Carbon taxes, Path Dependency and Directed Technical Change : Evidence from the Auto Industry*

Philippe Aghion†, Antoine Dechezlepretre‡

David Hemous§, Ralf Martin¶ and John Van Reenen‖

December 10, 2010

Abstract

We show a strong impact of tax-adjusted fuel prices in inducing directed technical change to mitigate climate change. A model of endogenous and path dependent directed innovation implies that carbon taxes can stimulate the development of "clean" technologies, but that the effect of this depends on a firm's innovation history. We exploit new firm-level panel data on the auto industry examining "dirty" (internal combustion engine) and "clean" (e.g. electric and hybrid) patents 1978-2007 across 80 countries. We show that firms tend to innovate relatively more in clean when they face a higher tax-adjusted fuel prices. Further, although there is path dependence in the type of innovation, tax-adjusted fuel prices induce more clean innovation for firms with a prior history of dirty innovations as predicted by our model.

JEL Classification: O3

Keywords: Climate Change, Directed Technical Change, automobiles

---

*We would like to thank Nick Stern for inspiration and many helpful comments. Financial support has come from the ESRC Centre for Economic Performance.

†Harvard University

‡Centre for Economic Performance, London School of Economics

§Harvard University

¶Centre for Economic Performance, London School of Economics

‖Centre for Economic Performance, London School of Economics and NBER

1 Introduction

With increasing concern about the possibility of human induced climate change there is now much interest in new technologies and innovation that will help reduce emissions of greenhouse gases such as Carbon Dioxide (CO_2). Most models of climate change assume exogenous technological change (e.g. Stern, 2006), but dealing with the challenge of global warming almost certainly requires some more climate change related innovation.

Standard models suggest that the market will generate insufficient climate change related innovation and too much R&D investment directed at "dirty" technologies. For example, in Acemoglu et al. (2009, henceforth AABH) an important feature is that there is path-dependence in the direction of technical change: namely, firms that have innovated a lot in dirty technologies in the past will find it more profitable to innovate in dirty technologies today. This path dependence feature combined with the environmental externality whereby firms do not factor in the loss in aggregate productivity or consumer utility induced by environmental degradation, will induce the a laissez-faire economy to produce and innovate too much in dirty technologies compared to the social optimum. This in turn calls for government intervention to "redirect" technical change.

A challenge to the path-dependence hypothesis is that there may be decreasing returns to each type (clean or dirty) of innovation, so that a firm that has innovated dirty a lot in the past would have more incentives to innovate clean today. In that case the market would do part of the job of redirecting technical change towards clean technologies, albeit temporarily (until firms run sufficiently into decreasing returns on innovations in clean technologies that they would choose to shift again towards dirty innovations).

In this paper, we exploit a new patent data set on innovations in the auto industry to examine the existence of directed technical change when carbon prices change and its relation to path dependence. Our data are drawn from the European Patent Office (EPO) World Patent Statistical database. These data cover the population of EPO patents (and their citations) since 1978 (over 80 countries took out such patents). In automobiles we estimate that around 12,000 patents in "clean" technologies (electric vehicles, hybrid vehicles, fuel cells,..) were filed and about 36,000 patents in "dirty" technologies which affect regular combustion engines. Moreover, our database reports the name of patent applicants which in turn allows us to match clean and dirty patents with distinct patent holders each of whom has her own history of clean versus dirty patenting.

Our main results can be summarized as follows: (i) higher tax-adjusted fuel price encourages clean innovation consistent with the directed technical change hypothesis; (ii) a firm's propensity to innovate in "clean" technologies is positively correlated with their lagged stock of clean patents and vice versa, consistent with the path-dependence hypothesis; (iii) finally, the positive marginal impact of a higher tax-adjusted fuel price on the propensity to clean innovation is stronger for firms with a higher stock of *dirty* patents and weaker for firms with a higher stock of *clean* patents.

These results have a number of potential implications. First, in addition to a beneficial effect on reducing consumer demand for carbon, higher carbon taxes induce relatively more clean innovation which magnifies the benefit of such a policy. Second, absent government intervention, firms that have innovated dirty in the past tend to get locked in the same type of innovative activities in the future. This makes the task of climate change reduction harder as the default option of the economy is to increase demand for carbon-using technologies. Third, pollution taxes redirect innovation towards clean mostly where this is needed the most, namely in firms with higher stocks of past dirty innovations.

Our research relates to a small empirical literature on the effect of energy prices on the direction of technical progress. In particular Popp (2002) uses aggregate U.S. patent data from 1970 to 1994 to study the effect of energy prices on energy-efficient innovations. Popp constructs two different measures of the knowledge stocks for the innovation regressions: (a) a simple stock of previously U.S. granted patents and (b) a quality-adjusted stock of patents weighted by the productivity estimates. In particular, he finds a significant impact from both, energy prices and the quality of the stock of knowledge available to the inventor, on directed innovations. This provides evidence in favor of directed technical change as a response to change in energy prices, but because the data is aggregate a concern is that there may other macro-economic shocks correlated with both innovation and the energy price. Because we have international firm-level data we can exploit differential policy-induced shocks to the energy price (e.g. fuel taxes) across countries to which firms are differentially exposed because of the market access. Further evidence of directed technical change applied to the context of saving energy can be found in Newell et al (1999) which focuses on the air-conditioning industry, or in Lanzi and Sue Wing (2010) who estimate the price elasticity of innovation in fossil versus non fossil fuel energy. However, none of these papers look at the effect of past clean versus dirty innovations on current innovation, and in particular they do not analyze whether there is path-dependence in the direction of technical change.

Other papers have also argued that ignoring directed technical change overstates the costs of environmental regulation (see Grubler and Messner (1998), Manne and Richels (2002), Messner (1997), Buonanno et al (2003), Nordhaus (2002), Sue Wing (2003)). The measure of the overstatement of costs depends on specific characteristics of the models found in these papers, namely, the possibility of crowding-out in R&D towards energy-saving innovations.[1] However, once again

[1]As an example, Popp (2004) modifies the standard Nordhaus model of climate change (DICE model) to allow for induced innovation in the energy sector. After some calibration and simulation exercises and allowing for the possibility of crowding-out across different kinds of R&D investment, he concludes that ignoring DTC overstates the

none of these papers deals with the path-dependence issue or how it interacts with direct technical change.

The paper is organized as follows. Section 2 develops a simple one-period model to guide our empirical analysis. Section 3 presents the data and the econometric methodology. In Section 4, we provide a description of the data. Section 5 presents the results and discusses their robustness. Section 6 concludes.

2 Model

2.1 Demand, production and innovation

We consider a one-period model of an industry populated by a mass 1 of different varieties. Each variety i is produced by a monopolist who faces an inverse demand curve of the form:

$$y_i = \left(p_i'\right)^{-\sigma} P^{\sigma-\beta} \tag{1}$$

where $\sigma > 1$ is the elasticity of substitution between the different varieties, ${p_i}'$ is the consumer (after tax) price, P is the price index (defined by: $P = \left(\int_0^1 \left(p_i'\right)^{1-\sigma} di\right)^{\frac{1}{1-\sigma}}$) and β is the elasticity of consumption of the composite good with respect to the overall price index.[2]

Good i is produced using an energy input with a linear technology. The energy input is produced one for one with labor (wages being normalized to 1), and comes in two forms, either clean or dirty. We denote by x_{ji} the amount of clean ($j = c$) or dirty ($j = d$) energy inputs used by the producer

welfare costs of an optimal tax policy by 9.4% in the base case (where partial - 50% - crowding-out is allowed).

[2]This demand structure for varieties can be generated by a quasi-linear utility function

$$u = C_0 + \frac{\beta}{\beta-1}\left(\int_0^1 c_i^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}\frac{\beta-1}{\beta}},$$

where C_0 is a homogenous good.

of variety i. Each producer has a productivity level for each input, denoted by A_{ji} (with $j = c, d$). Variety i is then produced according to the following technology:

$$y_i = A_{ci}x_{ci} + A_{di}x_{di}. \tag{2}$$

The use of the dirty energy input generates pollution: more specifically, if the producer uses a quantity x_{di} of dirty energy input, the atmospheric emissions are given by ξx_{di} $(\xi > 0)$.[3]

Before production occurs a firm has the opportunity to innovate in clean and/or dirty technologies. By hiring z_{ji} workers the producer can increase his productivity with input j by a factor $(1 + \eta_j z_{ji})$ $(j = c, d)$. We denote by A^0_{ji} the initial productivity level, so that the end of period productivity is given by

$$A_{ji} = (1 + \eta_j z_{ji})\, A^0_{ji}.$$

At the beginning of the period, the government can implement two types of environmental policies: a subsidy to research in the clean sector q (so that the effective cost of hiring z_{ci} workers to innovate on the clean technology is $(1 - q)\, z_{ci}$), and a tax τ per unit of pollution. The relationship between the consumer and the producer (p_i) prices, is then given by

$$p'_i = p_i + \tau \frac{\xi x_{di}}{y_i}. \tag{3}$$

The timing of moves within the period can be summarized as follows. First, the government decides about research subsidies and pollution tax. Then, producers decide how much to invest in clean and/or dirty innovation. Then, production takes place.

[3] Alternatively one could assume that pollution is proportional to $A_{di}x_{di}$. All our results carry through in this case.

