Indirectly Induced Innovation: Consequences for Environmental Policy Analysis

(Job Market Paper)

Steve Miller*

Bren School of Environmental Science and Management, University of California, Santa Barbara

November 2014

Abstract

Environmental regulation is increasingly judged on its ability to stimulate innovation. While prior studies focus primarily on the direct effect on innovation by regulated firms, this paper examines the importance of indirectly induced innovation by both regulated and unregulated firms. Two channels for indirect effects are considered: pass-through of regulation-imposed costs into output prices and knowledge spillovers. Both are well-documented phenomena on their own, but their consequences for innovation by unregulated firms are largely ignored. This paper uses simple models to illustrate such effects and the resulting bias in standard policy analysis, then employs dynamic count models to empirically test these claims in an application to the European Union Emissions Trading System. Estimates suggest that indirect effects of regulation on innovation are at least as large as commonly estimated direct effects. Indirectly induced innovation may thus form an important part of the overall innovative response to environmental regulation.

Keywords: induced innovation, technical change, emissions trading

JEL Classification: O3, Q55, Q5

*This paper has benefited greatly from the feedback of others, most notably Chris Costello, Bob Deacon, and Olivier Deschênes. I am also grateful for the valuable comments offered by Raphael Calel, Gary Libecap, Paulina Oliva, Chris Severen, and participants in the environmental economics lunch at UCSB. Any errors are, of course, my own.

1 Introduction

Many discussions of climate policy emphasize the prospects for technological change to keep mitigation and adaptation costs at manageable levels. This has led to increased focus on the ability of different climate policies to spur innovation in reducing production and increasing capture of greenhouse gases (Milliman and Prince 1989; Fischer et al. 2003; Requate and Unold 2003; Fischer and Newell 2008). The rationale for this emphasis is straightforward: since environmental policy often increases the costs of polluting, firms will seek ways to avoid those costs, including development of new technologies that reduce pollution. The idea that changes in relative prices should influence innovation traces back to Hicks (1932), and the application of those ideas to environmental policy intensified with a set of provocative claims by Porter and van der Linde (1995). Those authors posited not only that environmental regulation would induce innovation, but also that innovation could lead to enhanced competitiveness for firms (or countries) subject to environmental regulation.

Empirical studies seeking to test these claims by quantifying the innovation induced by environmental policy have thus far produced mixed results. Early studies found suggestive evidence for a link between policy stringency and innovation (Lanjouw and Mody 1996), but a follow-up study provided no evidence that link was causal (Jaffe and Palmer 1997). Subsequent work again reversed the conclusions, finding that policy stringency may increase pollution-relevant innovation but has a negligible effect on overall patent rates (Brunnermeier and Cohen 2003). More recent studies have echoed the importance of outcome variable choice: Johnstone et al. (2010) find that renewable energy policies have heterogeneous effects on innovation across different types of renewable technologies (e.g. wind vs solar).

Most recently, the scale and scope of the European Union's Emissions Trading Scheme (EU ETS) has generated great interest in whether the policy is aiding in the transition to a lower-carbon economy. Analyses thus far have suggested that, while regulated firms have increased their low-carbon patenting activities substantially, the program accounts for only a small fraction of the increase in low-carbon patenting in the EU (Calel and Dechezleprêtre 2014). Other factors, such as fuel costs and country-specific renewable energy policies, may be driving the bulk of low-carbon patenting (Hoffmann 2007). The EU ETS may be influencing other types of innovation, such as (disembodied) process innovation like fuel switching (Delarue et al. 2008), but the current best estimates of the effect of the EU ETS on product innovation indicate only a small effect.

In this paper, I argue that empirical studies focused solely on induced innovation by regulated firms may be missing important causal effects of regulation on innovation, and thus will produce biased estimates of the total effect of environmental policy. Commonly-used treatment effect estimators (e.g. difference-in-difference, matching, and propensity score approaches) identify the *direct* effects of policy on innovation by regulated firms. However, when a regulated firm responds to a policy through innovation or other means, that firm's actions may change innovation incentives for other regulated and even unregulated firms, thereby creating *indirect* policy effects.[1] If those indirect effects are present, the direct and total policy effects will diverge, and even unbiased estimators of direct effects will produce biased estimates of the total effects of a policy. I focus on two channels for such indirect effects: knowledge spillovers that augment the knowledge stock available to all firms, and pass-through of regulatory costs into the price of outputs used by unregulated firms as inputs (e.g. electricity). Both knowledge stocks and energy prices have been shown to impact patenting output of firms (e.g. Popp 2002), so any impact of a policy on those intermediate quantities should ultimately affect innovation.

The analysis here draws upon prior work pertaining to treatment effect spillovers. Such spillovers have consequences for identification and bias in many settings, including studies of social effects (Manski 1993) and medical treatments (Miguel and Kremer 2004). In those other contexts, a few authors have also suggested estimation strategies that can account for certain types of indirect effects (Hudgens and Halloran 2008; Aronow and Samii 2012). However, the issue remains largely ignored in studies of innovation induced by environmental regulation. One of the aforementioned studies, that by Calel and Dechezleprêtre (2014), correctly acknowledges the potential for indirect effects. Their analysis of indirect effects is, however, limited to unregulated firms with a history of co-patenting with regulated firms. I aim to build upon these prior studies to offer a more extensive treatment of indirectly induced innovation, combining both theory and empirics in a way that I hope sheds light on the prevalence and consequences of indirect effects of environmental regulations.

With this background, my goals for this study are threefold. First, I use simple models to illustrate the potential bias in standard treatment effect estimators and to determine when indirect innovation effects are likely to matter, and when they are not. The qualitative conclusions from these models apply, in principle, to innovation in response to many forms of regulation. My second goal is to develop an alternate estimation strategy (using dynamic count models) that can account for and quantify indirect effects of environmental regulation on innovation. Finally, I apply these methods to quantify the *relative* importance of indirect effects in the context of the EU ETS using a firm-level panel data set. There is good reason to think that, due both to low emissions prices and the partial coverage of the EU ETS, the *absolute* amount of innovation induced by the EU ETS is small, and so too are any indirect

[1] In the treatment effects literature, such indirect effects are sometimes referred to as interference (Cox 1958).

effects. However, getting a better handle on the relative importance of indirect effects may provide important insight into the potential innovation effects of stricter climate policy in the EU or elsewhere.

The EU ETS provides a useful test of the relative importance of indirect effects for two primary reasons.[2] First, evidence suggests that regulated power producers are indeed passing along permit costs into the price of electricity, such that one of the hypothesized channels giving rise to innovation by unregulated firms may be active. Casual observation of the movement of prices for both electricity and emissions in Germany suggests a relationship between the two (Figure 1), and more formal analyses (e.g. Sijm et al. 2006; Sijm et al. 2008; Fabra and Reguant 2013) imply that the relationship is causal.[3] In an early example, Sijm et al. (2006) regress the difference between power and coal prices on the permit price and repeat that analysis with the difference between the power and gas prices as the response, finding pass-through rates of 60% (base load) and 117% (peak load) using year-ahead power market data. A more recent study of the Spanish wholesale market estimates pass-through rate of 80% (base load) to 100% (peak load), though Spanish retail prices are unlikely to respond in kind due to heavy price regulation in that market (Fabra and Reguant 2013). Other studies point to a link between CO2 prices and electricity prices, but do not explicitly estimate pass-through rates (e.g. Keppler and Mansanet-Bataller 2010). Finally, Convery et al. (2008) suggest that carbon costs could comprise anywhere from 2% to 9% of electricity prices. Given the prevalence of electricity as an input to production, such cost pass-through is likely to impact innovation incentives for a large number of unregulated firms.

Second, in order to empirically identify indirect effects, some variation in the exposure of unregulated firms to the hypothesized channels is desirable, and in the EU the potential for cost pass-through varies geographically. While many wholesale electricity markets in Europe have undergone liberalization, the degree of retail price regulation varies by market. I use such cross-country differences in retail electricity price regulation to attempt to isolate the role that carbon cost pass-through plays in indirectly induced innovation.[4] In particular,

[2]I highlight features of the ETS that I think make indirect effects both relevant and estimable. There are plenty of other benefits to studying the ETS, including its size, quality of available data, and the existence of a carefully constructed estimate of the *direct* innovation effect from Calel and Dechezleprêtre (2014) that can act as a reference point.

[3]A consistent but indirect source of evidence of pass-through is a positive link between CO2 prices and stock performance of large, ETS-regulated electric utilities(Veith et al. 2009; Oberndorfer 2009; Bushnell et al. 2013). Such a positive link implies that the carbon market must convey some sort of benefit to those firms, one plausible source of which is cost pass-through.

[4]Within-country or firm-level variation in price data would of course be preferable, in order to isolate the effects of price regulation vs other cross-country differences, such as generation mix. Still, even if pass-through rates were identical, the prevalence of nuclear generation in France should lead to smaller indirect effects in that country. I am currently working on leveraging within-country price variation in Italy due to network congestion that may provide corroborating evidence.

I contrast the mostly market-based pricing in Germany with the still high degree of retail price regulation in France. If carbon cost pass-through does indeed induce innovation among unregulated firms, there should be a stronger response in Germany than in France.[5]

In the context of the European carbon market, my estimates suggest that the total effect of environmental regulation on innovation is at least twice as large as the direct effect on regulated firms alone. As such, consideration of indirect effects may be important for accurate estimation of effects of environmental policy on innovation. As suspected, even after accounting for these indirect effects, I find that the total effect of the policy on low-carbon innovation remains small in absolute terms. However, the primary purpose of that empirical exercise and this paper as a whole is to illustrate relative effect sizes. As I show theoretically, the indirect effects of policy on innovation by unregulated firms are likely to scale with the size of direct effects on regulated firms. As such, if the total policy effect can be twice the size of the direct effect, in settings where the direct effect of a policy is larger, the bias of studies that consider only direct effects may be economically significant.

These indirect effects are likely to be present for many environmental regulations, since the channels through which indirect effects operate are common to many such policies. For example, cost pass-through has been raised as an issue in the context of a number of other environmental policies, including but not limited to NO_x reductions in the United States (Burtraw et al. 2001), standards for management of livestock waste (Vukina 2003), and air pollution regulation under the Clean Air Act (Gianessi et al. 1979). Similarly, fully regulated utilities may have a cost-pass through allowance for new investments needed to comply with stricter environmental regulation, such as water quality standards (Cowan 1993). As such, the potential for standard estimators to produce biased estimates of (total) induced innovation extends to a potentially large number of policies.

The magnitude of these indirect effects of any one policy on innovation will, of course, depend upon the strength of the hypothesized channels. For example, when regulated firms are able to pass through costs because of a lack of unregulated competition, the effects are likely to be larger. In contrast, sectors whose competitiveness may be threatened by regulation due to the presence of competition from an unregulated sector (e.g. foreign firms) are less likely to pass through costs, providing fewer incentives for innovation by consumers of that sector's output. Similarly, the strength of intellectual property rights in the jurisdiction of a particular environmental policy will impact the degree of knowledge spillovers, which will

[5] In a broader sense, Nesta et al. (2014) also consider the interaction between market liberalization and policies to promote renewable energy in the stimulation of renewable energy technology, but their focus is again on firms subjected directly to an environmental policy. Their argument for more innovation in liberalized markets is based on entry of new energy producers, whereas here, I focus on innovation by downstream users in response to higher prices resulting from pass-through of emissions costs.

influence the way in which higher patent output by regulated firms affects the productivity of R&D for unregulated firms.

The remainder of this paper is organized as follows. The next section more formally presents the standard approaches taken in estimating induced innovation, highlighting the prevalence of the assumption that unregulated firms do not innovate in response to regulation, and the consequences when that assumption is violated. The third section uses simple models to illustrate two common channels through which regulation is likely to have indirect effects on innovation by unregulated firms: knowledge spillovers and cost pass-through. The fourth section presents the empirical application to the EU ETS, including development of an alternative estimation approach using dynamic count models that account for both indirect channels, as well as a description of the data used for estimation. Section five presents the results of estimation, and the final section concludes.

