affect our conclusions. First, 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, 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 rematch 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 previous difference-in-differences, we can then estimate the upper bound of the treatment effect.[23] This procedure produces point estimates of 13 $(3.01, 39.99)$ for low-carbon and 5.75 $(4, 10.99)$ 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 0 patents after 2005. Subtracting a large number of patents from each firm and accounting for censoring at 0 therefore does not add as many patents as the higher point estimates perhaps might suggest. The new estimates translate into 500 (230.37, 866.9) additional low-carbon patents and 2,005.75 (1,558, 3,144.95) additional other patents, or increases of 29.6 (11.8%, 65.6%) and 3.1% (2.4%, 5.0%) 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 2.3% (1.0%, 4.0%) increase in low-carbon patenting and a 0.26% (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 together would have to have filed 1,062 additional low-carbon patents over 2005 to 2009. Subtracting our best estimate of 183 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 879 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 to have filed 0.32 additional low-carbon patents in the same period. The out-of-sample response would have to be ten times greater than the in-sample response. Even if we use the upper-bound estimate (in-sample firms filed 500 additional low-carbon patents), the out-of-sample firms would have to have filed 562 extra low-carbon patents, or at

[23] This bound is analogous to the sharp bounds derived by Manski (2007) 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."

least 0.2 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 therefore seems 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. 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 rematched our EU ETS firms to similar firms in Norway, Switzerland, Bulgaria, and Romania (four countries that did not launch the EU ETS in 2005 and two of which have remained outside). We also rematched our EU ETS firms to similar U.S. 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 rematched our EU ETS firms using 2003 as the treatment year instead of 2005. The treatment effect for the period 2003 to 2004 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. 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 0.[24]

Table 5.—Summary of Treatment Effect Estimates

	Low-Carbon	Other
Original estimate	2	1
	(1, 5)	(1, 1.99)
Alternative specifications		
Excluding opt-ins	2	1
	(1, 5.99)	(1, 1.99)
Matching with 2005 turnover	2.75	1
	(1, 5.99)	(1, 1.99)
Expanded low-carbon definition	1.75	1
	(1, 3.99)	(1, 1.99)
Non–EU ETS firms from Norway, Switzerland, Romania, and Bulgaria	1	2
	(0, 1.99)	(1, 3)
Non–EU ETS firms from United States	−1	0
	(−1.99, 0.99)	(−0.99, 0.99)
Treatment years 2003–2004	−1.75	−1
	(−∞, 1.99)	(−4, −0.01)
Upper bounds		
Assuming 1% boost to EPO	20.4	–
Low-carbon patenting	–	–
Assuming all patents of unmatched EU ETS firms are additional	13	5.75
	(3.01, 39.99)	(4, 10.99)

For convenience, table 5 summarizes the results from robustness tests that involved reestimating the treatment effect under alternative assumptions. More information and a few additional robustness tests can be found in appendix 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.[25]

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

[25] One must be careful also because some of the tests we have used to investigate these alternative explanations, though addressing one potential source

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.

V. 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 that we would expect the indirect impact to be comparatively small. First, 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 & Prince, 1989; Fischer et al., 2003). The asymmetry 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.

Second, EU ETS firms have filed over 120,000 patents with the EPO since 2000; approximately 2.5% of them 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.

Third, 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 firm to either provide technologies to EU ETS firms once the regulations came into force or 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, 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 purchasers 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 over 1978 to 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.[26] Figure 7 and table 6 show the properties of our matched sample.[27]

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 the counterfactual, it would still translate into a negligible increase relative to the number of low-carbon patents filed at the EPO (0.2%). Extrapolating the number to all 11,603 co-patenters would naturally make

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, 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.

[26] Compared 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.