2.2 Equilibrium profits

To simplify the expressions we drop the subscript i in this subsection. The producer chooses the amount of energy inputs in order to maximize his profits:

$$\Pi = p(y)\, y - x_c - x_d. \tag{4}$$

Using (1), (2) and (3), (4) can be rewritten as:

$$\Pi = P^{\frac{\sigma-\beta}{\sigma}} \left(A_c x_c + A_d x_d\right)^{\frac{\sigma-1}{\sigma}} - x_c - (1+\tau\xi)\, x_d.$$

Because the clean and dirty energy inputs are perfectly substitutes the producer only uses the input with the most cost effective technology, thus he uses the clean energy input only if and only if

$$A_c > \frac{A_d}{1+\tau\xi},$$

and the dirty energy input only when $A_c < \frac{A_d}{1+\tau\xi}$.

The optimal input productions, are then given by:

$$x_c = P^{\sigma-\beta} \left(\frac{\sigma-1}{\sigma}\right)^{\sigma} A_c^{\sigma-1},\ x_d = 0 \text{ if } A_c > \frac{A_d}{1+\tau\xi}, \tag{5}$$

$$x_c = 0,\ x_d = P^{\sigma-\beta} \left(\frac{\sigma-1}{\sigma}\right)^{\sigma} \frac{A_d^{\sigma-1}}{(1+\tau\xi)^{\sigma}} \text{ if } A_c < \frac{A_d}{1+\tau\xi}. \tag{6}$$

These expressions show clearly that the higher A_d, or the lower τ, the higher the level of pollution as measured by ξx_d.

Equilibrium profits are then given by

$$\Pi = P^{\sigma-\beta} \frac{(\sigma-1)^{\sigma-1}}{\sigma^{\sigma}} \max\left\{A_c^{\sigma-1}, \left(\frac{A_d}{1+\tau\xi}\right)^{\sigma-1}\right\}. \tag{7}$$

2.3 Innovation decision

Moving back one step, the producer will decide to invest in clean innovation up to the point where[4]

$$\frac{(\sigma-1)^{\sigma}}{\sigma^{\sigma}} P^{\sigma-\beta}\left(A_c^0\right)^{\sigma-1} \eta_c\left(1+\eta_c z_c\right)^{\sigma-2}=1-q, \tag{8}$$

if it is profitable to innovate clean, and he will invest in dirty innovation up to the point where

$$\frac{(\sigma-1)^{\sigma}}{\sigma^{\sigma}} P^{\sigma-\beta}\left(\frac{A_d^0}{1+\tau\xi}\right)^{\sigma-1} \eta_d\left(1+\eta_d z_d\right)^{\sigma-2}=1 \tag{9}$$

if it is profitable to innovate dirty. In particular, we see that clean R&D investment z_c is weakly increasing in the clean research subsidy q, and in the initial clean productivity A_c^0, whereas the dirty R&D investment z_d is weakly decreasing in the rate of pollution tax τ but increasing in the initial dirty productivity A_d^0.

Now, comparing between Π_c and Π_d we see that producers will innovate clean whenever

$$A_c^0(1+\eta_c z_c)>\frac{A_d^0}{1+\tau\xi}(1+\eta_d z_d),$$

where the left hand side (resp. right hand side) is the final productivity conditional upon innovating clean (resp. innovating dirty).

Using (8) and (9), this last condition simply boils down to

$$\frac{A_c^0}{A_d^0}>\frac{\eta_d}{\eta_c}\frac{(1-q)}{(1+\tau\xi)}. \tag{10}$$

In particular this expression shows that producers are more likely to innovate clean when $\eta_c >> \eta_d$ or the larger q and/or the larger τ, or the larger the initial productivity ratio A_c^0/A_d^0. Also,

[4]Here, we implicitly restrict attention to the case where the producer's profit is concave in the technology, which in turn requires that σ be less than 2. This latter assumption can be dispensed with if we assume an innovation cost which is sufficiently convex (instead of linear).

starting from a situation where (10) is violated, a *small increase* in q or τ will have no effect on clean innovation since that condition will remain violated so that producers will keep innovating dirty; yet, the increase in τ will reduce the amount of dirty innovation z_d by (9) - holding the price index P constant-. However, a sufficiently *large increase* in q or τ will make (10) become satisfied, so that all R&D investment will go into clean.

We now compute the price index to investigate the general equilibrium effect. Using (1), (2), (5) and (6), we can express the price index as:

$$P = \frac{\sigma}{\sigma - 1} \left(\int_0^1 \left(\min \left(A_{ci}^{-1}, (1 + \tau\xi) A_{di}^{-1} \right) \right)^{1-\sigma} di \right)^{\frac{1}{1-\sigma}}.$$

or

$$P = \frac{\sigma}{\sigma - 1} \left(\mu_{dirty} (1 + \tau\xi)^{1-\sigma} E \left(A_{di}^{\sigma-1} | dirty \right) + \mu_{clean} E \left(A_{ci}^{\sigma-1} | clean \right) \right)^{\frac{1}{1-\sigma}}$$

where μ_{clean} is the mass of firms producing with the clean technology, μ_{dirty} the mass of firms producing with the dirty technology, and $E \left(A_{di}^{\sigma-1} | dirty \right)$ (resp. $E \left(A_{ci}^{\sigma-1} | clean \right)$) is the expectation of $A_{di}^{\sigma-1}$ conditional upon the firm producing dirty (resp. the expectation of $A_{ci}^{\sigma-1}$ conditional upon the firm producing clean).

Therefore, the price index is increasing in the tax rate and decreasing in the technological levels. A high price index itself favors both types of innovation as long as the elasticity of substitution between varieties, σ, is higher than the price elasticity of the composite good ,β, and it is detrimental to both types of innovation if $\sigma < \beta$ (from equation (7)). Overall, a small increase in the tax rate pushes towards less dirty innovations when most firms have a large stock of clean innovations or when $\beta > 1$.[5] A small increase in the tax rate pushes towards more clean innovation in firms

[5] The LHS of (9) is increasing in τ if and only if:

$$\mu_{dirty} (1 - \beta) E(A_{di}^{\sigma-1} | dirty) (1 + \tau\xi)^{1-\sigma} + \mu_{clean} (1 - \sigma) E(A_{ci}^{\sigma-1} | clean)$$

producing clean as long as $\sigma > \beta$, but towards less innovation in clean when $\sigma < \beta$. In the latter case, the income effect from increasing the carbon tax τ dominates the substitution effect, so that clean and dirty consumption and innovation in dirty and clean goods decrease, even though consumption and innovation in dirty goods still decrease by more than for clean goods.[6]

2.4 Summarizing our main predictions

Abstracting from general equilibrium effects working through P, which, as we have just seen, are a priori ambiguous, our main predictions are:

1. producers have a higher propensity to innovate clean the larger q and/or the larger τ
2. producers have a higher propensity to innovate clean the higher the initial productivity ratio A_c^0/A_d^0, i.e. the higher the stock of clean vs. dirty innovations
3. starting from a situation where (10) is violated so that firms produce and innovate dirty, a *small increase* in q or τ will will reduce the amount of dirty innovation but will have little effect on clean innovation
4. a sufficiently *large increase* in q or τ will push all R&D investment into clean.

Part 3 points to a positive interaction effect between fuel price and the stock of dirty innovation, whereas part 4 points to a positive effect of fuel price squared on the propensity to innovate clean.

is positive.

[6] There are further general equilibrium effects working through the endogenous response of firms' technological levels. These can be ignored when innovation rates are sufficiently small (so that the technological levels don't change much). For instance, an increase in clean research subsidy increases innovation in clean research and therefore decreases the price index, which reduces the amounts of both types of innovation when $\sigma > \beta$ and increases these two amounts when $\sigma < \beta$.

3 Econometric Implementation and Data

3.1 Econometric specification

Our simplest dependent variable is the firm's propensity to currently innovate clean rather than dirty, which we capture by the measure

$$RPAT_{it} = \ln(1 + PATC_{it}) - \ln(1 + PATD_{it})$$

where $PATC_{it}$ and $PATD_{it}$ are the flows of clean and dirty patents filed by firm i in year t. $RPAT_{it}$ is approximately equal to the (log) ratio between clean and dirty patents. This variable allows us to use (log) linear panel data models. We also present results that use the flow of clean and dirty patents separately as dependent variables and estimate such models using both the OLS and fixed effect count data methods.

We regress relative patents $RPAT_{it}$ on

1. Various measures of government policy, P_{it}. The primary measure we consider is carbon pricing for cars. As no country has established meaningful carbon pricing yet, we use tax-adjusted fuel prices in the various countries in our sample, exploiting cross-country variations in taxation and market differentiation. We then test the robustness of our results to explicitly using fuel taxes as the policy variable. Because different firms operate in different markets (for example GM operates primarily on the US market whereas Toyota operates primarily on the Japanese market) they are differently exposed to tax changes in different countries. To take this heterogeneity into account we construct a firm-specific government policy variable. More specifically, we use the firm's pre-sample history of patent filing to assess the relative importance of the various markets the firm is operating in and construct firm-specific weights

on carbon prices from the corresponding markets (we discuss this in more details in the data section).

2. Firm i's lagged clean and dirty patent stocks, which we denote respectively by $KPATC_{it-1}$ and $KPATD_{it-1}$.We construct stocks using the perpetual inventory method, but then show robustness to using non-parametric distributions of patent flows and to considering alternative assumptions over knowledge depreciation rates.