2 Estimation of policy effects and potential bias

To illustrate why innovation by unregulated firms in the EU or elsewhere might pose a problem for assessment of policy effects, it is useful to review the objectives of and assumptions underlying most studies of induced innovation. Many empirical studies of environmental policy effects, including those on innovation, seek to quantify the effect of regulation on a particular outcome of interest. From a policy perspective, the relevant outcome of interest is typically an aggregate quantity, such as total low-carbon patenting across all firms. To get at that total policy effect (TPE), most studies first estimate firm-level average responses, such as the average treatment effect (ATE) or average treatment effect on the treated (ATT), then sum those effects across regulated firms. There are two potential problems with this approach when unregulated firms respond to policy, which may result in biased estimates of the TPE. The first is that unregulated firms do not form a suitable counterfactual for regulated firm behavior in the absence of a policy. The second is that the responses by unregulated firms go uncounted. This section briefly illustrates these issues in the treatment effects framework introduced by Rubin (1974).

To fix ideas, suppose that the outcome of interest is aggregate innovation in a particular technological field and the total policy effect of interest is the change in such innovation (as measured by patents) caused by new environmental regulation. The TPE can theoretically be calculated by comparing the total patenting under the actual policy regime with total patenting in the absence of the policy. The key challenge in such studies is construction of a credible counterfactual: how many patents would have been generated if the policy had not been implemented?

To answer that question, the treatment effects literature makes use of the idea of potential outcomes. Suppose that firm i's regulatory status is $T_i \in \{0,1\}$, with $T_i = 1$ if firm i is subject to environmental regulation and $T_i = 0$ if not. Let $R \in \{0,1\}$ be a binary indicator of whether or not any firms are regulated, i.e. $R = \mathbb{1}(\sum_i T_i > 0)$.[6] Finally, let patenting by firm i with regulatory status T within overall policy regime R be denoted y_{iTR}. For example, i's patent output if regulated under a policy is y_{i11}, while that same firm would output y_{i00} patents if it not regulated and no other firms were regulated.

This notation permits a concise expression of the total policy effect. The total policy effect is the change in expected total patenting caused by the introduction of regulation. Note that the introduction of regulation changes R from 0 to 1 for all firms, while it changes T_i from 0 to 1 only for firms directly subject to regulation. As a result, we may write the TPE as

$$TPE = \sum_i y_{iT1} - y_{i00}, \tag{1}$$

where the sum is over all firms that might be affected in some way by the policy, not just firms directly subject to regulation.[7] For each firm, the TPE compares patenting given a firm's actual treatment status and the presence of regulation (y_{iT1}) with what that firm's patenting output would have been in if it were not regulated and no other firm were regulated (y_{i00}). For firms actually regulated, this difference corresponds to $y_{i11} - y_{i00}$, while for unregulated firms it corresponds to $y_{i01} - y_{i00}$.

The principal practical challenge in estimating the TPE is that we only observe one outcome for each firm at a given point in time. When regulation is in place, we observe y_{i11} for regulated firms, but we do not observe y_{i01} or y_{i00}, both of which enter into (3). Similarly, for unregulated firms, we do not observe y_{j00}. Since we cannot estimate such unobserved outcomes for a single individual, a natural approach is to replace firm-level potential outcomes with their population averages, frequently after conditioning on covariates, e.g. to replace y_{iTR} with $E[y_{iTR}|x_i]$. Here, the expectation is be taken over unobserved components

[6]This could be extended to allow T_i, R, or both to be continuous variables measuring the intensity of treatment/regulation. More generally, R could even be an N-dimensional vector capturing the treatment status for all firms. For practical purposes, and looking ahead to the empirical application to the EU ETS, the largest firms in covered sectors are regulated, such that a binary indicator for R should be a useful approximation. I use a binary treatment T at the firm-level for comparison with existing studies.

[7]This sum should include all entities that could potentially respond to the policy in some way. This may include individuals, or universities, and may include firms outside of the jurisdiction of the policy in question. For practical purposes, in the empirical application to the EU ETS I consider the universe of responding entities to be those falling within a particular country, regardless of regulatory status.

of the potential outcome. Making these replacements yields:

$$TPE = \sum_i E[y_{iT1} - y_{i00}|x_i] \tag{2}$$

where, to save on notation, I simply re-define TPE to refer to this new expression.

The total policy effect can be usefully decomposed into direct and indirect effects on both regulated and unregulated firms. In particular, we may rewrite (2) as:

$$TPE = \underbrace{\sum_{i:T_i=1} E[y_{i11} - y_{i01}|x_i]}_{\substack{\text{Direct Effect on} \\ \text{Regulated Firms} \\ \gamma^R_{Direct}}} + \underbrace{\sum_{i:T_i=1} E[y_{i01} - y_{i00}|x_i]}_{\substack{\text{Indirect Effect on} \\ \text{Regulated Firms} \\ \gamma^R_{Indirect}}} + \underbrace{\sum_{j:T_j=0} E[y_{j01} - y_{j00}|x_j]}_{\substack{\text{Indirect Effect on} \\ \text{Unregulated Firms} \\ \gamma^U_{Indirect}}}. \tag{3}$$

The total policy effect thus consists of three components, corresponding to the three sums above. The first (γ^R_{Direct}) is the direct effect of regulation on regulated firms: what are the consequences of regulating a firm given that some regulation of other firms is in place (i.e. holding R fixed at 1)? The second term ($\gamma^R_{Indirect}$) is the indirect effect of regulation on regulated firms, capturing ways in which decisions made by other regulated firms might impact firm i even if firm i were not regulated. The final term ($\gamma^U_{Indirect}$) captures a similar indirect effect on firms that are actually unregulated. This decomposition of the TPE into direct and indirect effects is similar to that in Hudgens and Halloran (2008); I further decompose the indirect effects into those on regulated firms and those on unregulated firms. Doing so emphasizes the two potential pitfalls of focusing only on direct effects: part of the response by regulated firms goes uncounted, and any response by unregulated firms is ignored.

From (3) we can begin to see how standard estimators relate to the total policy effect. Most studies of induced innovation seek to quantify the Average Treatment Effect (ATE) or Average Treatment Effect on the Treated (ATT) (e.g. Lanjouw and Mody 1996; Jaffe and Palmer 1997; Brunnermeier and Cohen 2003). In what follows I focus on estimators of the ATE, but similar logic applies for the ATT.[8] The ATE conditioned on x_i is defined as:

$$ATE(x_i) = E[y_{i11} - y_{i01}|x_i], \tag{4}$$

which corresponds precisely to the summand in the first sum in (3).[9] As such, estimators

[8]Specifically, let T_i be an element of x_i. Then the first expectation in (3) is conditioned on $T_i = 1$.

[9]Technically, ATE definitions often are written under SUTVA and thus do not include the R subscript on potential outcomes, and would simply be written as $E[y_{i1} - y_{i0}|x_i]$. Practically speaking, though, most ATE estimators are implemented to estimate (4) since they use outcomes for unregulated firms in the presence of regulation to estimate $E[y_{i0}|x_i]$.

that seek to estimate the ATE are useful in estimating the direct effect γ^R_{Direct} of a policy on the outcomes of regulated firms. To make progress, it is common to assume that potential outcomes are conditionally mean-independent of a firm's regulatory status given some vector of covariates x. In particular, this implies that if i is regulated, j is unregulated, and $x_i = x_j$, then $E[y_{i01}|x_i] = E[y_{j01}|x_j]$. In other words, after conditioning on covariates, the expected outcome a firm would have if it were not regulated but some regulation were in place is the same for all firms, regardless of their actual regulatory status. That assumption allows unbiased estimation of $ATE(x_i)$, and by summing over regulated firms, unbiased estimation of γ^R_{Direct}. Regression and matching estimators, including difference-in-difference and propensity score approaches, take this general approach to estimating the ATE.

Since estimators of the ATE quantify the *direct* effects of policy, a key question is whether those same estimators provide unbiased estimates of the *total* policy effect. Mathematically, we wish to know when $TPE = \gamma^R_{Direct}$. From (3) a sufficient condition for $TPE = \gamma^R_{Direct}$ is for there to be no indirect effects of any kind. That condition corresponds to the Stable Unit Treatment Value Assumption (SUTVA), under which the regulatory status of other firms has no effect on firm i's patent output. Thus, if SUTVA holds, average treatment effect estimators provide an unbiased estimate of the total policy effect (when scaled by the number of regulated firms). If SUTVA does not hold, an estimator of the direct effect may still provide an unbiased estimate of the total policy effect, but only in the extreme case where the indirect effects on regulated and unregulated firms exactly cancel. A secondary issue arising from the presence of indirect effects is that the total effect of the policy on regulated firms, $TPE^R = \gamma^R_{Direct} + \gamma^R_{Indirect}$, cannot be derived from the ATE alone. Estimators of the ATE only capture the differential (direct) effect of a policy on regulated vs unregulated firms, which may not be the same as the total effect on regulated firms.

By this logic, prior studies of regulation-induced innovation that estimate only direct effects on regulated firms will provide biased estimates of the total policy effect unless the indirect policy effects are zero. Those studies may still provide unbiased estimates of the direct effects they set out to quantify, but they offer an incomplete picture of the total innovation effects of environmental policy. Further, as I will illustrate theoretically, in many environmental policy settings, SUTVA is indeed likely to fail: both cost pass-through and knowledge spillovers give rise to indirect effects of regulation on innovation by regulated firms, unregulated firms, or both.

The direction of bias in standard estimators depends upon the sign of the indirect effects. If both of those indirect effect terms are positive, estimators corresponding to the direct effect alone will underestimate the total policy effect. Both indirect effects that I consider are likely to be positive because of positive external innovation incentives stemming from higher prices of intermediate goods and knowledge spillovers. As a result, prior studies are likely to under-estimate the patenting effects of many environmental policies. To formalize this intuition, I next develop a simple theoretical framework that incorporates cost pass-through and spillovers into a model of induced innovation.

3 Theoretical framework

Consider a two-sector economy in which one of the sectors is subject to environmental regulation, and that sector produces a good used by the unregulated sector. The N^R regulated firms use a number of primary inputs to produce a single output good X^R, and in the process, generate a harmful pollutant e. The regulated sector is subject to an emissions permit scheme with permit price τ, which regulated firms take as given.[10] The N^U unregulated firms use the regulated sector's output as an input to production of a final good X^U. A simple example of this is a regulated sector producing electricity and an unregulated sector using electricity to produce textiles.

In this two-period model, firms have two decisions to make. In the second period, a firm must choose how much of its output to produce, with firms competing in Cournot fashion. In the first period, each firm may invest in R&D to alter its production technology for the second period. In particular, regulated firms may invest in R&D that lowers emissions, while unregulated firms may invest in R&D that reduces the use of X^R. R&D is costly, and to maintain focus, considered to be deterministic.

The primary purpose of the model is to demonstrate how the presence of τ alters R&D by unregulated firms through the indirect channels of knowledge spillovers and cost pass-through. For a related treatment of R&D without such indirect channels, focusing on the effects of τ on R&D by regulated firms only, see Baker and Shittu (2006). For a general theoretical treatment of cost pass-through (tax incidence) without discussion of R&D, see Weyl and Fabinger (2013). Knowledge spillovers have been studied extensively in the broader R&D literature; see, for example, Leahy and Neary (1997). The model here combines elements of all of these approaches to highlight the role that indirect effects may play in R&D responses to environmental policy.

I begin by considering firms' second-period output decisions, taking first period knowledge investment decisions (and thus second period knowledge stocks) as given. The analysis then turns to the first period R&D decisions, effectively working backwards.

[10] Even if regulated firms have market power in their output market, imagine that the permit scheme covers other firms that are economically and technologically disconnected from the two sectors considered here.

3.1 Output decisions

Let profits for representative regulated and unregulated firms be written as π^R and π^U, with

$$\pi^R = P^R\left(X^R\right)x_i^R - c^R(x_i^R, w) - \tau e(x_i^R, k_i^R), \tag{5}$$

$$\pi^U = P^U\left(X^U\right)x_j^U - c^U(x_j^U, w) - w^R D^R(x_j^U, k_j^U). \tag{6}$$

Here, $P^S(X^S)$, $S \in \{R, U\}$ is the inverse demand facing a firm in sector S, x_i^S is the output quantity chosen by that firm (with sector-wide total output $X^S = \sum_i x_i^S$), and production costs for firms in sector S are $c^S(x_i^S, w)$. Factor prices are denoted by w, with the endogenously determined price of the regulated firm's output separated out as w^R. Emissions $e(x_i^R, k_i^R)$ depend upon both output and the firm's knowledge k_i^R. Unregulated firms do not emit, but use production technology with factor demand for X^R denoted by $D^R(x_i^U, k_i^U)$, which depends on the unregulated firm's output x_i^U and knowledge k_i^U. I make standard assumptions that demand is non-negative and downward sloping ($P^{U'} < 0$, $P^U(X^U) \geq 0$) and costs are increasing and convex in output ($c_x^S > 0, c_{xx}^S > 0$). Emissions are increasing in output ($e_x > 0$) and decreasing and convex in knowledge ($e_k < 0, e_{kk} > 0$), and marginal emissions are decreasing in knowledge ($e_{xk} < 0$). Finally, the production technology for unregulated firms is such that demand for the regulated firm output is increasing in output ($D_x^R > 0$) and decreasing and convex in knowledge ($D_k^R < 0, D_{kk}^R > 0$), and marginal demand for regulated firm output is decreasing in knowledge ($D_{xk}^R < 0$).