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

FIGURE 7.—COMPARISON OF MATCHED CO-PATENTERS AND NON–CO-PATENTING FIRMS

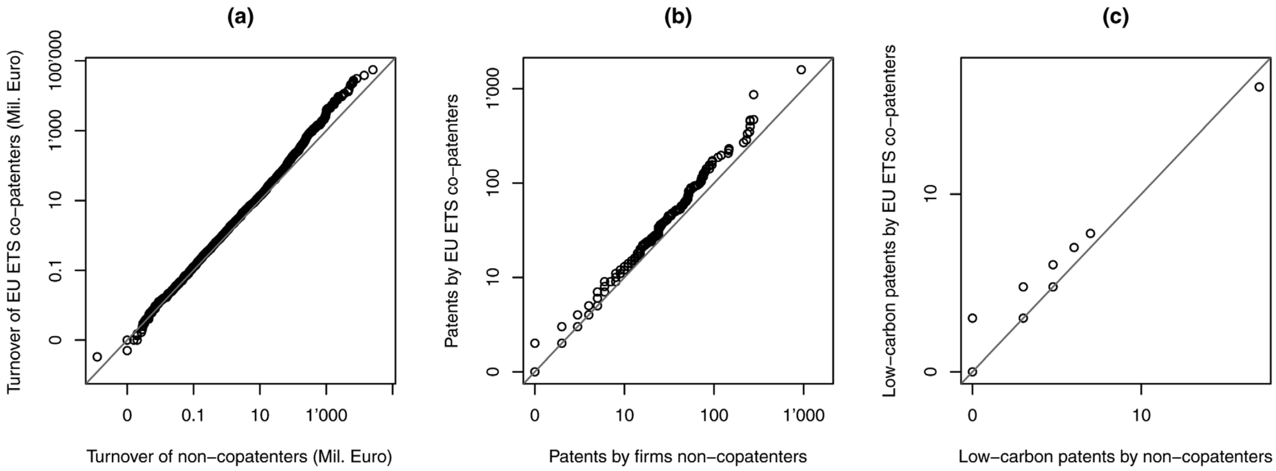

TABLE 6.—EQUIVALENCE TESTS FOR MATCHED CO-PATENTERS AND NON-CO-PATENTING FIRMS

	Median Difference between EU ETS and Non–EU ETS Firms	Equivalence Range	Critical Equivalence Range (5% significance level)
Turnover (in thousands of euros)	14.90	± 304,382.80	± 1,421.00
Patents	0	± 7.07	± < 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	± 1,613.82	± 20.66

See the table 2 footnote 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 0. 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 are also included at the bottom of this table.

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 0, 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 or to demand new technologies from EU ETS firms once the EU ETS launched. Yet, taken together, co-patenters appear to behave no different from 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 overinterpreted. 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.

VI. Discussion

The EU ETS launched in 2005 amid promises and pessimism. An important objective of carbon market programs like the EU ETS is to 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 five 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 behavior 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 behavior 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 had they not been regulated under the EU ETS. What is more, our estimates suggest that the system has also encouraged EU ETS firms to increase 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 eighteen countries, the EU ETS would account for a 9.1% increase in low-carbon patenting and a 0.83% 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 (0.83% for the full sample) and a meager 0.041% boost to patenting for other technologies (0.072% for the full sample). We should nevertheless remain cognizant of the fact that patent counts tend to emphasize technological changes and do not fully reflect the development of new operational strategies or capital investments and divestments as they relate to already available technologies. Other measures may provide a better understanding of the System's impact on other such aspects of innovation.

To test whether our focus on EU ETS firms 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 provide only 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 purchasers. If data on patent licensing agreements could be obtained, researchers in the future may 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 & Mody, 1996; Brunnermeier & Cohen, 2003; Popp, 2002, 2003, 2006; Arimura et al., 2007; Lanoie et al., 2007). In contrast, studies of previous emissions trading programs, like the U.S. Acid Rain Program, at best unearth evidence of very small impacts on directed technological change (Popp, 2003; Lange & 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 overallocated. 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 nonreaction 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 program'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 behavior 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 & Schultze, 1975; Pizer & Popp, 2008). Our results, read in combination with the findings of the broader literature, suggest that environmental policy instruments may differ also in the distribution of impacts on directed technological change. This could be potentially significant because of the positive spillovers 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, we have also looked at what types of technologies those patents protect, conditional on the estimated treatment effect. Most of them appear to protect alternative energy and energy storage, with the remaining ones focusing on energy efficiency. Most of these additional low-carbon patents belong to chemicals manufacturers, energy companies, and automobile manufacturers (see appendix D for details). 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 and different allocation rules, for example. 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. 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 road map to a low-carbon economy in 2050 (European Commission, 2011). Policymakers 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 such operational changes as fuel switching rather than technological changes, much as 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. 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 among a relatively small group of firms, their response appears 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., "Semiparametric Difference-in-Differences Estimators," *Review of Economic Studies* 72:1 (2005), 1–19.

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

Aghion, P., A. Dechezleprêtre, D. Hemous, R. Martin, and J. V. Reenen, "Carbon Taxes, Path Dependency and Directed Technical Change: Evidence from the Auto Industry," NBER working paper 18596 (2012).

Ambec, S., M. Cohen, S. Elgie, and P. Lanoie, "The Porter Hypothesis at 20: Can Environmental Regulation Enhance Innovation and Competitiveness?" *Review of Environmental Economics and Policy* 7 (2013), 2–22.