3. The interaction between P_{it} and the stocks of clean and dirty patents. Here we want to analyze the prediction that a firm's innovation response to a change in carbon pricing, depends upon the extent to which the firm is already locked-in with dirty technologies.

4. A vector of controls X_{it} including total country-wide GDP and per capita GDP to control for market size effects which could increase the demand for innovation, or to control for the type of innovation being country-specific (for example, richer countries might be a priori biased towards clean rather than dirty innovation). We also include country fixed effects and year fixed effects and check that our results are robust to the inclusion of country-by-year fixed effects. [7]

5. Firm fixed effects (η_i) to control for correlated unobserved heterogeneity.

We thus consider the regression equation

[7] The firm and time specific nature of the policy variable P_{it} enables us to also condition on a set of country by year fixed effects and still identify the effect of the policy variable on induced innovation.

$$RPAT_{it} = \beta P_{it-k} + \alpha_1 KPATC_{it-1} + \alpha_2 KPATD_{it-1} + \gamma_1 (KPATC_{it-1} * P_{it-k}) \quad (11)$$
$$+ \gamma_2 (KPATD_{it-1} * P_{it-k}) + \Omega . X_{it} + \eta_i + u_{it}$$

where our conjecture is that $\alpha_1 > 0, \alpha_2 < 0, \beta > 0$, $\gamma_1 < 0$, and $\gamma_2 > 0$. We lag the policy variable by k periods as we expect its impact on patenting not to be immediate. In our baseline estimation we let $k = 1$ but show the robustness of other results to increasing the lag.

Subsequently, we separately regress the flows of clean and dirty innovations. In the count data versions of these regressions we estimate Poisson models of the form:

$$PATC_{it} = \exp(\beta^C P_{it-k} + \alpha_1^C KPATC_{it-1} + \alpha_2^C KPATD_{it-1} + \gamma_1^C (KPATC_{it-1} * P_{it-k}) \quad (12)$$
$$+ \gamma_2^C (KPATD_{it-1} * P_{it-k}) + \Omega^C . X_{it} + \eta_i^C + u_{it}^C)$$

and

$$PATD_{it} = \exp(\beta^D P_{it-k} + \alpha_1^D KPATC_{it-1} + \alpha_2^D KPATD_{it-1} + \gamma_1^D (KPATC_{it-1} * P_{it-k}) \quad (13)$$
$$+ \gamma_2^D (KPATD_{it-1} * P_{it-k}) + \Omega^D . X_{it} + \eta_i^D + u_{it}^D)$$

To deal with fixed effects in these non-linear Poisson models, we use the results in Blundell, Griffith and Van Reenen (1999) who argue that using a pre-sample mean scaling estimator is an attractive way of controlling for correlated unobserved heterogeneity in dynamic innovation models[8]. We also compare this method with the Hausman, Hall and Griliches (1984) approach, although

[8]See also Blundell, Griffith and Windmeijer (2002) and Blundell, Griffith and Van Reenen (1995).

this requires strict exogeneity, which is inconsistent with dynamic models with lagged dependent variables on the right-hand side as we have in equations (12) and (13).

3.2 Data

3.2.1 Main dataset

Our data are drawn from the World Patent Statistical Database (PATSTAT) maintained by the European Patent Office. Patent documents are categorized using the international patent classification (IPC) and national classification systems. We have extracted all the patents filed from 1978 to 2007 at the EPO pertaining to "clean" (C) and "dirty" (D) technologies in the automotive industry. "Dirty" includes patents related to the internal combustion engine. "Clean" includes patents specifically related to clean car technologies such as electric or hydrogen vehicles. Our selection of relevant IPC codes for clean technologies relies heavily on previous work by the OECD.[9] The precise description of the IPC codes used to identify relevant patents can be found in Annex 1. The data set includes 12,438 "clean" and 37,103 "dirty" patent applications. In addition to these patents, we have extracted all other patents filed by holders of at least one clean or dirty patent. This represents a total of 746,564 patent applications.

The PATSTAT database reports the name of patent applicants. A common problem with patent data is that the name of patentees often varies, because of spelling mistakes, typographical errors and name variants. To uniquely identify patent holders we use the OECD "HAN" database, which provides a dictionary of harmonized patent applicants' names produced through a computer

[9]See www.oecd.org/environment/innovation

algorithm.[10] We complement the HAN database with manual improvements.[11] As a result, we are able to reduce the number of distinct patent holders from 11,334 to 6,827, 4,366 of which are companies and 2,461 are individuals.[12] For every patent holder we subsequently identify the number of clean, dirty and "other" (i.e. neither clean nor dirty) patent applications filed every year.

3.2.2 Patents as an indicator of innovation

To measure innovation, we use counts of patent applications. The advantages and limitations of patenting as a measure of innovation, have been discussed at length in the literature [13]. For our purpose, the main advantage of using patent data is that these are available at a highly disaggregated level. R&D cannot be disaggregated by type of innovation in this way. Further, R&D is not reported for many small and medium sized firms, especially in Europe (in the US privately listed firms are also exempt from the accounting requirement to report material R&D). In particular we can map innovations in the automotive industry according to specific technologies, such as control systems specially designed for hybrid vehicles. Moreover, the car industry is a large, R&D intensive industry where patents are perceived as an efficient means of protection against innovation. This high propensity to patent innovations makes patenting data a good indicator of innovative activity in the sector.

Patent-based indicators suffer from a number of limitations. The first is that patents are not the only way to protect innovations, although a large fraction of the most economically significant innovations appear to have been patented (Dernis et al., 2001). Another problem with patent-based

[10]The HAN database is only available for patent applicants at the EPO. This is we restricted our data sample to patents applied for at the EPO.

[11]This allows us to match for example Ford Motor Company with Ford Werke, its German subsidiary.

[12]We are also able to match 1736 of these companies with the Orbis database, a rich company characteristics database.

[13]See Griliches (1990) and, for a recent overview, OECD (2009)

indicators is that the value of individual patents is heterogeneous, and the overall distribution of individual patent values is skewed, with most patents having a very low valuation, thus the number of patents does not perfectly reflect the aggregate innovative output. To deal with this problem, we use citation-weighted patenting data in some of our robustness specifications. FinaIly, the number of patents that are granted for a given innovation varies significantly across patent offices. Our decision to consider only patents filed at the European Patent Office is meant to avoid this potential problem.

3.2.3 Constructing tax-adjusted fuel prices

To estimate the impact of a carbon tax on innovation in clean and dirty technologies, we use information on fuel prices and fuel taxes. Data on tax-adjusted fuel prices are available from the International Energy Agency for 25 countries (including some non-OECD countries), from 1978 onwards.[14] Since data are available for both diesel and gasoline fuels, we construct a time-varying country-level fuel price defined as the average of diesel and gasoline prices.

An important issue noted in the previous section is that data on fuel prices are available only at the country level, whereas our empirical analysis requires the policy variable to vary across firms (as we include time dummies). Another issue is that the car market is a global market where government policies abroad might be at least as important for firms' innovation decisions as domestic policies in the country where the firm operates. To deal with these issues we construct a firm-level tax-adjusted fuel price variable for each firm as a weighted average of fuel prices across countries where the firm sells. The weight of each country is in turn determined by the importance of that country as a market outlet for that particular firm. To construct those weights and measure the exposure

[14]The IEA reports some incomplete data for an additional 13 countries.

of a company to a specific market we make use of information on patent families. For every patent applied for at the European Patent Office, we know that the patenting firm has paid the cost of legal protection in a discrete number of countries. For example, a firm may choose to enforce its rights in all EU countries or only in a subset of countries, say Germany and the UK. Similarly, the firm may decide to apply for patent protection in the US but not in smaller markets. Assuming that the country distribution of a firm's patent portfolio is a good indicator of the firm's exposure to the various markets, we can use this distribution information to construct a firm-specific fuel price whose value is computed as the weighted geometric mean of the fuel prices in the relevant markets, with weights equal to the shares of the corresponding countries in the firm's patent portfolio. In addition, in order to make sure that the computed exposures are a (weakly) exogenous source of variation across firms, the weights are calculated using the patent portfolio of each company over the 1978-1985 "pre-sample" period, whereas we run regressions over the period 1986-2007. We then perform robustness tests using different pre-sample periods.

3.2.4 Patent stocks

Following Peri (2005) and Cockburn and Griliches (1988), the patent stock is calculated using the perpetual inventory method. We use the recursive formula:

$$KPATC_{it} = (1 - \delta)KPATC_{it-1} + PATC_{it}$$

We use the same formula for the stocks of dirty patents. We take δ, the depreciation of R&D capital, to be equal to 15%, as is commonly assumed in the literature, but we check the robustness of our results to other values of . [15].