Both regulated and unregulated firms choose output to maximize profits, taking prices, knowledge, and actions of other firms as given. Equilibrium output choices (assumed interior) are defined by the first order conditions for maximization of (5) and (6). Specifically:

$$P^{R'}(X^R)x_i^{R*} + P^R(X^R) = c_x^R(x_i^{R*}, w) + \tau e_x(x_i^{R*}, k_i^R), \tag{7}$$

$$P^{U'}(X^U)x_j^{U*} + P^U(X^U) = c_x^U(x_j^{U*}, w) + w^R D_x^R(x_j^{U*}, k_j^U). \tag{8}$$

Let the equilibrium output of an unregulated firm implicitly defined by those first order conditions be denoted x_j^{U*}. Focusing on symmetric equilibrium, total demand for X^R facing regulated firms is then $N^U D^R(x_j^{U*}, k_j^U)$, and the inverse demand function $P^R(\cdot)$ is defined such that $P^R(N^U D^R(x_j^{U*}, k_j^U)) = w^R$. Similarly, first order conditions for regulated firms yield implicitly defined output choices x_i^{R*}.

As in Weyl and Fabinger (2013), the first order conditions defining equilibrium output by regulated firms can be implicitly differentiated to yield the rate at which regulated firms pass-through the costs of regulation (τ) into the price of their output (w^R). In particular, it can be shown (see Appendix) that

$$\frac{\partial w^R}{\partial \tau} = P^{R'} N^R \frac{e_x}{N^R P^{R''} x_i^{R*} + (N^R + 1) P^{R\prime} - c^R_{xx} - \tau e_{xx}}, \tag{9}$$

where all derivatives are evaluated at equilibrium output levels. The fraction corresponds to the change in a firm's equilibrium output in response to a marginal increase in the tax, which, when multiplied by $P^{R'}(X^{R*})N^R$, yields the resulting change in the price w^R. That change in output is determined by the relative size of the two effects a tax increase has on the marginal profit of output. First, holding output constant, a higher tax directly alters marginal profits in proportion to marginal emissions, as captured in the numerator. Second, a tax increase will cause all firms to alter their output decisions, which will affect both prices and the firm's own marginal costs, as represented by the denominator. The ratio of the direct tax effect and the output effect determines how much output will change in response to a tax increase.

The pass through effect described in (9) will, in general, depend on both market characteristics and technology, as is often discussed in the literature on tax incidence. Note in particular that market power is not required for cost pass-through to occur. As the market becomes competitive, i.e. $N \to \infty$ and consequently $x_i^{R*} \to 0$, the pass-through rate from (9) converges to a constant e_x, which depends upon regulated firm production technologies (often in the literature unit choices and technology with constant emissions per unit output imply $e_x = 1$). If there were a third, unregulated production sector competing with regulated firms, then prospects for pass-through begin to erode. In such a case, $P^{R'}$ could be seen as a residual demand facing the regulated sector, and the larger that competing sector is, the flatter that residual demand would be ($P^{R'} \to 0$), driving pass-through to zero. As such, the regulated sector as a whole must possess power in its output market, but firms within that sector need not possess market power for cost pass-through to occur.

Provided that the equilibrium is stable in the sense of Seade (1980),[11] $\frac{\partial w^R}{\partial \tau} > 0$ and the effect of a tax increase on the price of regulated firm output is positive (see Appendix). Thus, for a given level of knowledge k_i^R, regulated firms facing a cost increase via τ will choose output such that w^R increases, and unregulated firms will thus face a higher relative price for X^R. It is intuitive that the higher price may affect R&D decisions for unregulated firms in earlier periods; I turn to this next.

[11]Essentially, reaction functions must be downward sloping.

3.2 Innovation decisions

In the first period, firms choose R&D investment with a goal of maximizing the sum of profits across the two periods. I ignore discounting given the short time horizon of the model; adding it offers little additional insight. R&D by a firm in sector S is denoted r^S, with each unit of R&D having unit cost. Such investment generates new knowledge according to knowledge production function $f^S(r_i^S, R^S, R^{-S}, k_i)$, where R^S captures total investment by other firms in i's sector, R^{-S} is total investment by firms in the other sector, and k_i is the firm's current knowledge stock. Knowledge production is assumed to be increasing and concave in both own R&D ($f_r^S > 0, f_{rr}^S < 0$) and own knowledge ($f_k^S > 0, f_{kk}^S < 0$). In general, $f^S(\cdot)$ may either increase or decrease in both R^S and R^{-S}. An increasing relationship reflects positive R&D spillovers ($f_R^S > 0$), while a decreasing one may reflect technological saturation ($f_R^S < 0$). The theory below does not depend on the sign of f_R^S, but when making predictions regarding the direction of expected effects for the empirical application, I assume $f_R^S > 0$ since the low-carbon technologies of interest are not yet mature enough that saturation should be an issue. New knowledge generated via these functions increments the knowledge stock additively:

$$k_i^S = k_{i0}^S + f^S(r^S, R^S, R^{-S}, k_{i0}^S). \tag{10}$$

With this notation, the first-period problems facing firms can be written as follows:

$$\max_{r_i^R} \left\{-r_i^R + \pi^{R*}(k_i^R, w, \tau, K^R, K^U)\right\} \tag{11}$$

$$\max_{r_j^U} \left\{-r_j^U + \pi^{U*}(k_i^U, w, \tau, K^U, K^R)\right\}. \tag{12}$$

Here, π^{S*} represents the equilibrium payoffs a firm in sector S will earn in the second period. Those future payoffs will depend upon both future prices (including the emissions price) and future knowledge stocks of all firms (K^R, K^U), which are affected by current period investments in R&D via (10). When choosing R&D, each firm takes the R&D decisions of all other firms as given. Further, consistent with the price-taking behavior described earlier, unregulated firms assume that their R&D will have no effect on future prices. However, unregulated firms do recognize that their R&D decisions may affect future output choices by other unregulated firms, due both to knowledge spillovers and the strategic nature of second period output decisions.

First order conditions corresponding to the problems above again implicitly define equilibrium firm choices r_i^{R*} and r_j^{U*}. I focus solely on first period R&D choices by unregulated firms, since the purpose of the theoretical framework is to highlight channels for induced innovation outside the regulated sector. Still, due to the parallel structure of (11) and (12), the analysis for regulated firms is analogous. Adapting the methods in Leahy and Neary (1997) (see Appendix), in a symmetric equilibrium the first order condition defining r_j^{U*} can be written:

$$-w^R D_k^R f_r^U \left(\underbrace{1}_{\text{R\&D cost effect}} + \underbrace{(N^U-1)P' x_j^{U*} \frac{\pi_{jj}^{Uj}\left(\frac{\pi_{lj}^{Uj}}{\pi_{jj}^{Uj}} - \frac{f_{R^U}^U}{f_r^U}\right)}{\left(\pi_{jj}^{Uj} - \pi_{lj}^{Uj}\right)\left(\pi_{jj}^{Uj} + (N^U-1)\pi_{lj}^{Uj}\right)}}_{\text{R\&D strategic effect}} \right) = 1, \quad (13)$$

where the notation π_{lj}^{Uj} is used to denote a second order partial derivative of firm j's profits with respect to output decisions by unregulated firms j and l. This condition equates the marginal benefit of R&D, which consists of a (direct) cost effect and a strategic effect via impacts to other firms' output decisions, with the (constant) marginal cost of R&D. Note that the sign of the strategic effect on optimal R&D depends on the strength of knowledge spillovers, captured here by $f_{R^U}^U$. Once $f_{R^U}^U$ grows larger than a threshold level, the strategic effect changes sign. The intuition is straightforward: if spillovers are small, R&D enables j to increase output, causing a strategic reduction in output by other firms, and providing additional marginal benefits of R&D to firm j. As spillovers get larger, however, R&D by j also reduces costs for other firms, causing them to increase output and reduce the marginal benefits to firm j. Above a spillover threshold, the net impact of those effects is negative.

This condition provides useful insight into the determinants of R&D by unregulated firms in equilibrium. First, and of greatest interest, the price of regulated firm output will clearly impact R&D decisions. Thus, any factor affecting the price of that good will, in turn, impact R&D by unregulated firms. In particular, factor prices w, pass through of τ into w^R, and R&D by regulated firms R^R (which affects regulated firm output X^R) will all impact w^R, and ultimately R&D by unregulated firms. Second, through $f_{R^U}^U$ and f_r^U, R&D by one unregulated firm is likely to depend upon R&D choices by other unregulated firms R^U. Thus, determinants of R^U and R^R, including the knowledge stocks of other firms K^U and K^R, will also affect the equilibrium value of r_j^{U*}. Finally, the marginal benefits of R&D scale with D_k^R, which corresponds to the (negative of the) marginal rate of technical substitution (MRTS) between X^R and k in an unregulated firm's production of its output X^U. If knowledge is not a good substitute for X^R, then D_k^R is small in magnitude, and regardless of the magnitude of cost pass-through, the marginal benefits of R&D will be small, and we should expect to see little R&D in that unregulated sector. The MRTS will be small in magnitude if past investment in R&D is high (so that k is large) since $D_k^R < 0$ and $D_{kk}^R > 0$. Similarly, the

MRTS will be small if little or no X^R is required in the production of X^U in the first place. That is, for a given level of pass-through (and hence a fixed w^R), we should expect to see the most innovation in a sector that is intensive in its use of X^R and immature with respect to technological advances to reduce use of X^R.

Determining the sign of the effect of τ on r_j^{U*} becomes quickly intractable: τ impacts the price w^R, R&D by regulated firms and other unregulated firms, and output choices of all unregulated firms, all of which affect r_j^{U*} via (13). Some additional insight can be gained if the unregulated sector is a monopoly ($N^U = 1$), in which case the strategic effect of R&D disappears ($N^U - 1 = 0$).[12] In that case, differentiation of (13) with respect to τ yields

$$\frac{\partial r_j^{U*}}{\partial \tau} = \frac{\partial w^R}{\partial \tau}\frac{\partial r_j^{U*}}{\partial w^R} = \frac{\partial w^R}{\partial \tau}\frac{f_r^U}{w^R}\frac{\left(-D_k^R - w^R D_{kx}^R \frac{\partial x_j^{U*}}{\partial w^R}\right)}{D_k^R f_{rr}^U + D_{kk}^R (f_r^U)^2}. \tag{14}$$

Casual inspection of this expression suggests that the change in R&D by unregulated firms due to regulation will depend upon the pass-through of the permit price $\frac{\partial w^R}{\partial \tau}$. Any market characteristics or policies (e.g. price regulation) that limit cost pass-through should therefore reduce the level of indirectly induced innovation.

This sign of the effect in (14) hinges upon the numerator of the final fraction, which captures the two effects of an increase in the price w_2^R on the marginal benefit of knowledge. The first effect ($-D_k^R$) makes marginal knowledge more valuable, since knowledge acts to reduce the use of X^R and thus lowers costs. The second effect ($-w^R D_{kx}^R \frac{\partial x_j^{U*}}{\partial w^R}$) concerns how contraction of output in response to a change in w^R might change the MRTS between k and X^R. The sign of $\frac{\partial r_j^{U*}}{\partial \tau}$ will depend on the relative magnitude of those two effects. If the MRTS is constant with respect to output, $D_{kx}^R = 0$ so the second effect disappears, and the first effect is positive, so that R&D by unregulated firms increases with τ. But in general, the magnitude and sign will depend on the production technology used by unregulated firms.