Anderson, B., F. Convery, and C. D. Maria, "Technological Change and the EU ETS: The Case of Ireland," IEFE working paper 43 (2011).

Anderson, B., and C. Di Maria, "Abatement and Allocation in the Pilot Phase of the EU ETS," *Environmental and Resource Economics* 48 (2011), 88–103.

Arimura, T. H., A. Hibiki, and N. Johnstone, "An Empirical Study of Environmental R&D: What Encourages Facilities to Be Environmentally Innovative?" N. Johnstone, ed., *Corporate Behaviour and Environmental Policy* (Cheltenham, UK: Edward Elgar in association with OECD, 2007).

Arora, A., M. Ceccagnoli, and W. M. Cohen, "R&D and the Patent Premium," *International Journal of Industrial Organization* 26 (2008), 1153–1179.

Brunnermeier, S. B., and M. A. Cohen, "Determinants of Environmental Innovation in US Manufacturing Industries," *Journal of Environmental Economics and Management* 45 (2003), 278–293.

Burtraw, D., and S. Szambelan, "US Emissions Trading Markets for SO_2 and NOx," *Resources for the Future discussion paper* 09-40 (2009).

Chay, K. Y., and J. L. Powell, "Semiparametric Censored Regression Models," *Journal of Economic Perspectives* 15:4 (2001), 29–42.

Cochran, W., and D. Rubin, "Controlling Bias in Observational Studies: A Review," *Sankhyā: The Indian Journal of Statistics, Series A* 35 (1973), 417–446.

Cohen, W. M., R. R. Nelson, and J. P. Walsh "Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not)," NBER working paper 7552 (2000).

Dechezleprêtre, A., M. Glachant, I. Haščič, N. Johnstone, and Y. Ménière, "Invention and Transfer of Climate Change–Mitigation Technologies: A Global Analysis," *Review of Environmental Economics and Policy* 5:1 (2011), 109.

Dehejia, R., and S. Wahba, "Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs," *Journal of the American Statistical Association* 94 (1999), 1053–1062.

Dekker, T., H. R. Vollebergh, F. P. de Vries, and C. A. Withagen, "Inciting Protocols," *Journal of Environmental Economics and Management* 64:1 (2012), 45–67.

Delarue, E., A. Ellerman, and W. D'haeseleer, "Short-Term CO_2 Abatement in the European Power Sector: 2005–2006," *Climate Change Economics* 1:2 (2010), 113–133.

Delarue, E., K. Voorspools, and W. D'haeseleer, "Fuel Switching in the Electricity Sector under the EU ETS: Review and Prospective," *Journal of Energy Engineering* 134:2 (2008), 40–46.

Dernis, H., D. Guellec, and B. V. Pottelsberghe, "Using Patent Counts for Cross-Country Comparisons of Technology Output," *STI Review* 27 (2001).

Ellerman, A., and B. Buchner, "Over-Allocation or Abatement? A Preliminary Analysis of the EU ETS Based on the 2005–06 Emissions Data," *Environmental and Resource Economics* 41 (2008), 267–287.

Ellerman, A. D., F. J. Convery, and C. de Perthuis, *Pricing Carbon: The European Union Emissions Trading Scheme* (Cambridge: Cambridge University Press, 2010).

European Commission, "EU Action Against Climate Change: EU Emissions Trading—An Open Scheme Promoting Global Innovation" (2005). http://bookshop.ec.europa.eu/ent/eu-emissions-trading-an-open-system-promoting-global-innovation-pbKH807426/.

——— "A Roadmap for Moving to a Competitive Low Carbon Economy in 2050," EU technical report COM 112 (2011).

——— "Emissions Trading: Annual Compliance Round-Up Shows Declining Emissions in 2011," press release IP/12/477 (2012).

Fischer, C., I. Parry, and W. Pizer, "Instrument Choice for Environmental Protection When Technological Innovation Is Endogenous," *Journal of Environmental Economics and Management* 45 (2003), 523–545.

Fowlie, M., "Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement," *American Economic Review* 100 (2010), 837–869.

Gagelmann, F., and M. Frondel, "The Impact of Emission Trading on Innovation—Science Fiction or Reality?" *European Environment* 15 (2005), 203–211.

Gerlagh, R., "A Climate-Change Policy Induced Shift from Innovations in Carbon-Energy Production to Carbon-Energy Savings," *Energy Economics* 30 (2008), 425–448.

Goulder, L., and S. Schneider, "Induced Technological Change and the Attractiveness of CO_2 Abatement Policies," *Resource and Energy Economics* 21 (1999), 211–253.