[15] Since the European Patent Office was created in 1978, we do not have patent data prior to 1978. We assume that the pre-1978 growth (g) in patent stock was 15% and assume that the initial patent level was in steady state

A common problem with patent data is that the value of patents is highly heterogeneous across firms and over time. Evidence also suggests that there may be diminishing returns to R&D (Popp, 2002). In order to deal with these problems, we construct alternative measures of patent stocks in which patents are weighted by the number of citations they received in subsequent patents. We should also stress that we rely exclusively on patents filed at the European Patent Office. Since patent filing at the EPO involves higher fees than filings at national offices, our patent stock measure is biased towards higher quality patents.[16]

4 Descriptive statistics

4.1 Aggregate statistics

4.1.1 Trend in innovation activity

Aggregate patenting in clean and dirty technologies has been rising over time. The number of patents in dirty technologies rose steadily between 1980 and 2001 and remained stable since then. The number of clean patents remained very low during 15 years before rising sharply between 1995 and 2002 at an average annual growth rate of 21%. It reached a record-high of 1,203 patents in 2003. The rate of innovation has been stable during the last available years. Figure 1 describes the ratio of clean to dirty patents between 1980 and 2007. While the number of clean patents represented only 10% of the number of dirty patents during the 1980s, this ratio has grown sharply since 1990. Today, around 4 clean patents are filed for every 10 dirty patents.

so we can approximate the patent stock by $PAT_0/(\delta + g)$. Note that, as we perform regressions over the 1986–2007 period, the influence of the discounted initial stocks is small.

[16] It has been empirically demonstrated that the number of countries in which a patent is filed is correlated with other indicators of patent value (see, for example, Lanjouw et al, 1998, Harhoff et al, 2003). On average, EPO patents are validated in 5 European countries and filed in 4 additional patent offices worldwide (Van Zeebroeck, 2010), which shows that the value of patents filed at the EPO is high.

Figure 1: Ratio of clean to dirty patents, 1978-2007

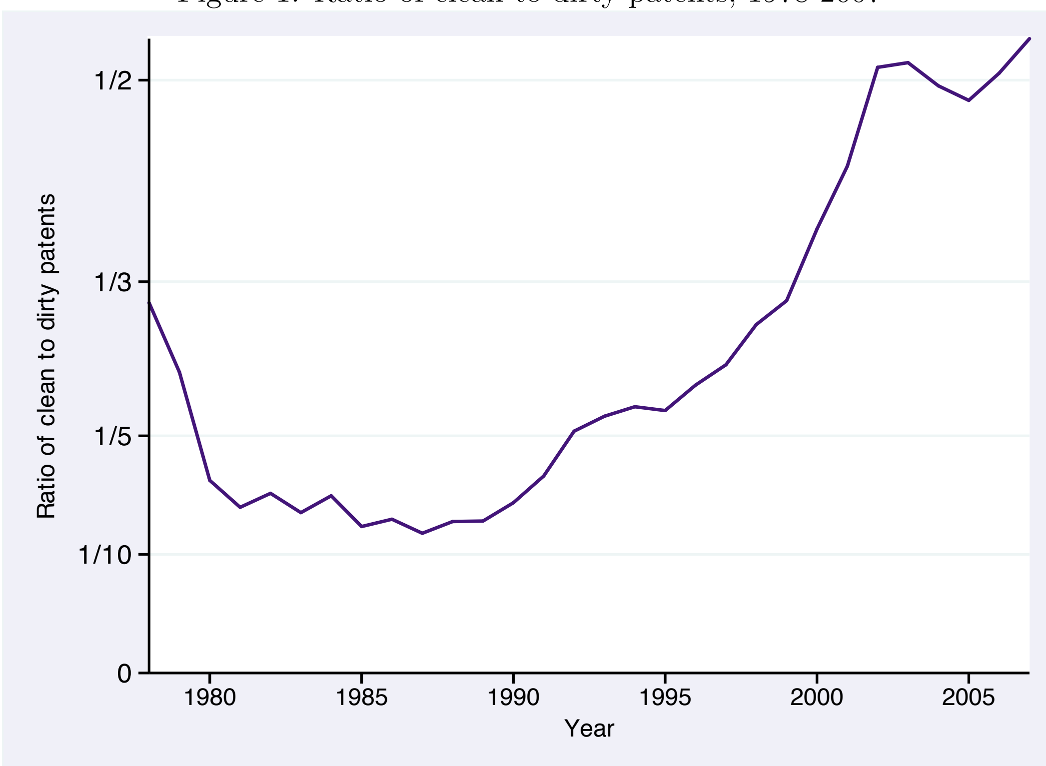

Table 1: Distribution of clean patents by IPC code

IPC code	Description	Number of patents
B60K 1	Arrangement or mounting of electrical propulsion units	373
B60K 6	Arrangement or mounting of hybrid propulsion systems comprising electric motors and internal combustion engines	1378
B60L 3	Electric devices on electrically-propelled vehicles for safety purposes; Monitoring operating variables, e.g. speed, deceleration, power consumption	333
B60L 7	Dynamic electric regenerative braking	80
B60L 11	Electric propulsion with power supplied within the vehicle	951
B60L 15	Methods, circuits, or devices for controlling the traction- motor speed of electrically-propelled vehicles	354
B60R 16	Electric or fluid circuits specially adapted for vehicles and not otherwise provided for	192
B60S 5	Supplying batteries to, or removing batteries from, vehicles	25
B60W 10	Conjoint control of vehicle sub-units of different type or different function	1174
B60W 20	Control systems specially adapted for hybrid vehicles	257
H01M 8	Fuel cells	8065

4.1.2 Technological distribution of patents

Clean and dirty patents are identified using a number of relevant International Patent Classification codes. The tables below provides a detailed breakdown of the patents included in our data set by core IPC code.

4.1.3 Where are clean and dirty inventions protected?

The PATSTAT database holds information about the set of countries where the same invention is patented. For every patent in our data set, we know whether the invention has also been filed

Table 2: Distribution of dirty patents by IPC code

IPC code	Description	Number of patents
F02B	Internal-combustion piston engines; combustion engines in general	9691
F02D	Controlling combustion engines	11392
F02F	Cylinders, pistons, or casings for combustion engines; arrangement of sealings in combusion engines	2637
F02M	Suplying combusion engines with combustible mixtures or constituents thereof	12200
F02N	Starting of combusion engines	1308
F02P	Ignition (other than compression ignition) for internal-combustion engines	1883

(prior to or following the filing of the European patent) at the Japanese Patent Office (JPO), at the US Patent Office (USPTO), or at any other patent office included in PATSTAT. Table 3 provides information on the geographical coverage of clean and dirty innovations. The Patent applications filed at the three main patent offices (EPO, JPO, and USPTO) are referred to as triadic patent families (Dernis and Kahn, 2004). Triadic patent families make up for 50% of dirty inventions and for 59% of clean inventions. Triadic patent families are naturally considered as the most valuable innovations. Table 3 shows that the average value of inventions in our data set is high. Interestingly, 31% of clean inventions are also patented in China. This is almost twice the rate for dirty inventions, 17% of which only make their way to China.

Table 3: Geographical coverage of patent protection

Type of technology	Share of inventions also patented in:			
	USA	Japan	USA & Japan	China
Clean	75%	66%	59%	31%
Dirty	66%	59%	50%	17%

Notes: The patents included in our data set are from the European Patent Office. The table reports the share of patents that are also filed in Japan, USA and China for each category.

4.1.4 Knowledge spillovers as indicated by citation patterns in clean and dirty technologies

When a patent is filed, it must include citations to earlier patents that are related to the new invention. Citations to earlier patents - or backward citations - are indicative of the accumulated knowledge used by the inventor to develop the new invention. We collect this information from the Patstat database. This represents 172,600 citations for all clean and dirty patents included in our data set, which amount to 3.34 citations for the average patent. Dirty patents cite 3.5 patents on average, while clean patents cite only 2.8. Moreover, 40% of clean patents have no citation, whereas the figure is only 29% for dirty patents (which partly reflects the fact that clean patents are on average of more recent vintage than dirty patents). Table 4 reports the distribution of citations between clean and dirty categories. We see that among the patents cited by clean patents, 55% are clean, whereas 4% are dirty. The remaining 41% refer to other - i.e. neither clean or dirty – patents. To get a sense of what these figures imply, suppose that spillovers between clean and dirty patent categories are uniform; i.e. a clean patent – on average – facilitates subsequent clean innovation no more than it would facilitate subsequent dirty innovation and vice versa. Considering that even at the end of our sampling period in 2008 there are about three times as many dirty patents as there

Table 4: Citation patterns

		Cited patent		
		Clean	Dirty	Other
	Clean	55.2%	3.7%	40.1%
Citing patent	Dirty	1.0%	67.7%	31.3%
	Other	0.3%	1.2%	98.5%

Notes: The table reports the distribution of citations across the dirty and clean categories; e.g. 55% of all citations found in clean patents refer to other clean patents whereas 4% refer to a dirty patent. The remaining 41% are citations of other patents.

are clean patents, we would expect that the likelihood of a clean or dirty patent citing a dirty patent be at least three times higher than that of the clean patent citing a clean patent. Interestingly, we find that the likelihood of a clean on clean citations (55%) is almost as high as the likelihood of dirty on dirty citations (68%), suggesting that within category spillovers are vastly higher than between category spillovers. This clearly reflects path-dependence in the direction of innovation.