Pulling this all together, we may write equilibrium R&D by an unregulated firm as a function of variables that firm treats as exogenous (w^R is endogenous, but unregulated firms are price takers):

$$r_j^{U*} = r^{U*}(w^R, w, \tau, k_{i0}^U, K_0^U, K_0^R). \tag{15}$$

These dependencies will be important in translating the theoretical framework to an empirical specification for the application to the EU ETS. As noted above, it is difficult to definitively sign the effect of changes in determinants of R&D on equilibrium levels of innovation without further restrictions on the primitive functions defining production of both

[12] The strategic term also vanishes in perfect competition, since $x_j^{U*} \to 0$.

output and knowledge.

I conclude this section by summarizing the expectations about what will impact innovation by unregulated firms based on the discussion above. The first two hypotheses follow directly from assumptions about the properties of the knowledge production function $f^U(\cdot)$:

H1 *Innovation by unregulated firms increases in own knowledge.*

H2 *Innovation by unregulated firms increases in total knowledge for newer technologies.*

While these knowledge effects do not directly pertain to unregulated firms' response to regulation, they apply in general and thus will mediate some of the indirect effects of regulation on unregulated firms.

The remaining hypotheses address unregulated firm responses to regulation more directly. First, while the sign of $\frac{\partial r_j^{U*}}{\partial \tau}$ cannot be determined in general, for a variety of production technologies:

H3 *Innovation by unregulated firms increases with higher costs of inputs produced by regulated firms.*

Further, the discussion of condition (13) includes additional hypotheses about market characteristics and technology. Specifically:

H4 *Innovation by unregulated firms is higher in sectors that are intensive in their use of regulated firm output and are technologically immature.*

H5 *Innovation by unregulated firms is higher in sectors facing higher cost pass-through from regulated firms (larger $\frac{\partial w^R}{\partial \tau}$).*

To test these hypotheses, I turn to an empirical application involving low-carbon innovation and the emissions trading program in the EU.

4 Empirical Application: EU ETS

The European Union launched its permit market for greenhouse gas emissions in 2005 covering a subset of CO2 emitting installations across several industries; currently over 11,000 installations in 31 countries are regulated under the scheme.[13] As the largest emissions permit market in the world, the EU ETS is under intense scrutiny for its ability to meet a variety of objectives. One important metric on which the EU ETS is being evaluated is its ability to stimulate low-carbon innovation, both in patentable product form and process innovations.

[13]For detailed background on the ETS, the reader is referred to any of the prior studies referenced earlier.

As discussed earlier, studies to date have found little evidence of a strong impact of the EU ETS on product innovation. In this empirical application, I revisit those findings in light of the potential for indirect effects as outlined in the theoretical framework.

While cost pass-through is possible for any output produced by regulated firms, I focus only on electricity, given its prevalence as an input used by unregulated firms and existing documentation of cost pass-through. Further, cursory examination of the growth in different types of low-carbon patents (Figure 2) suggests that electricity-related patents are growing at a faster rate than other types of low-carbon patents since 2004, and thus may be a good place to look for indirectly induced innovation. Finally, from an empirical perspective, electricity markets offer two sources of variation that will prove useful in identification of cost pass-through as a driver of indirectly induced innovation. Specifically, power producers face different emissions costs due to the use of different technologies, and geographic variation in price regulation and market characteristics will influence those firms' ability to pass through emission costs into the price of electricity. Although the EU as a whole has made strides toward market-based pricing for electricity, there is still substantial variation in scope for carbon price pass-through all the way to end users. Such variation in cost pass-through should, in turn, influence the strength of the hypothesized indirect effects of regulation on innovation.

Based on that observation, this paper focuses on two countries with quite different electricity markets and price regulations that give rise to different pass-through opportunities: Germany and France.[14] To summarize the regulatory difference of interest, Germany's electricity market is relatively more liberalized, leaving firms more exposed to variation in electricity prices, including pass-through of carbon prices. In contrast, many French firms have the option (and most take it) to buy electricity under a regulated tariff scheme, such that even if carbon prices are passed through to wholesale prices, retail prices will be relatively unaffected.[15] Further, differences in the generation mix across countries (i.e. the nuclear capacity in France) should lead to higher regulation-caused electricity price increases in Germany, even if the pass-through rates were identical. In short, both regulatory and technological differences should lead to higher incidence of indirectly induced innovation in Germany.

In addition to their contrasting technological mixes and degrees of price regulation, Ger-

[14] I have also begun investigating within-country, regional variation within Italy, which would help rule out other cross-country differences that could drive the results here. Congestion in the Italian transmission networks frequently allows prices to vary regionally, and differences in the generation mix (i.e. carbon intensity) across region should lead to different pass-through prospects by region. In principle, this should lead to different levels of indirect effect by region.

[15] That generalization comes with an important caveat: the largest electricity-consuming firms in France mostly buy electricity through market offers much as in Germany, but the bulk of small to medium-sized firms and household consumers are part of the regulated tariff system.

many and France possess other characteristics that are useful for this study. The two countries lead the EU in terms of electricity production, and Germany leads countries in low-carbon patenting at the European Patent Office (EPO) while France ranks 3rd among EU nations and 5th overall (Figure 3). The combination of these characteristics and the different degree of retail price regulation in the two countries should lead to a detectable difference in indirect innovation effects of the EU ETS. In particular, we should see a stronger positive impact of the EU ETS on low-carbon patenting by unregulated firms in Germany than we do for such firms in France. There still may be an indirect effect of the EU ETS on innovation by unregulated firms in France due to both the fact that some customers do buy electricity on open markets in France and the potential for knowledge spillovers from regulated firms.

To operationalize these ideas, I next detail my estimation strategy.

4.1 Empirical Model

We can use the theoretical framework outlined earlier to define an estimation strategy for studying the effects of the EU ETS on low-carbon patenting. In particular, the knowledge production functions from (10), evaluated at equilibrium levels of output and investment, provide the basis for estimation. In light of the discrete nature of patent counts generated via those knowledge production functions, all models considered are count models. The preferred model uses the pre-sample mean estimator introduced by Blundell et al. (1995), which uses pre-sample information on patenting by firms in the data set to proxy for unobserved heterogeneity. I also contrast this approach with a standard fixed effects poisson model (estimated via concentrated maximum likelihood) and a simple pooled poisson estimator that ignores unobserved heterogeneity. All of the estimated models include both direct and indirect effects from (3). The latter effect is identified through variation in the carbon price and patenting over time, relying upon inclusion of other time-varying variables to control for other factors which might include patenting output.

To make these count models operational, it is necessary to specify functional forms for the knowledge production functions $f^R(\cdot)$ and $f^U(\cdot)$. I assume both take a Cobb-Douglas form, as is common in many models of innovation (e.g. Griliches 1979; Popp 2002). Evaluating the knowledge production function at equilibrium levels of R&D as defined by (13) yields

$$pats_{it}^S = \alpha_i^0 \alpha_{ct} (r_{it}^{S*})^{\alpha_r^S} (R_t^{S*})^{\alpha_S^S} (R_t^{-S*})^{\alpha_{-S}^S} k_{it}^{\alpha_k^S} \nu_{it}, \tag{16}$$

where $\nu_{it} \geq 0$ is a shock to patenting productivity. Recall that R_t^{S*} captures R&D by other firms within i's sector, while R_t^{-S*} captures R&D by firms outside i's sector. The α parameters capture various aspects of innovative productivity: α_i^0 represents firm-specific productivity, α_{ct} is a time-varying patenting productivity parameter that varies at the country level c,[16] α_r^S is the productivity of own R&D, α_S^S is the productivity of within-sector other R&D, and α_{-S}^S is the productivity of cross-sector other R&D. Thus, α_S^S and α_{-S}^S capture spillovers within and across sectors. The superscripts on coefficients indicate that the effects of determinants on patent output may vary with regulatory status.

This equation can be re-expressed in a form more familiar in count models with exponential mean:

$$pats_{it}^S = \alpha_i^0 \alpha_{ct} exp\left(\alpha_r^S ln(r_{it}^{S*}) + \alpha_S^S ln(R_t^{S*}) + \alpha_{-S}^S ln(R_t^{-S*}) + \alpha_k^S ln(k_{it})\right) \nu_{it}, \tag{17}$$

Since R&D decisions are endogenous, I replace $ln(r_{it}^{S*})$, $ln(R_t^{S*})$, and $ln(R_t^{-S*})$ with predicted determinants of R&D based on (13) and (15) and the analogous condition for regulated firms. Those determinants include factor prices w, the price of regulated firm output w^R, own knowledge k_i, and knowledge stocks from other firms, K. Further, to reflect the fact that equilibrium and R&D in the regulated sector will differ once regulation is introduced, I allow R&D investment to depend upon a treatment status dummy variable T_{it}, which takes the value one if a firm is in the regulated sector and the policy is active, and zero otherwise (i.e. $T_{it} = \mathbb{1}(S_i = R \ \& \ t \geq \tilde{t})$, where $\tilde{t}$ is the first year the policy is in effect). In particular, I assume

$$\begin{aligned} ln(r_{it}^{S*}) =& \theta_i^0 + \theta_R T_{it-1} + \sum_{m \in M} \theta_l^S ln(w_{ct-1}^m) + \theta_{elec}^S ln(w_{lct-1}^{elec}) \\ &+ \theta_k ln(k_{it}^S) + \theta_{K^S}^S ln(K_t^S) + \theta_{K^{-S}}^S ln(K_t^{-S}) + \epsilon_{it} \end{aligned} \tag{18}$$

Here, w_{ct-1}^m is the lagged price for commodity m in country c (with the set of all commodities included denoted M). The commodities I focus on for estimation are the prices of coal and natural gas. Both are commonly used as inputs to production processes in conjunction with electricity, and may be thought of more generally as capturing the effect of movements in other input prices. T_{it-1} is a (lagged) dummy variable taking the value 1 if firm i is regulated under the ETS, and 0 if that firm is not regulated. I use one year lags of prices and regulatory status to allow for prices to affect R&D with some delay. K^S and K^{-S} reflect the overall knowledge stocks in the firm's own sector and the other sector. Expressions for $ln(R_t^{S*})$ and $ln(R_t^{-S*})$ can be obtained by summing (18) across firms.

Substituting the expressions for $ln(r_{it}^{S*})$, $ln(R_t^{S*})$, and $ln(R_t^{-S*})$ into (17) and combining

[16] Note a firm may operate in several countries. For the purposes of analysis, a firm operating in multiple countries is treated as multiple firms.

coefficients yields

$$pats_{ict} = \beta_i \beta_{ct} exp(\beta_S T_{it-1} + \beta_k^S ln(k_{it}) + \beta_{K^S}^S ln(K_t^S) + \beta_{K^{-S}}^S ln(K_t^{-S}) + \sum_{m \in M} \beta_m^S ln(w_{ct-1}^m) + \beta_{elec}^S ln(w_{ct-1}^{elec}))u_{it}. \quad (19)$$

Before this equation can be estimated, a few practical considerations must be dealt with. First, knowledge stocks cannot be measured directly, and so I will proxy for a firm's knowledge stock with that firm's past patenting output. At time t a given firm may not have patented prior to t, which introduces an undefined $ln(0)$ into (19). To account for firms with no prior patenting history, I replace $ln(k_{it})$ with a combination of a dummy variable indicating if the firm has not patented in the past, and the log of past patents if the firm has a nonzero patent history (as in Aghion et al. 2012). Second, because one of the variables of interest, electricity price, is only observed at the country by year level, estimating β_{ct} using a standard dummy variable approach would preclude identification of β_{elec}^S. Since β_{elec}^S is a coefficient of interest, I instead proxy for β_{ct} with observed total patenting rates $Pats_{ct}$ at the country level across all technology types. Finally, I assume that the industry knowledge stocks K_t^S and K_t^{-S} are truly public, such that they can be considered a single public knowledge stock K_t. Making these substitutions yields the main model for estimation:

$$\begin{aligned} pats_{ict} = \beta_i exp(\beta_S T_{it-1} + & & \textit{Treatment effect} \\ & \beta_{k,0}^S \mathbb{1}(k_{it} = 0) + \beta_{k,>0}^S \mathbb{1}(k_{it} > 0) ln(k_{it}) + & \textit{Own knowledge effect} \\ & \beta_{K^S}^S ln(K_t) + & \textit{Knowledge spillover effect} \\ & \beta_{elec}^S ln(w_{ct-1}^{elec}) + & \textit{Electricity price effect} \\ & \beta_{Pats}^S ln(Pats_{ct}) + & \\ & \sum_{m \in M} \beta_m^S ln(w_{ct-1}^m)) u_{it}. & (20) \end{aligned}$$

Taking the expectation of equation (20) defines the conditional mean number of patents for firm i, which forms the basis for the three count model variants described earlier. All estimators are variants of quasi-poisson models estimated via quasi-maximum likelihood, with differing approaches to handling unobserved heterogeneity β_i. To relax the strong mean-variance equality assumption in Poisson models, these quasi-poisson models adjust standard errors for over-dispersion.