Greenstone, M., and T. Gayer, "Quasi-Experimental and Experimental Approaches to Environmental Economics," *Journal of Environmental Economics and Management* 57:1 (2009), 21–44.

Griliches, Z., *R&D, Patents, and Productivity* (Chicago: University of Chicago Press, 1984).

Grubb, M., C. Azar, and U. Persson, "Allowance Allocation in the European Emissions Trading System: A Commentary," *Climate Policy* 5:1 (2005), 127–136.

Hall, B., A. Jaffe, and M. Trajtenberg, "Market Value and Patent Citations," *RAND Journal of Economics* 36:1 (2005), 16–38.

Harhoff, D., F. Narin, F. Scherer, and K. Vopel, "Citation Frequency and the Value of Patented Inventions," this REVIEW 81 (1999), 511–515.

Harhoff, D., F. Scherer, and K. Vopel, "Citations, Family Size, Opposition and the Value of Patent Rights," *Research Policy* 32 (2003), 1343–1363.

Heckman, J., H. Ichimura, J. Smith, and P. Todd, "Characterizing Selection Bias Using Experimental Data," *Econometrica* 66 (1998), 1017–1098.

Heckman, J. J., H. Ichimura, and P. Todd, "Matching as an Econometric Evaluation Estimator," *Review of Economic Studies* 65 (1998), 261–294.

Hicks, J. R., *The Theory of Wages* (New York: Macmillan, 1932).

Ho, D. E., K. Imai, G. King, and E. A. Stuart, "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference," *Political Analysis* 15 (2007), 199–236.

Hoffmann, V. H., "EU ETS and Investment Decisions: The Case of the German Electricity Industry," *European Management Journal* 25 (2007), 464–474.

Hoppmann, J., M. Peters, M. Schneider, and V. H. Hoffmann, "The Two Faces of Market Support: How Deployment Policies Affect Technological Exploration and Exploitation in the Solar Photovoltaic Industry," *Research Policy* 42 (2013), 989–1003.

Jaffe, A. B., R. G. Newell, and R. N. Stavins, "Technological Change and the Environment" (pp. 461–516), in Karl-Göran Mäler and Jeffrey R. Vincent, eds., *Handbook of Environmental Economics* (New York: Elsevier, 2003).

Jaffe, A. B., and K. Palmer, "Environmental Regulation and Innovation: A Panel Data Study," this REVIEW 79 (1997), 610–619.

Johnstone, N., I. Haščič, and D. Popp, "Renewable Energy Policies and Technological Innovation: Evidence Based on Patent Counts," *Environmental and Resource Economics* 45 (2010), 133–155.

Kaufer, E., *The Economics of the Patent System* (New York: Routledge, 1989).

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

Kossoy, A., and P. Guigon, *State and Trends of the Carbon Market 2012: Annual Report* (Washington, DC: World Bank, 2012).

Kossoy, A., K. Oppermann, R. C. Reddy, M. Bosi, S. Boukerche, N. Höhne, N. Klein, et al., "Mapping Carbon Pricing Initiatives: Developments and Prospects," World Bank and Ecofys technical report (2013).

Lange, I., and A. Bellas, "Technological Change for Sulfur Dioxide Scrubbers under Market-Based Regulation," *Land Economics* 81 (2005), 546–556.

Lanjouw, J., and A. Mody, "Innovation and the International Diffusion of Environmentally Responsive Technology," *Research Policy* 25 (1996), 549–571.

Lanoie, P., J. Laurent-Lucchetti, N. Johnstone, and S. Ambec, "Environmental Policy, Innovation and Performance: New Insights on the Porter Hypothesis," CIRANO working paper (2007).

List, J. A., D. L. Millimet, P. G. Fredriksson, and W. W. McHone, "Effects of Environmental Regulations on Manufacturing Plant Births: Evidence from a Propensity Score Matching Estimator," this REVIEW 85 (2003), 944–952.

Manski, C. F., *Identification for Prediction and Decision* (Cambridge, MA: Harvard University Press, 2007).

Markussen, P., and G. Svendsen, "Industry Lobbying and the Political Economy of GHG Trade in the European Union," *Energy Policy* 33 (2005), 245–255.

Martin, R., M. Muûls, and U. Wagner, "Climate Change, Investment and Carbon Markets and Prices: Evidence from Manager Interviews," *Climate Strategies, Carbon Pricing for Low-Carbon Investment Project* (2011).