4.2 Describing companies' patent portfolios

4.2.1 Who are the top inventors in clean and dirty technologies?

Table 5 displays the top 10 inventor companies in clean technologies between 1978 and 2007. In particular it shows the predominance of Japanese and German companies. Table 6 focuses on the last three years of the data sample and shows that the ranking is quite stable over time. That said there are a number of new leaders in clean technology patenting including Peugeot Citroën and BASF Fuel Cells. An interesting finding from tables 5 and 6 is that the vast majority of top innovators in clean technologies are not strictly specialized in this field. Most top companies' patent

Table 5: Main clean patent holders 1978-2007

Company	Clean patents	Dirty patents	Other patents	Total patents
Toyota	719	1578	3414	5711
Nissan Motor	521	818	2602	3941
Honda	405	1044	3216	4665
Siemens	349	1832	33201	35382
Robert Bosch	235	4746	14530	19511
Hitachi	191	815	9685	10691
Ballard Power Systems	183	0	49	232
DaimlerChrysler AG	173	669	3727	4569
Panasonic	162	3	7938	8103
Zahnradfabrik Friedrichshafen	135	57	2826	3018

Notes: The table reports the top 10 clean patent holders between 1978 and 2007. We also report the number of dirty patents and the number of total patents (including clean, dirty and other patents) held by these applicants.

portfolios include both clean and dirty patents. The only exception are Ballard Power System, a fuel cell manufacturer, and Samsung SDI, a battery specialist. Table 7 displays the top 10 inventor companies in dirty technologies between 1980 and 2007. Again we see a predominance of Japanese and German companies.[17] While it is clear that there a number of big companies active in both clean and dirty automotive patenting, computing a Herfindahl Index (HHI) for patenting over 2005 to 2007 for clean innovation we find a HHI of 0.019 and for dirty we find a HHI of 0.031, both of which reflect a low degree of concentration. The top 10 patent holders in clean account for 30% of patents over 2005 to 2007 whereas the corresponding figure is 41% for dirty, suggesting that innovation in dirty is slightly more concentrated than innovation in clean.

[17]Recall that this is based on patents filed at the European Patent Office and US companies tend to file disproportionally more patents in the US than in Europe. This explains why companies such as General Motors - the top patenter in dirty technologies at the US patent office - is not among the top 10 patenters at the EPO.

Table 6: Main clean patent holders 2005-2007

Company	Clean patents	Dirty patents	Other patents	Total patents
Toyota	337	517	1143	1997
Nissan Motor	166	152	633	951
Robert Bosch	123	1058	5187	6368
Samsung Electronics	98	0	4381	4479
Siemens	69	329	7454	7852
Honda	53	226	807	1086
Zahnradfabrik Friedrichshafen	51	8	945	1004
Panasonic Corporation	48	1	3157	3206
Toshiba	44	0	1105	1149
BASF	42	4	4172	4218

Notes: The table reports the top 10 clean patent holders between 2005 and 2007. We also report the number of dirty patents and the number of total patents (including clean, dirty and other patents) filed by these applicants during the same period.

Table 7: Main dirty patent holders 1978-2007

Company	Dirty patents	Clean patents	Other patents	Total patents
Robert Bosch	4746	235	14530	19511
Siemens	1832	349	33201	35382
Toyota	1578	719	3414	5711
Honda	1044	405	3216	4665
Ford	878	88	2825	3791
Nissan Motor	818	521	2602	3941
Hitachi	815	191	9685	10691
DaimlerChrysler	669	173	3727	4569
Renault	644	128	2002	2774
Delphi Technologies	619	89	2696	3404

*Notes:*The table reports the top 10 clean patent holders between 1978 and 2007. We also report the number of dirty patents and the number of total patents (including clean, dirty and other patents) held by these applicants.

4.2.2 Complementarity between clean and dirty innovation

From Table 6 we saw that top clean patent holders also innovate in dirty technologies. Is this result true for all companies or only for leaders in the field? Figure 2 shows the number of firms with respectively at least one dirty patent in their portfolio, at east one clean patent, or at least one clean and one dirty patent. As the figure shows, the proportion of firms innovating both in clean and dirty technologies is small. Interestingly, this proportion varies across technologies. Among firms active in dirty innovation, only 6% on average are also active in clean innovation. However, among clean innovators, the proportion of firms also active in dirty technologies accounts for 29%.[18] This might suggest that the skills needed for the production of clean and dirty inventions are not necessarily complementary. These results also suggest that the likelihood that a firm active in clean innovation be also active in dirty innovation is higher than the likelihood that dirty innovators also work on clean technologies. Clean patents make up 42% of the average patent portfolio of firms having at least one clean patent whereas firms with at least one dirty patent have 65% of dirty patents in their portfolio.

5 Results

In this section we present and discuss the results from estimating the model in equation 11 and then we present some robustness checks.

[18]Note that interpretation of this result is subject to caution as the patent classifications used for each technology may not be equally inclusive. For example, if we identify 90% of all clean vehicles patents but only 50% of all dirty patents, we will underestimate the proportion of clean innovators also active in dirty innovation.

Figure 2: Companies with clean and/or dirty patents

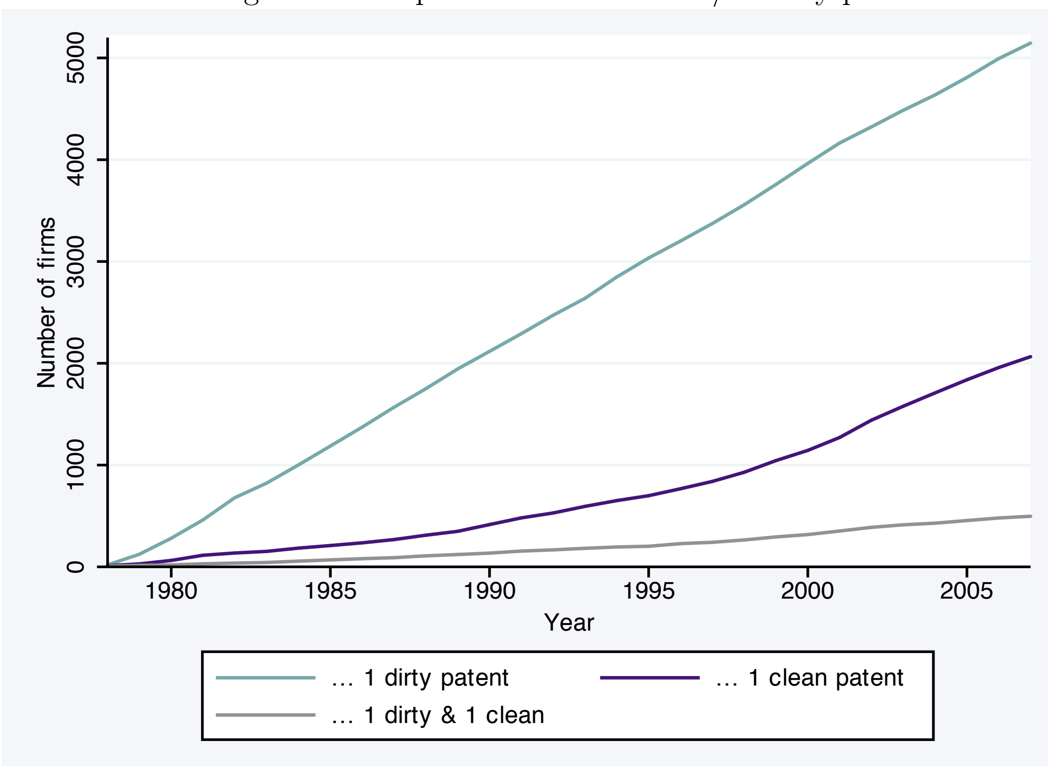

*Notes:*The figure shows the number of applicants that hold respectively at least 1 dirty patent, at least 1 clean patent, and at least 1 clean and 1 dirty patent, across time. Most companies are specialized in either clean or dirty patenting.

5.1 Main Results

Our main results are shown in Table 8 covering 6,422 patent holders.[19] We estimate our panel over 22 years from 1986 to 2007 and use pre-sample information for the 1978-1985 period for the weights (and calculation of the stocks). Note that all estimates include firm fixed effects and country by year dummies. In column (1), we include only the tax inclusive fuel price. This policy variable shows a positive and significant coefficient which is consistent with the idea of endogenous directed technical change: higher carbon prices induce firms to switch their innovation portfolio towards clean and away from dirty innovation. Column (2) includes the lagged patent stocks and suggests strong path dependence in the direction of innovation: namely, a higher stock of clean patents makes future clean innovations relatively more likely, and a higher stock of dirty patents makes future clean innovations less likely. In other words, firms build on their existing stock of technology-specific knowledge to develop new innovations, which in turn can lead to technological lock-in. Column (3) includes both the patent stocks and the fuel price, and shows that all three terms are significant.

In order to account for possible correlations between the fuel price and other macro-level variables, in column (4) we add GDP and GDP per capita as control variables. Both variables are constructed in the same way as the fuel price variable, using the same time-invariant applicant-level weights based on 1978-1985 patent portfolios. Our results remain robust to including these controls. In the last column where we interact the fuel price with the patent stocks variables. Here, the regression results suggest that firms with higher stocks of dirty patents react more strongly to an increased fuel price than firms with no patents at all. And conversely firms with higher stocks of clean patents respond less to an increased fuel price than firms with no patents as indicated by the negative coefficient on the interaction between the stock of clean patents and fuel price. This

[19]Applicants who only filed patents before 1986 are dropped from the regression.

is in line with our third prediction, namely that an increase in fuel price will lower the incentive to pursue dirty innovation to a greater extent for firms that are locked in dirty production and innovation, while it will have a smaller effect on firms that have already accumulated a higher stock of clean technologies.

To summarize: (i) tax-adjusted fuel prices (our proxy for a carbon tax) appear to induce directed technical change towards "clean" innovation; (ii) there is path dependence in the direction of innovation; (iii) tax-policies aimed at inducing clean innovation are, at the margin, more effective for those firms with a higher stock of dirty innovations.