Note that this specification does not directly include the price of emissions permits τ. The treatment dummy variable T_{it-1} captures a discrete effect of regulation on regulated firms. I use that dummy variable rather than a price for direct comparison with existing studies which use binary treatment status. Further, per the theoretical framework, any indirect effects operate through both knowledge stocks and the electricity price, both of which are included in the model. As a result, the evidence for indirect effects in this model is itself partly indirect, requiring combination of the parameter estimates from (20) and prior evidence of cost pass-through.

To provide a more direct test of the presence of indirect effects, I also estimate a modified version of (20) that includes the permit price. Since one of the key hypothesized channels for indirect effects is pass-through of the emissions price into the electricity price, a natural step would be to replace the electricity price with its determinants, including the emissions price. However, because price data in this model only enter at the country by year level (as opposed to the higher frequency data used in prior pass-through studies), such an approach is likely to provide an imprecise estimate of any indirect effect mediated by cost pass-through.

With this in mind, I instead include both the emissions price and electricity price, but remove other determinants of the electricity price. The coefficient on the emissions price, β_τ^{Tc}, then captures the joint effect of an increase in the permit price and a simultaneous change in other determinants of the electricity price such that the electricity price stays the same. If the emissions price does not impact the electricity price, then no other price determinants must change, and β_τ^{Tc} should be zero. If instead, the permit price impacts the electricity price, then β_τ^{Tc} should be positive, as it effectively increases the price of electricity relative to its other determinants. Because β_τ^{Tc} is allowed to vary by country, this setup permits a direct evaluation of the hypothesis that electricity market differences between Germany and France should lead to greater cost pass-through in Germany, and as a result, more innovation in the unregulated sector in Germany due to those pass-through effects. In other words, I expect $\beta_\tau^{U,Germany} > \beta_\tau^{U,France}$.

The resulting model is

$$
\begin{aligned}
pats_{ict} = \beta_i exp(&\beta_S T_{it-1}+ && \textit{Treatment effect}\\
&\beta_{k,0}^S \mathbb{1}(k_{it}=0) + \beta_{k,>0}^S \mathbb{1}(k_{it}>0) ln(k_{it})+ && \textit{Own knowledge effect}\\
&\beta_K^S ln(K_t)+ && \textit{Knowledge spillover effect}\\
&\beta_{elec}^S ln(w_{ct-1}^{elec})+ && \textit{Electricity price effect}\\
&\beta_\tau^{Sc} \tau_{t-1}+ && \textit{Emissions price effect}\\
&\beta_{Pats}^S ln(Pats_{ct})+ &&\\
&\sum_{m\in M'} \beta_m^S ln(w_{ct-1}^m))u_{it}. && (21)
\end{aligned}
$$

where other determinants of the electricity price have been removed from the set M of factor prices to produce subset M'.

Since the coefficient β_{τ}^{Sc} captures the joint effect of raising the emissions price but lowering other determinants of the electricity price (e.g. coal or natural gas prices) such that the electricity price stays constant, the coefficient on β_{τ}^{Sc} will provide a biased estimate of the impact of a marginal change in the emissions price alone. This can be seen as a form of double counting, so I interpret the estimates as loose upper bounds. Nonetheless, the magnitude of the point estimates are not the primary purpose of estimating this second model; instead, my goal is to test whether the emissions price impact on patenting varies by electricity market (in particular with each market's scope for cost pass-through). If so, there is some evidence that permit prices are affecting patenting indirectly via electricity price pass-through.

While the coefficient estimates from both (20) and (21) provide evidence as to the presence or absence of indirect effects, they do not correspond directly to the direct and indirect policy effects in (3) that are of most interest. With that in mind, I next describe the two approach I use to construct estimates of the direct and indirect policy effects using estimated coefficients.

4.2 Constructing Estimated Direct and Indirect Policy Effects

To construct estimates of the direct and indirect effects of policy in (3), I employ a simulation approach combining my fitted models with estimates of cost pass-through rates from the literature. Because both (20) and (21) are dynamic specifications, computing policy effects requires simulation using each fitted model rather than simply interpreting coefficients. The estimate of β_T, for example, provides insight into the sign and significance of the direct effect of regulation on regulated firms, but the magnitude of that effect depends on a combination of that coefficient and knowledge effects. In particular, to construct estimates for both $E[y_{is1}]$ and $E[y_{i01}]$, it is necessary to compute expected patent output for the first year in which regulation could have impacted patent output, use those fitted outputs to update the knowledge stocks for the second year, and so on. I next describe these simulations in more detail, beginning with computation of direct effects, then describing adaptations to compute indirect effects. The steps are written with reference to (20), but I use the same procedure for (21) except where noted.

4.2.1 Estimates of the direct effect

To estimate the direct effect of regulation on regulated firms (γ_{Direct}^{R}), I perform two related simulations. The first is designed to provide an estimate $\hat{y}_{iT1}$ of $E[y_{iT1}]$, while the second provides an estimate $\hat{y}_{i01}$ of $E[y_{i01}]$.

The steps for producing $\hat{y}_{iT1}$ are as follows:

S1. For the first year of regulation, simulate patent output per firm using (20) and estimated parameters.

S2. Update knowledge stocks for $t+1$ using the simulated patent output in year t.

S3. Repeat steps S1 and S2 for each year of regulation in sequence.

S4. Sum the patent output per firm from S1 across all years of regulation.

Denote the sum produced in $S4$ by $\hat{y}_{iT1}$. These steps compute predicted patent output per firm, and sum over the years of regulation such that $\hat{y}_{iT1}$ represents an estimate of total expected patent output under the actual policy and regulatory status of firms.

To compute $\hat{y}_{i01}$, I simulate a regulated firm's i's patent output if it were not regulated by setting its regulatory status to unregulated and simulating in a similar fashion to above. However, for this simulation, I hold estimated output by other firms $j \neq i$ fixed at the levels simulated during computation of $\hat{y}_{iT1}$. This reflects the fact that the direct effect is intended to capture only the effect of a firm's own regulatory status on its own patent output. This process is then repeated for all regulated firms. In particular, I perform the following steps to compute $\hat{y}_{i01}$:

D1. Simulate patent output under observed regulation according to S1-S4 above.

D2. For firm i, set $T_{it} = 0$ for all t.

D3. For the first year of regulation, simulate patent output for firm i only using the modified data and estimates from (20).

D4. Update firm i's knowledge stock and total knowledge stocks for year $t+1$.

D5. Repeat steps D3 and D4 for each year of regulation in sequence.

D6. Sum the patent output computed in step D3 for each firm across all years of regulation, and denote the sum by $\hat{y}_{i01}$.

D7. Repeat steps D2-D6 for each regulated firm i.

Once estimates of $\hat{y}_{is1}$ and $\hat{y}_{i01}$ have been computed according to the simulations above, the estimate $\hat{\gamma}^R_{Direct}$ of the direct effect γ^R_{Direct} is computed as:

$$\hat{\gamma}^R_{Direct} = \sum_{i:T_i=1} (\hat{y}_{iT1} - \hat{y}_{i01}). \tag{22}$$

4.2.2 Estimates of the indirect effects

Estimation of the indirect effects of the ETS on patenting combines each fitted model with existing estimates of cost pass-through rates from prior studies. Both of the indirect effects in (3) contain the expected change in patenting when a given firm remains unregulated but regulation is introduced for other firms. To construct an estimate $\hat{y}_{i00}$ of the counterfactual patent output in which a firm is unregulated and no regulation exists ($E[y_{i00}]$), I use the estimated model to simulate patenting output when the electricity price is modified according to previously estimated pass-through rates, the regulatory status of all firms is unregulated, and, for the alternative specification (21), the emissions permit price is set to zero.

Letting PTR_c denote the pass-through rate for carbon and electricity prices in the EU ETS, the simulation of the counterfactual proceeds as follows:

I1 Set $T_{it} = 0$ for all i and all t

I2 Modify w_t^{elec} to be $w_t^{ALTelec} = w_t^{elec}/(1 + PTR_c)$ for all t.

I3 If the specification is (21), set $\tau_t = 0$ for all t.

I4 Simulate patent output according to steps S1-S4 above using the modified data, and denote the sum from S4 as $\hat{y}_{i00}$

The simulation procedure produces an estimate $\hat{y}_{i00}$ of $E[y_{i00}]$, the patent output of firm i if no regulation existed. That estimate can be subtracted from estimates of $\hat{y}_{i10}$ and $\hat{y}_{i00}$ computed in S1-S4 and D1-D7. to produce the desired estimates of indirect effects. In particular:

$$\hat{\gamma}^R_{indirect} = \sum_{i:T_{it}=1} (\hat{y}_{i01} - \hat{y}_{i00}) \tag{23}$$

$$\hat{\gamma}^U_{indirect} = \sum_{i:T_{it}=0} (\hat{y}_{i01} - \hat{y}_{i00}) \tag{24}$$

On a final note, interpretation of the estimates $\hat{\gamma}^R_{indirect}$ and $\hat{\gamma}^U_{indirect}$ requires some care. As discussed above, β^{Sc}_{τ} in the alternative model (21) captures a joint effect of simultaneous changes in the permit price and other determinants of the electricity price. As a result, I interpret estimates $\hat{\gamma}^R_{indirect}$ and $\hat{\gamma}^U_{indirect}$ derived from that model as upper bounds. That alternative model is quite useful in pinning down the role that cost pass-through plays in innovation, but I focus my discussion of the estimated magnitude of indirect and total effects on those derived from the main specification (20).

4.3 Data

In line with the specifications presented above, I construct a panel dataset of low-carbon patenting in the EU at the firm level. The study period ranges from 1992 to 2010, covering activity both before the EU ETS (1992-2004) and during (2005-2010). I begin in 1992 based on the availability of price data described below and end in 2010 to avoid well-documented problems of truncation in patent counts (due to delays in processing of patent applications). For the pre-sample mean estimator, I also use patent data going back to 1985.

The primary outcome of interest is low-carbon patenting activity at the European Patent Office (EPO). I obtained patent records from the Worldwide Patent Statistical Database (PATSTAT), with patents pertaining to reduction of carbon emissions identified using the Cooperative Patent Classification (CPC) categories associated with each patent. As in Calel and Dechezleprêtre (2014), I consider any patent labeled with a CPC category beginning with *Y02* to be a low-carbon invention. Within that classification, I also further sub-divide low-carbon patents into those that are related to electricity (CPC categories H01-H05) and those that are not. Knowledge stocks based on these patent counts are simply assumed to be equal to the patenting output in year $t-1$. This is equivalent to assuming complete depreciation of the knowledge stock after one year.[17]

Price data come from several different sources. Emissions allowance prices are from the European Energy Exchange (EEX), and represent the price (in €) of a year-ahead future contract for a permit granting the right to emit one ton of CO2. Coal prices come from the International Energy Agency[18] and represent the annual average per-ton price of Colombian coal. Electricity and natural gas prices are from Eurostat and represent within-year average prices for those commodities by country. The natural gas prices used are those for mid-size industrial consumers.[19]

Regulatory status per firm is based on matching of patent data to the European Union Transaction Log (EUTL) associated with the EU ETS. To match firms, I apply the name harmonization process used in PATSTAT (Magerman et al. 2006) to account operator names from the EUTL. Firms are then matched between datasets using the harmonized name and country from both datasets.

To motivate the empirical analysis, Table 1 provides brief summary statistics of low-carbon, electricity related patent output broken out by country, ETS status, and period

[17]In future iterations, I intend to examine the robustness of the results to this means of constructing a knowledge stock.

[18]Table 4 of http://www.iea.org/media/training/presentations/statisticsmarch/CoalInformation.pdf

[19]Specifically, gas prices are from tables nrg_pc_203 and nrg_pc_203_h. Data from 1991-2007 are from nrg_pc_203_h, and data from 2008-2010 are from nrg_pc_203. Prices are for band I3, which is for customers using between 10,000 and 100,000 GJ.

(1992-2004 vs 2005-2010).