Milliman, S., and R. Prince, "Firm Incentives to Promote Technological Change in Pollution Control," *Journal of Environmental Economics and Management* 17 (1989), 247–265.

Nemet, G., "Demand-Pull, Technology-Push, and Government-Led Incentives for Non-Incremental Technical Change," *Research Policy* 38 (2009), 700–709.

Newell, R., A. Jaffe, and R. Stavins, "The Induced Innovation Hypothesis and Energy-Saving Technological Change," *Quarterly Journal of Economics* 114 (1999), 941–975.

OECD, "OECD Patent Statistics Manual," OECD technical report (2009).

Petsonk, A., and J. Cozijnsen, "Harvesting the Low-Carbon Cornucopia: How the European Union Emissions Trading System (EU-ETS) Is Spurring Innovation and Scoring Results" (Environmental Defense Fund 2007). http://www.edf.org/sites/default/files/harvesting-the-low-carbon-cornucopia-march2007.pdf.

Pizer, W. A., and D. Popp, "Endogenizing Technological Change: Matching Empirical Evidence to Modeling Needs," *Energy Economics* 30 (2008), 2754–2770.

Popp, D., "Induced Innovation and Energy Prices," *American Economic Review* 92 (2002), 160–180.

——— "Pollution Control Innovations and the Clean Air Act of 1990," *Journal of Policy Analysis and Management* 22 (2003), 641–660.

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

——— "International Innovation and Diffusion of Air Pollution Control Technologies: The Effects of NO_X and SO_2 Regulation in the US, Japan, and Germany," *Journal of Environmental Economics and Management* 51 (2006), 46–71.

——— "Innovation and Climate Policy," NBER working paper 15673 (2010).

Popp, D., and R. Newell, "Where Does Energy R&D Come From? Examining Crowding Out from Energy R&D," *Energy Economics* 34 (2012), 980–991.

Popp, D., R. Newell, and A. Jaffe, "Energy, the Environment, and Technological Change" (pp. 837–937), in Bronwyn Hall and Nathan Rosenberg, eds., *Handbook of the Economics of Innovation* (Orlando, FL: Academic Press/Elsevier, 2010).

Porter, M. E., "Essay: America's Green Strategy," *Scientific American* 264 (1991).

Rosenbaum, P., "Sensitivity Analysis for Certain Permutation Inferences in Matched Observational Studies," *Biometrika* 74 (1987), 13–26.

——— *Design of Observational Studies* (New York: Springer, 2009).

Rosenbaum, P., and J. Silber, "Amplification of Sensitivity Analysis in Matched Observational Studies," *Journal of the American Statistical Association*, 104 (2009), 1398–1405.

Schleich, J., and R. Betz, "Incentives for Energy Efficiency and Innovation in the European Emission Trading System" (pp. 1495–1506), in *Proceedings of the 2005 ECEEE Summer Study: What Works and Who Delivers?* (2005).

Schmalensee, R., P. Joskow, A. Ellerman, J. Montero, and E. Bailey, "An Interim Evaluation of Sulfur Dioxide Emissions Trading," *Journal of Economic Perspectives* 12 (1998), 53–68.

Smith, J., and P. Todd, "Does Matching Overcome LaLonde's Critique of Nonexperimental Estimators?" *Journal of Econometrics* 125 (2005), 305–353.

Stavins, R., "A US Cap-and-Trade System to Address Global Climate Change," Regulatory Policy Program working paper RPP-2007-04 (2007).

Taylor, M. R., "Innovation Under Cap-and-Trade Programs," *Proceedings of the National Academy of Sciences* 109 (2012), 4804–4809. PMID: 22411797.

Tomás, R., F. R. Ribeiro, V. Santos, J. Gomes, and J. Bordado, "Assessment of the Impact of the European CO_2 Emissions Trading Scheme on the Portuguese Chemical Industry," *Energy Policy* 38 (2010), 626–632.

Trajtenberg, M., "A Penny for Your Quotes: Patent Citations and the Value of Innovations," *Rand Journal of Economics* 21 (1990), 172–187.

van der Zwaan, B., R. Gerlagh, and L. Schrattenholzer, et al., "Endogenous Technological Change in Climate Change Modelling," *Energy Economics* 24:1 (2002), 1–19.

van Zeebroeck, N., "The Puzzle of Patent Value Indicators," *Economics of Innovation and New Technology* 20:1 (2011), 33–62.

Veefkind, V., J. Hurtado-Albir, S. Angelucci, K. Karachalios, and N. Thumm, "A New EPO Classification Scheme for Climate Change Mitigation Technologies," *World Patent Information* 34 (2012), 106–111.