5.2 Extensions

5.2.1 Disentangling clean and dirty innovation

Above we used the ratio between clean and dirty patenting as our dependent variable. Alternatively, we may want to look separately at how clean and dirty innovation respond to government policies and to the stocks of clean and dirty patents, as formalized in equations (12) and (13). One advantage of this latter approach is that we can implement count data models that more properly reflect the nature of the dependent variable. In Table 9 we start by reproducing the OLS log linear models

for clean and dirty patents respectively. Again we see evidence of path dependence in the direction of innovation, with stronger spillovers from past innovation within innovation type (clean versus dirty) than between innovation types. Moreover, a fuel price increase has a positive impact on the level of clean patenting and a negative impact on the level dirty patenting. In Table 10 we repeat the same regressions using a Poisson model without fixed effects (Columns 1 and 2), and in columns (3) and (4) we include fixed effects using the Blundell et al (1999) approach. While there are some changes to the values of point estimates, the qualitative implications are the same in all

Table 8: Regressions of the ratio between clean and dirty patenting

	(1)	(2)	(3)	(4)	(5)
Dep. Variable	Difference between Clean and Dirty Patent applications ln(1+Pc)-ln(1+Pd)				
Fuel Price (including tax)	1.688***		1.235***	0.838***	0.498**
$\ln P_{it-1}$	(0.246)		(0.225)	(0.201)	(0.194)
Stock of clean patents		0.161***	0.159***	0.158***	0.144***
$\ln(1+KPATC_{it-1})$		(0.014)	(0.014)	(0.014)	(0.015)
Stock of dirty patents		-0.085***	-0.084***	-0.080***	-0.046**
$\ln(1+KPATD_{it-1})$		(0.013)	(0.013)	(0.014)	(0.019)
Stock of clean patents X Fuel Price					-0.029
$\ln(1+KPATC_{it-1})$ X $\ln P_{it-1}$					(0.046)
Stock of dirty patents X Fuel Price					0.131***
$\ln(1+KPATD_{it-1})$ X $\ln P_{it-1}$					(0.032)
Controls for population & GDP	no	no	no	yes	yes
Firm Fixed Effects	yes	yes	yes	yes	yes
Year Fixed Effects	yes	yes	yes	yes	yes
Observations	141284	141284	141284	141284	141284
Firms	6422	6422	6422	6422	6422

*Notes:**=significant at the 10% level, **=significant at the 5% level, ***significant at the 1% level. The dependent variable in all columns is the log of the ratio of the number of clean to the number of dirty patents. All columns estimated by OLS with standard errors in parentheses (clustered by firm). All columns include a full set of country dummies and year dummies and a set of company fixed effects. The tax adjusted fuel price variable is a weighted average of 26 country-specific fuel prices where the firm-specific weights are constructed according to the firm's patent portfolio in these countries (see Appendix 1). To construct the weights we use patent data 1978-1985. Regressions estimated 1986-2007. ln(GDP) and ln(GDP per capita) are constructed using the same time-invariant firm-specific weights.

Table 9: Separate regressions for clean and dirty

Dep.Variable	Number of Patent Applications					
	Clean			Dirty		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Fuel Price (including tax)	0.564***	0.307***	-0.006	-0.671***	-0.531***	-0.504***
$\ln(P_{it-1})$	(0.068)	(0.077)	(0.078)	(0.086)	(0.097)	(0.098)
Stock of clean patents	0.216***	0.216***	0.201***	0.057***	0.057***	0.057***
$\ln(1+KPATC_{it-1})$	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
Stock of dirty patents	0.036***	0.039***	0.072***	0.120***	0.120***	0.118***
$\ln(1+KPATD_{it-1})$	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Stock of clean patents X Fuel Price			-0.040***			-0.011
$\ln(1+KPATC_{it-1}) \times \ln(P_{it-1})$			(0.006)			(0.007)
Stock of dirty patents X Fuel Price			0.125***			-0.006
$\ln(1+KPATD_{it-1}) \times \ln(P_{it-1})$			(0.004)			(0.005)
Controls for GDP & Population	no	yes	yes	no	yes	yes
Firm Fixed Effects	yes	yes	yes	yes	yes	yes
Year Controls	yes	yes	yes	yes	yes	yes
Observations	141284	141284	141284	141284	141284	141284
Firms	6422	6422	6422	6422	6422	6422

*Notes:**=significant at the 10% level, **=significant at the 5% level, ***significant at the 1% level. The dependent variable in is the log of one plus the number of clean (respectively dirty) patents. The fuel price variable is a weighted average of 26 country fuel prices where the firm-specific weights are constructed according to the firm's patent portfolio in these countries. To construct the weights we use patent data in the 1978-1985 period and estimate regressions from 1986. GDP and GDP per capita are constructed using the same time-invariant firm-level weights.

specifications. In particular, fuel prices have a positive effect on clean innovation and a negative effect on dirty innovation, whereas our previous regressions where we used relative innovation as the dependent variable, would not rule out the possibility that the effect of a higher fuel price on clean innovation also be negative as long as it would be even more negative on dirty innovation.

Table 10: Regressions with count data models

Dep.Variable	Patent counts			
	Clean	Dirty	Clean	Dirty
	(1)	**(2)**	**(3)**	**(4)**
Fuel Price (including tax)	1.239***	-0.828***	1.360***	-0.329***
$\ln(P_{it-1})$	(0.139)	(0.089)	(0.140)	(0.086)
Stock of clean patents	1.456***	-0.102***	1.462***	-0.070***
$\ln(1+KPATC_{it-1})$	(0.008)	(0.006)	(0.008)	(0.006)
Stock of dirty patents	-0.010*	1.352***	-0.020**	1.431***
$\ln(1+KPATD_{it-1})$	(0.006)	(0.004)	(0.009)	(0.005)
Firm Fixed Effects	no	no	yes	yes
Year Controls	yes	yes	yes	yes
Observations	141284	141284	141284	141284
Firms	6422	6422	6422	6422

*Notes:**=significant at the 10% level, **=significant at the 5% level, ***significant at the 1% level. The dependent variable in is the number of clean (respectively dirty) patents. Columns 1 and 2 are estimated using a poisson count data model. In columns 3 and 4 we add the average level of patenting over the 1978 to 1985 period for both both, clean and dirty patents, as well as a dummy variable capturing if a firm has any patents at all in this period. In this we follow the approach proposed by Blundell et al (1995), to account for fixed differenced in the propensity to patent in the presence of lagged endogenous variables. The fuel price variable is a weighted average of 26 country fuel prices where the firm-specific weights are constructed according to the firm's patent portfolio in these countries. To construct the weights we use patent data in the 1978-1985 period and estimate regressions from 1986. GDP and GDP per capita are constructed using the same time-invariant firm-level weights.

5.2.2 Computing elasticities

The regressions for clean patents alone are closest to the results reported in Popp (2002) who also focuses on clean patenting. Therefore it makes sense to use the coefficients found in the above regressions for a basic comparison of magnitudes. Popp reports coefficients of a regression of the share of clean patents in total patenting on (log) price. If we assume that changes in the price have no effect on total patenting, then these coefficients can be interpreted as the short run price elasticity of clean patenting. Popp reports coefficients ranging from 2.8 to 6%. To compute comparable figures we have to take account of two things. First, Popp's figures refer to the aggregate economy and second, because of the construction of our dependant firm level variable and because we interact prices with patent stocks, our econometric model does not assume that elasticities are constant across firms. Note that from differentiation of our dependant variable with respect to fuel price we find that

$$\frac{\partial ln\left(1+PATC_i\right)}{\partial lnP_i}=\frac{1}{1+PATC_i}\frac{\partial PATC_i}{\partial lnP_i}=\beta_P+\beta_{PC}KPATC_i+\beta_{PD}KPATD_i$$

Hence, we can compute the marginal effect on patenting of a 1% change in price at the firm level as

$$\frac{\partial PATC_i}{\partial lnP_i}=\left(\beta_P+\beta_{PC}KPATC_i+\beta_{PD}KPATD_i\right)\left(1+PATC_i\right)$$

Thus, to get the aggregate price elasticitiy we simply have to aggregate this across firms and divide by the total number of patents:

$$\frac{\frac{\Delta PATC_A}{PATC_A}}{\frac{\Delta P}{P}}=\frac{\sum_i\left(\beta_P+\beta_{PC}KPATC_i+\beta_{PD}KPATD_i\right)\left(1+PATC_i\right)}{\sum_i PATC_i}$$

Doing this in the last year of our sample period and using the parameter estimates reported in Column 3 of Table 9 suggests that the aggregate price elasticity is 11.7%; i.e. somewhat higher than the range of values found by Popp.

5.2.3 Non linear effects

In order to explore nonlinearities in the inducement effect of fuel prices, we include a quadratic price term in our estimations. The results are presented in Table 11. As shown in column (1), the coefficient on the squared fuel price is statistically significant and positive.[20] This finding is consistent with our model which suggests that larger price changes induce non-marginal changes in firm behavior leading to more dramatic responses in innovation outcomes. In column (3) we interact patent stocks with both linear and quadratic price terms. As before we find a stronger response for firms with higher stocks of dirty patents. This general finding of a stronger effect for larger changes in the fuel price persists when we allow for a more flexible functional form for price, but our sense is that the inclusion of the quadratic fuel price variable already captures non-linear effects. Interestingly, when we analyze the non-linear effect of fuel prices separately on clean and on dirty innovations[21] , we find that the quadratic term has a statistically significant effect only on clean innovation. This shows that a strong increase in the price of fuel will strongly encourage clean innovation but does not have much additional effect on dirty innovation.