A number of observations emerge from these summary statistics. First, the average number of low-carbon electrical patents per firm per year increases during years in which the ETS is active. This is consistent with the idea of induced innovation. Second, the increase in patenting among unregulated firms in Germany that occurs with the introduction of the ETS is larger than the corresponding change in France, which is consistent with the hypothesis that indirect innovation effects are stronger in Germany than in France. However, attributing either of these observations to the introduction of regulation is complicated by a number of factors. Both electricity and fuel prices increase, and the patent rate for regulated firms differs from that for unregulated firms even prior to the introduction of regulation. Thus, any estimation strategy must account for both unobserved heterogeneity and the impact of factors other than the ETS on low-carbon patenting. The preferred specification does both.

5 Results and Discussion

5.1 Models with electricity price only

I first estimate (20) using the three estimation strategies outlined earlier, then compute estimates of the direct, indirect, and total policy effects according to Sections 4.2.1 and 4.2.2.[20] The results in Table 2 indicate that being subjected to ETS regulation increases patent output, higher electricity prices lead to additional patents, and that lagged own knowledge increases patent output. In contrast, evidence for strong cross-firm knowledge spillovers in this context is weak, with significant estimates occurring only in the fixed effects model, and not in the preferred presample model.

Estimates of the direct, indirect, and total policy effects computed according to Sections 4.2.1 and 4.2.2 are reported at the bottom of Table 2. Calculation of those estimated effect sizes requires specification of a pass-through rate relating the emissions permit price to the *retail* electricity price. Very few estimates of the retail pass-through rates exist; most studies concern wholesale electricity prices due to the frequency and availability of the data. In one of the only studies attempting to quantify retail pass-through rates, Sijm et al. (2008) offer several such estimates for Germany, but none for France. To provide conservative estimates of the indirect and total effects of interest, I use the lower of their two estimates of retail pass-through for Germany (4.8%) that are derived from empirical estimates of wholesale pass-

[20]In current estimates, I adopt the restriction $\beta_{\cdot}^{R} = \beta_{\cdot}^{U}$, i.e. predictor variables impact all firms identically (except the treatment status variable). See Section 5.3 for relaxation of this restriction for the electricity price coefficient.

through rates.[21] For France, due to retail price regulation and in keeping with my desire to provide conservative estimates, I assume that while wholesale pass-through rate estimates are positive, there will be no pass-through to retail prices (0%). For the preferred presample mean specification, the estimated direct effect of the ETS on patenting, captured by $\hat{\gamma}^{R}_{Direct}$ is 151 additional patents. For comparison, Calel and Dechezleprêtre (2014) report an increase between 84 and 188 additional patents, depending on whether they consider only matched firms or extrapolate to all ETS firms using their estimated treatment effect. The estimates of $\hat{\gamma}^{R}_{Indirect}$ and $\hat{\gamma}^{U}_{Indirect}$ are of greater interest, as those indirect effects are the central focus of this paper. The indirect effect on regulated firms is minor (10 additional patents), while the indirect effect on unregulated firms is substantial (172 additional patents).

These results suggest that indirect effects may be at least as large as the direct effects, such that the total effect of regulation on innovation may be twice the size of frequently reported estimates of direct effects. It is this relative size of the direct and total policy effects that I wish to emphasize. While the absolute level of innovation spurred by the EU ETS remains modest, for higher emission permit prices, the size of the indirect effect would increase in absolute terms.

5.2 Models including emissions permit price

The results in the previous section provide a two-part argument that the EU ETS has increased low-carbon patenting by unregulated firms: the ETS increased electricity price, and an increased electricity price leads to more low-carbon patenting. Estimates of (21) should offer more direct evidence of such an effect, since the carbon price appears directly in the model. Estimates of three variants of that model (pooled, fixed effects, and presample mean) are reported in Table 3, with standard errors again adjusted for over-dispersion.

The results in Table 3 suggest that again, own knowledge and ETS regulatory status are important determinants of patent output. In addition, the estimates provide some evidence that passed-through carbon costs contribute positively to patent output for German firms, including firms that are not subject to the regulation. In contrast, there is no such positive effect for French firms. This discrepancy in effects across the two countries is consistent with the hypothesis that permit prices impact innovation through cost pass-through, since

[21] In that case, those authors assume the *level* of increase in retail prices is the same as that in wholesale prices, but they estimate the wholesale pass-through rates. Since retail prices are typically higher than wholesale prices, the pass-through rates as a % of retail prices will be lower than those estimated in wholesale prices. They offer other estimates based on alternate assumptions of 1) any difference in pre-tax retail spreads being due to emissions permits or 2) emissions permit costs being fully passed through to both wholesale and retail prices. The retail pass-through estimates I use, which are based on empirical wholesale pass-through estimates, appear to require the fewest assumptions.

the scope for pass-through is much higher in the more liberalized German retail market. If permit prices were to affect innovation through a channel other than the electricity price, we should still see an effect of the permit price on patenting in France. If, instead, permit prices do not affect innovation outside the regulated sector, we should see a zero effect in Germany. The fact that we see an effect only in Germany is consistent with the hypothesis that indirectly induced innovation is mediated by cost pass-through.

5.3 Additional tests

While the main results presented above are consistent with the ETS having indirect effects on low-carbon innovation, it is possible they reflect some other driver of innovation. To address that possibility, I attempt to falsify the results through estimation of variants of (20). First, since the stylized theoretical framework is based on unregulated firms being users of the output produced by regulated firms (i.e. electricity), we should see a differentially stronger patenting response by unregulated firms if we allow the coefficient on electricity price, β^S_{elec} to vary by regulatory status. Doing so yields the results in Table 4, which suggest that low-carbon electrical patenting by unregulated firms is indeed more elastic with respect to the electricity price than that by regulated firms.[22] I have also estimated variants of both (20) and (21) allowing β_{elec} and β_τ to vary by industrial sector rather than regulatory status. Lending credibility to the specification, firms in sectors that use substantial amounts of electricity, such as chemical, metals, machinery, electronics, and telecommunications, exhibit strong, positive patenting responses to higher electricity prices and pass-through.

Second, the hypothesized cost pass-through effect should primarily affect incentives for electricty-related low-carbon innovation. If cost pass-through is indeed one of the mechanisms through which the ETS impacts low-carbon innovation by unregulated firms, we should not see a strong patenting response in other types of low-carbon technologies that are unrelated to electricity.[23] To examine whether non-electrical, low-carbon patenting responds to electricity prices, I estimate (20) with both the response variable and knowledge stocks redefined to represent low-carbon patents that are not related to electricity. The results in Table 5 fail to reject the null hypothesis that increases in the electricity price have no impact on non-electrical low-carbon patenting. That null result lends some credibility to interpretation of the main model estimates as a causal effect of the electricity price on low-carbon electrical patenting as opposed to the effect of some omitted driver of low-carbon patenting.

[22]Even though many regulated firms are electricity users, a number of larger manufacturers generate their own electricity, and thus will be less responsive to changes in the electricity price.

[23]To the extent that higher electricity prices have any impact on other types of low-carbon patenting, we might expect there to be a negative effect if R&D budgets are binding and firms choose to reallocate R&D budgets from other low-carbon technologies to electricity related ones.

5.4 Discussion

In light of the fact that both estimation approaches suggest indirect effects of the EU ETS are non-negligible, it is worth examining why such indirect effects appear here but not in carefully done studies such as that by Calel and Dechezleprêtre (2014). Those authors acknowledged the possibility of indirect effects but found little evidence of them.

At least two differences in the estimation strategy used by those authors and the approach taken here may help explain this discrepancy. First, holding the set of firms considered constant, the matching estimator they use identifies differential effects, whereas the approach here is designed to capture total effects. For example, their estimator will correctly pick up any differential effect of knowledge spillovers on regulated firms (e.g. if regulated firms have greater technology absorption capabilities). However, it will not include any baseline impact of knowledge spillovers on both regulated firms and their matched unregulated counterparts.

A second potential reason I find stronger support for indirect effects concerns the set of unregulated firms that might be responding indirectly to regulation. Calel and Dechezleprêtre (2014) consider unregulated firms that co-patent with regulated firms, with an eye toward those co-patenters being technology *suppliers* to regulated firms. In contrast, the model introduced here focuses on *consumers* of regulated firm output, and how their innovation incentives are affected by regulation through cost pass-through. For example, the vast majority of firms use electricity in their daily activities, yet most of them are not regulated, not suitable matches for regulated firms, and not technology suppliers to regulated firms. Thus any estimator focused on innovation that directly reduces emissions is likely to miss any response by those electricity users. This suggests a broader definition of low-carbon innovation to include technology related to efficient use of goods produced by regulated firms (e.g. electrical efficiency). The concept of "embedded carbon" is familiar in lifecycle analysis but it merits more careful consideration in studies of induced innovation.

6 Conclusions

Increasingly, a stated objective of environmental policy is to stimulate the development of new technology that will make achievement of environmental goals less costly. This is especially true for climate policy, for which the social costs of not achieving environmental targets may be quite large. Traditional policy analysis seeks to quantify the innovation effects of a particular environmental policy by estimating a treatment effect on firms regulated under that policy. This paper has argued that the total effects of a policy may diverge from the direct effects, using both a simple modeling framework and an empirical application to the

European Union Emissions Trading System. In that empirical application, estimates suggest that the total effect of the EU ETS on low-carbon electrical patents is at least twice as large as the direct effect that would be estimated via standard methods based upon the treatment effects framework. The absolute size of the impact of the EU ETS on low-carbon innovation remains small, but the relative size of the direct and total policy effects suggests that for policies imposing higher costs on regulated firms, failure to account for indirectly induced innovation may have more important consequences.

Further, I argue the general mechanism for indirect effects is likely to apply to a broad array of environmental policies. As seen in the simplified theoretical model, pass-through of policy-imposed costs is likely to occur whenever regulated firms face little unregulated competition in their output market(s). This can occur when all firms are regulated, or when only some firms are regulated, but those firms possess market power. The latter scenario is quite common in environmental policies. Firms with market power are often included in environmental regulation: they are frequently large emitters, and the concentration of emissions among a smaller number of firms may lower monitoring costs for the regulator. As a result, the mechanisms presented here may be more broadly applicable to a large number of environmental policies, though the strength of the indirect policy effects on innovation will differ by policy, industry, and market conditions.

A Appendix

A.1 Deriving cost pass-through

Since w^R is an endogenously determined price that depends upon regulated firm output, the effect of τ on w^R acts solely through output decisions of regulated firms. Further, since $w^R = P(X^{R*})$, we may write

$$\frac{\partial w^R}{\partial \tau} = \sum_{i:S_i=R} \frac{\partial w^R}{\partial x_i^{R*}} \frac{\partial x_i^{R*}}{\partial \tau} = N^R P'(X^{R*}) \frac{\partial x_i^{R*}}{\partial \tau}. \tag{25}$$

Next, we can derive an expression for $\frac{\partial x_i^{R*}}{\partial \tau}$ through total differentiation of (7). Rearranging the result of that differentiation yields $\frac{\partial x_i^{R*}}{\partial \tau}$, which can be substituted into (25) to yield (9).

A.2 Signing cost pass-through

Differentiate (7) with respect to the output of some other firm j:

$$P''(X^{R*})x_i^{R*} + P'(X^{R*})\frac{\partial x_i^{R*}}{\partial x_j^R} + P'(X^{R*}) = c_{xx}(x_i^{R*}, w)\frac{\partial x_i^{R*}}{\partial x_j^R} + \tau e_{xx}(x_i^{R*}, k_i^R)\frac{\partial x_i^{R*}}{\partial x_j^R}$$

and solve for $\frac{\partial x_i^{R*}}{\partial x_j^R}$:

$$\frac{\partial x_i^{R*}}{\partial x_j^R} = -\frac{P''(X^{R*})x_i^{R*} + P'(X^{R*})}{P'(X^{R*}) - c_{xx}(x_i^{R*}, w) - \tau e_{xx}(x_i^{R*}, k_i^R)}$$

Equilibrium stability requires the entire expression to be negative: reaction functions should be downward sloping (see Seade 1980). Since the denominator is negative under the maintained assumptions, stability requires $P''(X^{R*})x_i^{R*} + P'(X^{R*}) < 0$.

Similarly, the firm's second-order condition for output requires for a maximum that

$$P''(X^{R*})x_i^{R*} + 2P'(X^{R*}) - c_{xx}(x_i^{R*}, w) - \tau e_{xx}(x_i^{R*}, k_i^R) < 0 \tag{26}$$

Summing $(N-1)$ times the stability condition plus the second order condition yields the denominator in (9). Since it is the sum of negative components, that denominator is negative. As a result, cost pass-through is positive provided that the equilibrium is stable. Note that the denominator captures the change in marginal profits that occur from a simultaneous marginal increase in quantity by all firms. If marginal profits increased in such a case at some equilibrium profile of quantity choices, the equilibrium could not be stable.

A.3 Derivation of optimal R&D choices

Derivation of R&D incentives follows Leahy and Neary (1997). The first-order condition for R&D by unregulated firm j is then

$$\frac{\partial \pi_j^{U*}}{\partial k_j^U}\frac{\partial k_j^U}{\partial r_j^U} + \sum_{l \neq j} \frac{\partial \pi_j^{U*}}{\partial x_l^{U*}}\frac{\partial x_l^{U*}}{\partial r_j^U} = 1. \tag{27}$$

Thus, the unregulated firm accounts for the direct effect of augmented knowledge on its own profits, but must also account for the effect of its R&D on second period output choices of other unregulated firms. Note, however, that this expression does not include an effect of r_j^U on profits through output choices of regulated firms. In keeping with the assumption that unregulated firms are price takers, unregulated firms assume that their R&D will not impact output decisions of regulated firms (and hence the equilibrium price w^R).

Note that by the envelope theorem, $\frac{\partial \pi_j^{U*}}{\partial k_j^U} = -w^R D_k^R$, and by the knowledge dynamics, $\frac{\partial k_j^U}{\partial r_j^U} = f_r^U$. In addition, $\frac{\partial \pi_j^{U*}}{\partial x_l^{U*}} = P' x_j^{U*}$. For a symmetric equilibrium, we may rewrite (27) as

$$-w^R D_k^R f_r^U + (N^U - 1) P' x_{j2}^{U*} \frac{\partial x_l^{U*}}{\partial r_j^U} = 1. \tag{28}$$

To derive $\frac{\partial x_l^{U*}}{\partial r_j^U}$, we start with the first order conditions defining x_l^{U*} and x_j^{U*}, and consider a shock that affects output decisions of all unregulated firms and R&D of unregulated firm j. Totally differentiate both first order conditions(except hold r_l^U, $l \neq j$, and w^R constant), substitute for x_j^{U*} from the differentiation of x_j^{U*} into the result for x_l^{U*}, and rearrange. Letting π_{lj}^{Uj} denote the second order partial derivative of profits by firm j with respect to output of unregulated firms l and j, differentiation yields:

$$\begin{aligned} \pi_{jj}^{Uj} dx_j^U + (N^U - 1)\pi_{lj}^{Uj} dx_l^U + \frac{\partial \pi_j^{Uj}}{\partial k_j^U} f_r^U dr_j^U &= 0, \\ \pi_{ll}^{Ul} dx_l^U + \pi_{jl}^{Ul} dx_j^U + (N^U - 2)\pi_{l'l}^{Ul} dx_{l'}^U + \frac{\partial \pi_l^{Ul}}{\partial k_j^U} f_{R^U}^U dr_j^U &= 0. \end{aligned}$$

Assuming symmetry of the equilibrium prior to the shock (i.e. $dx_l^U = dx_{l'}^U$, $\pi_{l'l}^{Ul} = \pi_{jl}^{Ul} = \pi_{lj}^{Uj}$, $\pi_{jj}^{Uj} = \pi_{ll}^{Ul}$, and $\frac{\partial \pi_j^{Uj}}{\partial k_j^U} = -w^R D_k^R$), solving these two equations and rearranging yields:

$$\frac{dx_l^U}{dr_j^U} = \frac{\pi_{jj}^{Uj} \left(f_{R^U}^U - \frac{\pi_{lj}^{Uj}}{\pi_{jj}^{Uj}} f_r^U \right) w^R D_k^R}{\left(\pi_{jj}^{Uj} - \pi_{lj}^{Uj} \right) \left(\pi_{jj}^{Uj} + (N^U - 1)\pi_{lj}^{Uj} \right)}. \tag{29}$$

Substituting (29) into (28) and factoring out $-f_r^U w^R D_k^R$ yields (13).

References

Aghion, P., A. Dechezleprêtre, D. Hemous, R. Martin, and J. V. Reenen (Dec. 2012). *Carbon Taxes, Path Dependency and Directed Technical Change: Evidence from the Auto Industry*. Working Paper 18596. National Bureau of Economic Research.

Aronow, P. M. and C. Samii (2012). "Estimating average causal effects under general interference." In: *Summer Meeting of the Society for Political Methodology, University of North Carolina, Chapel Hill, July*. Citeseer, pp. 19–21.

Baker, E. and E. Shittu (2006). "Profit-maximizing R&D in response to a random carbon tax." In: *Resource and Energy Economics* 28(2), 160–180.

Blundell, R., R. Griffith, and J. V. Reenen (1995). "Dynamic Count Data Models of Technological Innovation." English. In: *The Economic Journal* 105(429), 333–344.

Brunnermeier, S. B. and M. A. Cohen (2003). "Determinants of environmental innovation in {US} manufacturing industries." In: *Journal of Environmental Economics and Management* 45(2), 278–293.

Burtraw, D., K. Palmer, R. Bharvirkar, and A. Paul (2001). "Cost-effective reduction of NOx emissions from electricity generation." In: *Journal of the Air & Waste Management Association* 51(10), 1476–1489.

Bushnell, J. B., H. Chong, and E. T. Mansur (2013). "Profiting from regulation: Evidence from the european carbon market." In: *American Economic Journal: Economic Policy* 5(4), 78–106.

Calel, R. and A. Dechezleprêtre (2014). "Environmental Policy and Directed Technological Change: Evidence from the European carbon market." In: *Review of economics and statistics*. (Forthcoming).

Convery, F., A. D. Ellerman, and C. De Perthuis (2008). *The European Carbon Market in Action: Lessons from the First Trading Period Interim Report*. Tech. rep. MIT Joint Program on the Science and Policy of Global Change.

Cowan, S. (1993). "Regulation of several market failures: the water industry in England and Wales." In: *Oxford Review of Economic Policy*, 14–23.

Cox, D. R. (1958). "Planning of experiments." In:

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

Fabra, N. and M. Reguant (Nov. 2013). *Pass-through of Emissions Costs in Electricity Markets*. Working Paper 19613. National Bureau of Economic Research.

Fischer, C. and R. G. Newell (2008). "Environmental and technology policies for climate mitigation." In: *Journal of environmental economics and management* 55(2), 142–162.

Fischer, C., I. W. Parry, and W. A. Pizer (2003). "Instrument choice for environmental protection when technological innovation is endogenous." In: *Journal of Environmental Economics and Management* 45(3), 523–545.

Gianessi, L. P., H. M. Peskin, and E. Wolff (1979). "The distributional effects of uniform air pollution policy in the United States." In: *The Quarterly Journal of Economics*, 281–301.

Griliches, Z. (1979). "Issues in Assessing the Contribution of Research and Development to Productivity Growth." English. In: *The Bell Journal of Economics* 10(1), pages.

Hicks, J. R. (1932). "The Theory of Wages." In: *and* 245, 117.

Hoffmann, V. H. (2007). "EU ETS and Investment Decisions:: The Case of the German Electricity Industry." In: *European Management Journal* 25(6). Business, Climate Change and Emissions Trading, 464–474.

Hudgens, M. G. and M. E. Halloran (2008). "Toward causal inference with interference." In: *Journal of the American Statistical Association* 103(482).

Jaffe, A. B. and K. Palmer (1997). "Environmental regulation and innovation: a panel data study." In: *Review of economics and statistics* 79(4), 610–619.

Johnstone, N., I. Haščič, and D. Popp (2010). "Renewable energy policies and technological innovation: evidence based on patent counts." In: *Environmental and Resource Economics* 45(1), 133–155.

Keppler, J. H. and M. Mansanet-Bataller (2010). "Causalities between CO2, electricity, and other energy variables during phase I and phase {II} of the {EU} {ETS}." In: *Energy Policy* 38(7). Large-scale wind power in electricity markets with Regular Papers, 3329–3341.

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

Leahy, D. and J. P. Neary (1997). "Public Policy Towards R&D in Oligopolistic Industries." English. In: *The American Economic Review* 87(4), 642–662.

Lund, P. (2007). "Impacts of {EU} carbon emission trade directive on energy-intensive industries — Indicative micro-economic analyses." In: *Ecological Economics* 63(4). Sustainability and Cost-Benefit Analysis, 799–806.

Magerman, T., B. Van Looy, and X. Song (2006). "Data production methods for harmonized patent statistics: Patentee name harmonization." In: *K.U. Leuven FETEW MSI Research report 0605*, 1–88.

Manski, C. F. (1993). "Identification of Endogenous Social Effects: The Reflection Problem." English. In: *The Review of Economic Studies* 60(3), pages.

Miguel, E. and M. Kremer (2004). "Worms: identifying impacts on education and health in the presence of treatment externalities." In: *Econometrica* 72(1), 159–217.

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

Nesta, L., F. Vona, and F. Nicolli (2014). "Environmental policies, competition and innovation in renewable energy." In: *Journal of Environmental Economics and Management* 67(3), 396–411.

Oberndorfer, U. (2009). "{EU} Emission Allowances and the stock market: Evidence from the electricity industry." In: *Ecological Economics* 68(4). Participation and Evaluation for Sustainable River Basin Governance, 1116–1126.

Popp, D. (2002). "Induced Innovation and Energy Prices." In: *American Economic Review* 92(1), 160–180.

Porter, M. E. and C. van der Linde (1995). "Toward a New Conception of the Environment-Competitiveness Relationship." In: *Journal of Economic Perspectives* 9(4), 97–118.

Requate, T. and W. Unold (2003). "Environmental policy incentives to adopt advanced abatement technology:: Will the true ranking please stand up?" In: *European Economic Review* 47(1), 125–146.

Rubin, D. B. (1974). "Estimating causal effects of treatments in randomized and nonrandomized studies." In: *Journal of educational Psychology* 66(5), 688.

Seade, J. (1980). "The stability of Cournot revisited." In: *Journal of Economic Theory* 23(1), 15–27.

Sijm, J., S. Hers, W. Lise, and B. Wetzelaer (2008). "The impact of the EU ETS on electricity prices." In: *Final report to DG Environment of the European Commission. ECN-E-08-007.*

Sijm, J., K. Neuhoff, and Y. Chen (2006). "CO2 cost pass-through and windfall profits in the power sector." In: *Climate Policy* 6(1), 49–72.

Veith, S., J. Werner, and J. Zimmermann (2009). "Economic Consequences of Emission Trading Schemes: Evidence from the European Power Sector." In: *Energy Economics* 31(4), 605–613.

Vukina, T. (2003). "The relationship between contracting and livestock waste pollution." In: *Review of Agricultural Economics* 25(1), 66–88.

Weyl, E. G. and M. Fabinger (2013). "Pass-through as an economic tool: principles of incidence under imperfect competition." In: *Journal of Political Economy* 121(3), 528–583.