5.2.4 The tax component of the fuel price

Our estimations have so far used fuel prices as the government policy variable. However the IEA data allow us to isolate the tax component of the fuel price. We compute a firm-level fuel tax variable

[20]This effect is however not robust to including the usual control variables in column (2).
[21]The results are presented in Appendix D.

Table 11: Regressions with quadratic price effects

	(1)	(2)	(3)
Dep. Variable	Ratio between Clean and Dirty Patent applications		
Fuel Price (including tax)	1.278***	0.684**	0.083
$\ln P_{it-1}$	(0.232)	(0.298)	(0.277)
Fuel Price squared	0.222***	-0.102	-0.211
$(\ln P_{it-1})^2$	(0.070)	(0.142)	(0.132)
Stock of clean patents	0.159***	0.158***	0.150***
$\ln(1+KPATC_{it-1})$	(0.014)	(0.014)	(0.016)
Stock of dirty patents	-0.080***	-0.081***	-0.048**
$\ln(1+KPATD_{it-1})$	(0.014)	(0.014)	(0.020)
Stock of clean patents X Fuel Price			-0.180**
$\ln(1+KPATC_{it-1}) \times \ln P_{it-1}$			(0.072)
Stock of dirty patents X Fuel Price			0.191***
$\ln(1+KPATD_{it-1}) \times \ln P_{it-1}$			(0.053)
Stock of clean patents X Fuel Price^2			-0.394***
$\ln(1+KPATC_{it-1}) \times (\ln P_{it-1})^2$			(0.153)
Stock of dirty patents X Fuel Price^2			0.118
$\ln(1+KPATD_{it-1}) \times (\ln P_{it-1})^2$			(0.092)
Controls for population & GDP	no	yes	yes
Firm Fixed Effects	yes	yes	yes
Year Fixed Effects	yes	yes	yes
Observations	141284	141284	141284
Firms	6422	6422	6422

*Notes:**=significant at the 10% level, **=significant at the 5% level, ***significant at the 1% level. The dependent variable in all columns is the log of the ratio of the number of clean to the number of dirty patents. All columns estimated by OLS with standard errors in parentheses (clustered by firm). All columns include a full set of country dummies and year dummies and a set of company fixed effects. The tax adjusted fuel price variable is a weighted average of 26 country-specific fuel prices where the firm-specific weights are constructed according to the firm's patent portfolio in these countries (see Appendix 1). To construct the weights we use patent data 1978-1985. Regressions estimated 1986-2007. ln(GDP) and ln(GDP per capita) are constructed using the same time-invariant firm-specific weights.

using the same weights as before. Arguably, fuel taxes, rather than the overall fuel price, reflect the likely impact of carbon pricing and may suffer less from endogeneity issues. Table 12 contains the results and shows that an increase in fuel taxes leads to relatively more clean innovation. The coefficients in Tables 8 and 12 are similar, suggesting that firms react similarly to increases in the price of oil and to increases in taxes. Moreover, we still obtain significant path-dependence in the direction of innovation.

5.3 Robustness tests

5.3.1 Controlling for country specific time effects

All our results include year dummies to account for global shocks that might influence patenting irrespective of price effects. A concern might be, however, that our results could be driven by other country specific shocks rather than the price or tax channel we suggest above. For example fuel price increases might be correlated with other supply side measures governments are undertaking such as subsidies for clean innovation. We examine this in Table 13 by reporting our basic regression specification while also including country specific year effects. Comparing the results in Table 13 to the ones in Table 8 we see that the estimates change very little.

5.3.2 Alternative categorization of clean versus dirty technologies

In our analysis so far, we took patents related to internal combustion engines to represent dirty technologies in the automotive industry. However, some of the technologies we took to be dirty, are in fact aimed at reducing fuel consumption in combustion engines vehicles. In particular, one can identify patents pertaining to fuel injection technologies, which are explicitly designed to reduce the amount of fuel burnt in combustion engines. This is clearly not the case for innovations leading

Table 12: Regressions using with energy taxes as explanatory variable

	(1)	(2)	(3)	(4)	(5)
Dep. Variable	Difference between Clean and Dirty Patent applications ln(1+Pc)-ln(1+Pd)				
Fuel tax	1.643***		1.116***	0.616**	0.26
$\ln P_{it-1}$	(0.308)		(0.287)	(0.312)	(0.294)
Stock of clean patents		0.161***	0.159***	0.159***	0.218***
$\ln(1+KPATC_{it-1})$		(0.014)	(0.014)	(0.014)	(0.053)
Stock of dirty patents		-0.085***	-0.083***	-0.080***	0.02
$\ln(1+KPATD_{it-1})$		(0.013)	(0.013)	(0.014)	(0.044)
Stock of clean patents X Fuel Tax					0.09
$\ln(1+KPATC_{it-1})$ X $\ln P_{it-1}$					(0.064)
Stock of dirty patents X Fuel Tax					0.113***
$\ln(1+KPATD_{it-1})$ X $\ln P_{it-1}$					(0.041)
Controls for population & GDP	no	no	no	yes	yes
Firm Fixed Effects	yes	yes	yes	yes	yes
Year Fixed Effects	yes	yes	yes	yes	yes
Observations	141284	141284	141284	141284	141284
Firms	6422	6422	6422	6422	6422

*Notes:**=significant at the 10% level, **=significant at the 5% level, ***significant at the 1% level. The dependent variable in all columns is the log of the ratio of the number of clean to the number of dirty patents. All columns estimated by OLS with standard errors in parentheses (clustered by firm). All columns include a full set of country dummies and year dummies and a set of company fixed effects. The tax adjusted fuel price variable is a weighted average of 26 country-specific fuel prices where the firm-specific weights are constructed according to the firm's patent portfolio in these countries (see Appendix 1). To construct the weights we use patent data 1978-1985. Regressions estimated 1986-2007. ln(GDP) and ln(GDP per capita) are constructed using the same time-invariant firm-specific weights.

Table 13: Regressions with country specific time effects

	(1)	(2)	(3)	(4)	(5)
Dep. Variable	Difference between Clean and Dirty Patent applications ln(1+Pc)-ln(1+Pd)				
Fuel Price (including tax)	1.782***		1.351***	0.958***	0.616***
ln P_{it-1}	(0.250)		(0.229)	(0.204)	(0.199)
Stock of clean patents		0.159***	0.157***	0.156***	0.142***
ln(1+$KPATC_{it-1}$)		(0.014)	(0.014)	(0.014)	(0.015)
Stock of dirty patents		-0.086***	-0.084***	-0.081***	-0.047**
ln(1+$KPATD_{it-1}$)		(0.013)	(0.013)	(0.014)	(0.019)
Stock of clean patents X Fuel Price					-0.030
ln(1+$KPATC_{it-1}$) X ln P_{it-1}					(0.046)
Stock of dirty patents X Fuel Price					0.130***
ln(1+$KPATD_{it-1}$) X ln P_{it-1}					(0.032)
Controls for population & GDP	no	no	no	yes	yes
Firm Fixed Effects	yes	yes	yes	yes	yes
Country by Year Fixed Effects	yes	yes	yes	yes	yes
Observations	141284	141284	141284	141284	141284
Firms	6422	6422	6422	6422	6422

*Notes:**=significant at the 10% level, **=significant at the 5% level, ***significant at the 1% level. The dependent variable in all columns is the log of the ratio of the number of clean to the number of dirty patents. All columns estimated by OLS with standard errors in parentheses (clustered by firm). All columns include a full set of country × year dummies and a set of company fixed effects. The fuel price variable is a weighted average of 26 country-specific fuel prices where the firm-specific weights are constructed according to the firm's patent portfolio in these countries (see Appendix 1). To construct the weights we use patent data 1978-1985. Regressions estimated 1986-2007. ln(GDP) and ln(GDP per capita) are constructed using the same time-invariant firm-specific weights.

to marginal improvements in pollution intensities and to marginal reductions in the running costs of internal combustion engines. In this section we test the robustness of our previous results to moving all fuel injection patents from the "dirty" to the "clean" category. This represents 7,532 patent applications. We then perform the estimations on the new data set. As shown in Table 14, all the results presented above hold using this alternative organization of our data set.

5.3.3 Dropping individual patent holders

Our data set includes 6,422 distinct patent holders: 4,181 companies and 2,241 individuals. In order to check that the results are not driven by individuals, we perform a regression where we drop individual patent holders from our data sample. Table 15 shows that all the above results continue to hold. In addition, we find that the impact of higher fuel prices is higher on companies than on the average patent applicants.

5.3.4 Alternative lags

As pointed earlier we have tried various lags for the price variable. Our findings are robust to using larger lags. As shown in Table 16 the coefficient on the fuel price is higher when we use a 2-years lag of the price than when we use a 1-year lag (see Table 8).