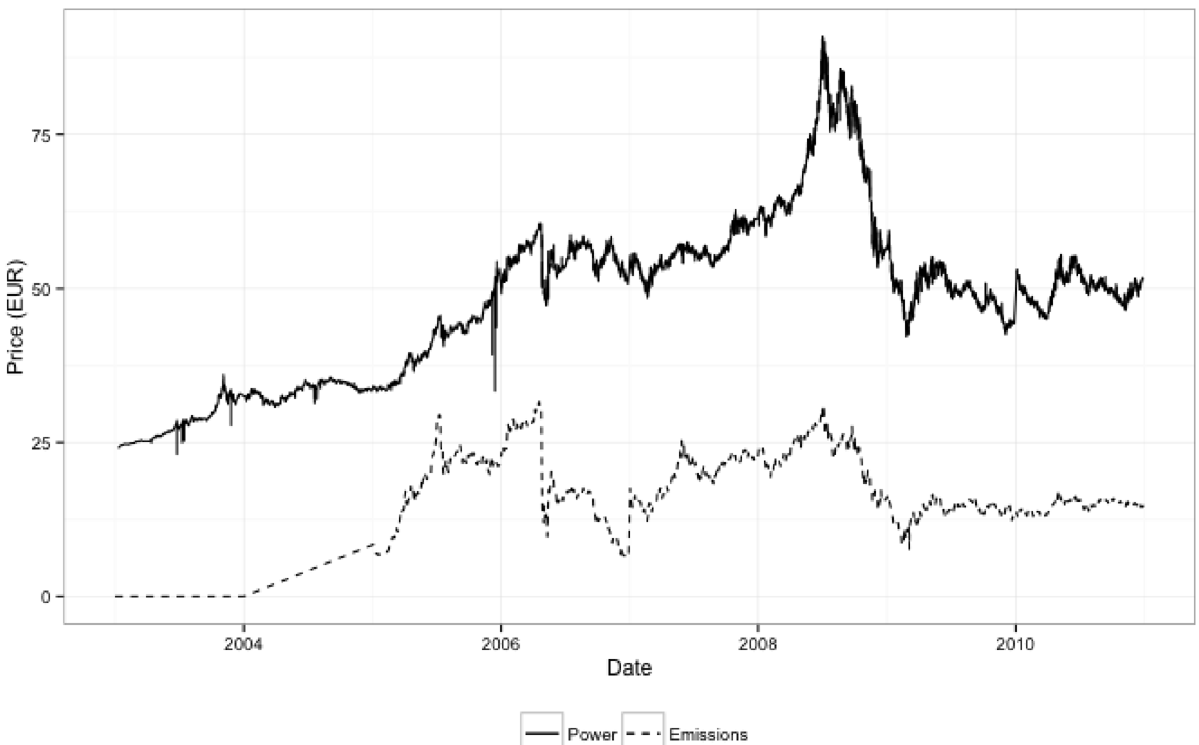

Figure 1: Prices over time for year-ahead power futures and year-ahead emissions permit futures. All prices are nominal and in EUR. Source: European Energy Exchange (EEX).

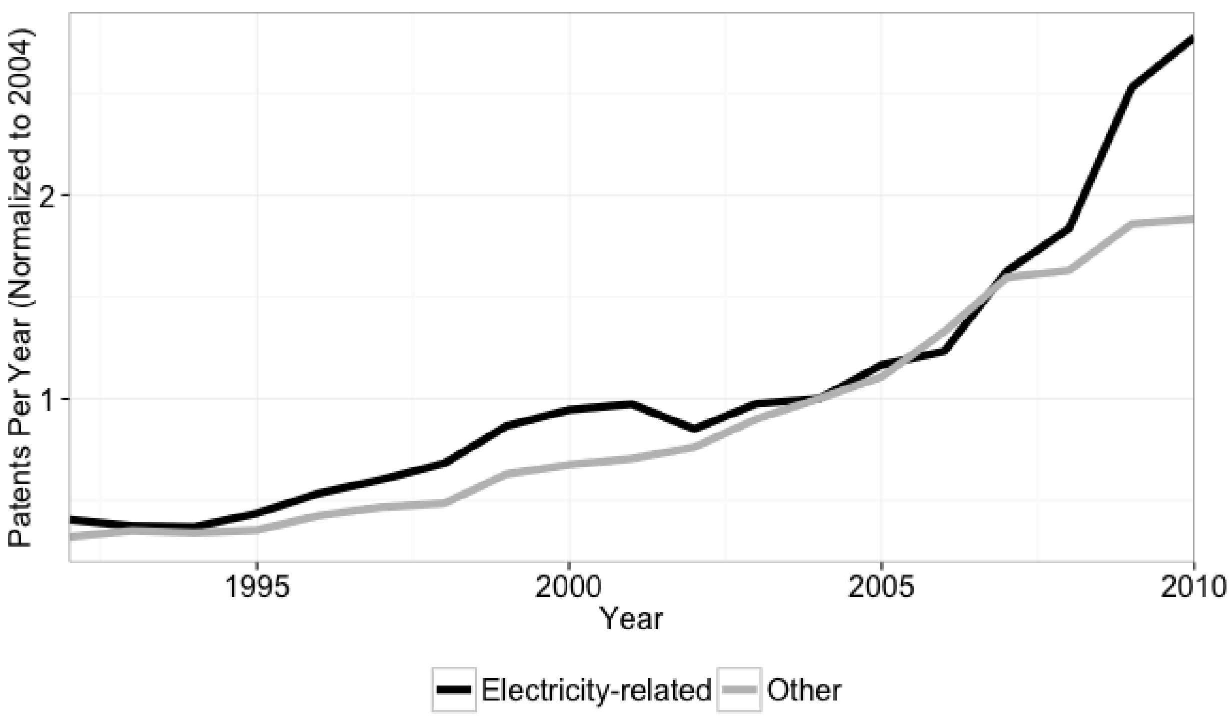

Figure 2: Low-carbon patents by type (electrical vs other) per year in Germany and France, normalized to 2004 levels. Source: European Patent Office (EPO).

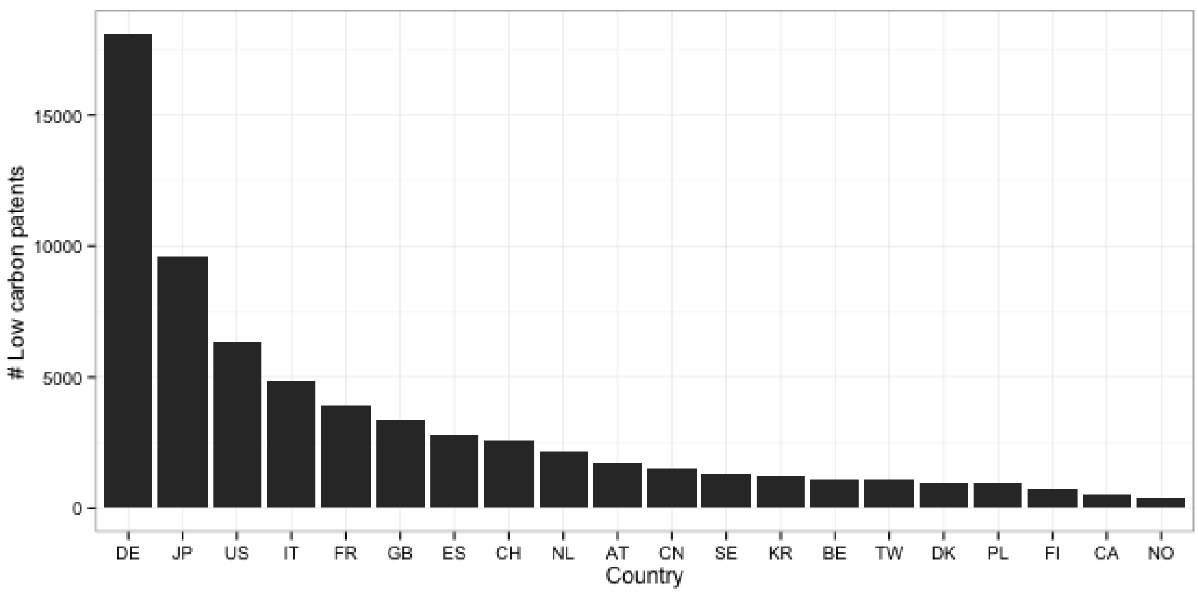

Figure 3: Number of low-carbon patents by country during 2000-2010 for the top 20 low-carbon patenting countries. Source: EPO.

Table 1
Summary statistics for German and French firms.

	Germany			France		
	All	ETS	Non-ETS	All	ETS	Non-ETS
# Patents pre-ETS	1952	151	1801	483	12	471
# Patents during ETS	3341	563	2778	425	35	390
# Mean patents pre-ETS	0.060	0.242	0.056	0.059	0.044	0.060
# Mean patents during ETS	0.221	1.955	0.187	0.113	0.278	0.107
Electricity price pre-ETS	98.38	–	–	70.74	–	–
Electricity price during ETS	119.62	–	–	73.08	–	–
Natural gas price pre-ETS	7.08	–	–	5.13	–	–
Natural gas price during ETS	13.36	–	–	10.18	–	–
# Firms	2518	48	2470	627	21	606

Table 2
Main regression results: low-carbon electrical patents as a function of predictors including electricity price. Estimates of coefficients corresponding to other covariates, including factor prices, presample mean, country effects, and total patent rate are omitted for emphasis on the effects of interest.

		Pooled	FE	Presample
Log total knowledge	$ln(K_t)$	-0.082	0.295***	0.047
		(0.093)	(0.019)	(0.091)
No own knowledge	$\mathbb{1}(k_{it} = 0)$	-2.733***	-0.890***	-2.712***
		(0.051)	(0.021)	(0.050)
Log own knowledge	$ln(k_{it})$	1.084***	0.421***	0.988***
		(0.017)	(0.006)	(0.020)
Log elec. price	$ln(w_{t-1}^{elec})$	0.513*	0.715***	0.528*
		(0.241)	(0.072)	(0.237)
Regulatory status	T_{it-1}	0.261***	0.532***	0.248**
		(0.072)	(0.032)	(0.076)
Policy effect estimates				
Direct effect	$\hat{\gamma}^R_{Direct}$	178	80	151
		(29, 478)	(60, 113)	(28, 390)
Indirect effect (reg)	$\hat{\gamma}^R_{Indirect}$	6	5	9
		(2, 9)	(3, 8)	(3, 17)
Indirect effect (unreg)	$\hat{\gamma}^U_{Indirect}$	99	129	155
		(39, 131)	(74, 242)	(49, 283)
Total Policy Effect	TPE	283	214	315
		(70, 618)	(137, 363)	(80, 690)

Signif. codes: 0 *** 0.001 ** 0.01 * 0.05 . 0.1 1

Table 3
Regression results: low-carbon electrical patents on predictors including emissions permit price.

		Pooled	FE	Presample
Log total knowledge	$ln(K_t)$	0.018	0.449***	0.138.
		(0.080)	(0.018)	(0.079)
No own knowledge	$\mathbb{1}(k_{it}=0)$	-2.735***	-0.893***	-2.714***
		(0.050)	(0.021)	(0.050)
Log own knowledge	$ln(k_{it})$	1.083***	0.417***	0.986***
		(0.018)	(0.006)	(0.020)
Log elec. price	$ln(w_{t-1}^{elec})$	0.377	0.501***	0.375
		(0.284)	(0.073)	(0.280)
Permit price	τ_{t-1}			
	Germany	0.009*	0.012***	0.008*
		(0.004)	(0.001)	(0.004)
	France	0.004	-0.002	0.001
		(0.006)	(0.003)	(0.006)
Regulatory status	T_{it-1}	0.259***	0.508***	0.247**
		(0.072)	(0.033)	(0.076)
Policy effect estimates				
Direct effect	$\hat{\gamma}_{Direct}^{R}$	113	46	103
		(24, 317)	(37, 57)	(24, 276)
Indirect effect (reg)	$\hat{\gamma}_{Indirect}^{R}$	31	16	33
		(13, 56)	(13, 19)	(14, 70)
Indirect effect (unreg)	$\hat{\gamma}_{Indirect}^{U}$	509	363	561
		(217, 934)	(265, 484)	(227, 1226)
Total Policy Effect	TPE	653	425	697
		(254, 1307)	(315, 560)	(265, 1572)

Signif. codes: 0 *** 0.001 ** 0.01 * 0.05 . 0.1 1

Table 4
Regression results: low-carbon electrical patents on predictors, allowing electricity price coefficient to vary with regulatory status. All estimates use presample mean estimator to account for unobserved heterogeneity.

Variable		Estimate	
Log total knowledge	$ln(K_t)$	0.049	
		(0.091)	
No own knowledge	$\mathbb{1}(k_{it} = 0)$	-2.700	***
		(0.050)	
Log own knowledge	$ln(k_{it})$	0.988	***
		(0.020)	
Log elec. price	$ln(w^{elec}_{t-1})$		
	Regulated firms	0.444	.
		(0.236)	
	Unregulated firms	0.600	*
		(0.236)	
Regulatory status	T_{it-1}	-0.076	
		(0.120)	

Signif. codes: 0 *** 0.001 ** 0.01 * 0.05 . 0.1 1

Table 5
Regression results: low-carbon non-electrical patents on predictors. All estimates use pre-sample mean estimator to account for unobserved heterogeneity.

Variable		Estimate	
Log total knowledge	$ln(K_t)$	0.063	
		(0.100)	
No own knowledge	$\mathbb{1}(k_{it} = 0)$	-2.075	***
		(0.035)	
Log own knowledge	$ln(k_{it})$	0.991	***
		(0.020)	
Log elec. price	$ln(w_{t-1}^{elec})$	0.076	
		(0.162)	
Regulatory status	T_{it-1}	0.019	
		(0.043)	

Signif. codes: 0 *** 0.001 ** 0.01 * 0.05 . 0.1 1