5.3.5 Checking for outliers

One needs to check that our results were not driven by outliers. As is commonly the case with patent data, the distribution of patents across applicants is highly heterogeneous with a few companies accounting for a large share of innovations. For this reason we considered trimming (dropping the top 1% of companies in both clean and dirty innovation) or winsorizing these extreme values. Our findings are robust to performing these changes. For example, the coefficient (standard error) on

Table 14: Regressions with alternative definition of clean

	(1)	(2)	(3)	(4)	(5)
Dep. Variable	Difference between Clean and Dirty Patent applications ln(1+Pc)-ln(1+Pd)				
Fuel Price (including tax)	1.334***		1.072***	0.755***	0.609***
$\ln P_{it-1}$	(0.213)		(0.197)	(0.181)	(0.171)
Stock of clean patents		0.126***	0.124***	0.124***	0.121***
$\ln(1+KPATC_{it-1})$		(0.012)	(0.012)	(0.012)	(0.014)
Stock of dirty patents		-0.049***	-0.047***	-0.045***	-0.031**
$\ln(1+KPATD_{it-1})$		(0.010)	(0.010)	(0.010)	(0.015)
Stock of clean patents X Fuel Price					0.005
$\ln(1+KPATC_{it-1}) \times \ln P_{it-1}$					(0.038)
Stock of dirty patents X Fuel Price					0.052*
$\ln(1+KPATD_{it-1}) \times \ln P_{it-1}$					(0.029)
Controls for population & GDP	no	no	no	yes	yes
Firm Fixed Effects	yes	yes	yes	yes	yes
Year Fixed Effects	yes	yes	yes	yes	yes
Observations	141284	141284	141284	141284	141284
Firms	6422	6422	6422	6422	6422

*Notes:**=significant at the 10% level, **=significant at the 5% level, ***significant at the 1% level. The dependent variable in all columns is the log of the ratio of the number of clean to the number of dirty patents. All columns estimated by OLS with standard errors in parentheses (clustered by firm). All columns include a full set of country dummies and year dummies and a set of company fixed effects. The tax adjusted fuel price variable is a weighted average of 26 country-specific fuel prices where the firm-specific weights are constructed according to the firm's patent portfolio in these countries (see Appendix 1). To construct the weights we use patent data 1978-1985. Regressions estimated 1986-2007. ln(GDP) and ln(GDP per capita) are constructed using the same time-invariant firm-specific weights.

Table 15: Regressions for company patent holders only

	(1)	(2)	(3)	(4)	(5)
Dep. Variable	Difference between Clean and Dirty Patent applications ln(1+Pc)-ln(1+Pd)				
Fuel Price (including tax)	2.016***		1.479***	0.993***	0.580**
$\ln P_{it-1}$	(0.314)		(0.284)	(0.265)	(0.255)
Stock of clean patents		0.174***	0.172***	0.172***	0.154***
$\ln(1+KPATC_{it-1})$		(0.014)	(0.014)	(0.014)	(0.016)
Stock of dirty patents		-0.120***	-0.119***	-0.115***	-0.078***
$\ln(1+KPATD_{it-1})$		(0.015)	(0.015)	(0.016)	(0.020)
Stock of clean patents X Fuel Price					-0.050
$\ln(1+KPATC_{it-1}) \times \ln P_{it-1}$					(0.048)
Stock of dirty patents X Fuel Price					0.156***
$\ln(1+KPATD_{it-1}) \times \ln P_{it-1}$					(0.033)
Controls for population & GDP	no	no	no	yes	yes
Firm Fixed Effects	yes	yes	yes	yes	yes
Year Fixed Effects	yes	yes	yes	yes	yes
Observations	91982	91982	91982	91982	91982
Firms	4181	4181	4181	4181	4181

*Notes:**=significant at the 10% level, **=significant at the 5% level, ***significant at the 1% level. The dependent variable in all columns is the log of the ratio of the number of clean to the number of dirty patents. All columns estimated by OLS with standard errors in parentheses (clustered by firm). All columns include a full set of country dummies and year dummies and a set of company fixed effects. The tax adjusted fuel price variable is a weighted average of 26 country-specific fuel prices where the firm-specific weights are constructed according to the firm's patent portfolio in these countries (see Appendix 1). To construct the weights we use patent data 1978-1985. Regressions estimated 1986-2007. ln(GDP) and ln(GDP per capita) are constructed using the same time-invariant firm-specific weights.

Table 16: Using 2-years lag of the price

	(1)	(2)	(3)	(4)	(5)
Dep. Variable	Difference between Clean and Dirty Patent applications ln(1+Pc)-ln(1+Pd)				
Fuel Price (including tax)	1.932***		1.400***	0.883***	0.459**
$\ln P_{it-2}$	(0.281)		(0.260)	(0.217)	(0.201)
Stock of clean patents		0.161***	0.159***	0.159***	0.142***
$\ln(1+KPATC_{it-1})$		(0.014)	(0.014)	(0.014)	(0.016)
Stock of dirty patents		-0.085***	-0.083***	-0.080***	-0.037*
$\ln(1+KPATD_{it-1})$		(0.013)	(0.013)	(0.014)	(0.021)
Stock of clean patents X Fuel Price					-0.032
$\ln(1+KPATC_{it-1})$ X $\ln P_{it-2}$					(0.049)
Stock of dirty patents X Fuel Price					0.143***
$\ln(1+KPATD_{it-1})$ X $\ln P_{it-2}$					(0.035)
Controls for population & GDP	no	no	no	yes	yes
Firm Fixed Effects	yes	yes	yes	yes	yes
Year Fixed Effects	yes	yes	yes	yes	yes
Observations	141284	141284	141284	141284	141284
Firms	6422	6422	6422	6422	6422

*Notes:**=significant at the 10% level, **=significant at the 5% level, ***significant at the 1% level. The dependent variable in all columns is the log of the ratio of the number of clean to the number of dirty patents. All columns estimated by OLS with standard errors in parentheses (clustered by firm). All columns include a full set of country dummies and year dummies and a set of company fixed effects. The tax adjusted fuel price variable is a weighted average of 26 country-specific fuel prices where the firm-specific weights are constructed according to the firm's patent portfolio in these countries (see Appendix 1). To construct the weights we use patent data 1978-1985. Regressions estimated 1986-2007. ln(GDP) and ln(GDP per capita) are constructed using the same time-invariant firm-specific weights.

the fuel price in the specification of 8 column (4) is equal to 1.160 (0.227) in the winsorized version.

5.3.6 Alternative sample period for constructing the weights for policy variable

The choice of the pre-sample period to construct the weights in somewhat arbitrary. Using more pre-sample information improves the precision with which one can compute the policy variable, however at the cost of having fewer remaining years for estimating the main equations. Yet, we have tried different pre-sample periods to calculate the weights. The results remain qualitatively similar. For example, instead of the base case where we use 1978-1985 data we experimented using 1978-1990 data. The coefficient (resp. standard error) on the fuel price in the specification of Table 8 column (4) is 0.921 (resp. 0.160) in the 1978-1990 version.

6 Conclusion

In this paper we used a unique patent data set to see whether there is path dependence in clean versus dirty innovation in the automotive industry. Our findings suggest that firms with past experience in dirty patenting are more likely to pursue dirty innovation activities in the future and conversely for firms that have been more active in clean patenting. We also explored how firms are reacting to exposure to fuel price increases in their product markets, and we found that higher fuel prices induce a bias towards cleaner innovation. We also find that this effect is heterogenous among firms: with prior focus on dirty technologies tend to react more strongly to fuel price increases than firms with prior focus on clean innovation: in line with the idea of technological lock-in, a fuel price increase has a small effect on firms already specialized in clean technologies.

Our analysis could be extended in several directions. One extension would be to see how the direction of innovation reacts to other instruments such as direct subsidies to clean innovation: one prediction of the model is that – contrary to the effect of fuel prices – the effect of clean research subsidies should be stronger for firms which have already innovated in clean technologies. Another extension would be to use micro data to estimate the relative efficiency of R&D investments in clean versus dirty innovation, and also the elasticity of substitution between the two types of production technologies (which the above model took to be infinite). As argued in AABH, these parameters play as important a role as the discount rate in characterizing the optimal environmental policy. These and other equally important extensions are left for further research.

References

Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2009). The Environment and Directed Technical Change. Working Paper 15451, National Bureau of Economic Research.

Blundell, R., Griffith, R., and Van Reenen, J. (1995). Dynamic Count Data Models of Technological Innovation. *The Economic Journal*, 105(429).

Blundell, R., Griffith, R., and Van Reenen, J. (1999). Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms. *The Review of Economic Studies*, 66(3):529–554.

Buonanno, P., Carraro, C., and Galeotti, M. (2003). Endegenous Induced Technical Change and the Costs of Kyoto. *Resource and Energy Economics*, 25(1):11–34.

Cockburn, I. and Griliches, Z. (1988). Industry Effects and Appropriability Measures in the Stock Market's Valuation of R&D and Patents. *American Economic Review*, 78(2):419–23.

Dernis, H., Guellec, D., and van Pottelsberghe de la Potterie, B. (2001). *Using patent counts for cross-country comparisons of technology output*, pages 129–146.

Dernis, H. and Khan, M. (2004). Triadic Patent Families Methodology. OECD Science, Technology and Industry Working Papers 2004/2, OECD, Directorate for Science, Technology and Industry.

Griliches, Z. (1990). Patent Statistics as Economic Indicators: A Survey. *Journal of Economic Literature*, 28(4):1661–1707.

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

Hausman, J., Hall, B. H., and Griliches, Z. (1984). Econometric models for count data with an application to the patents-R&D relationship. *Econometrica*, 52(4):909–938.

Newell, R. G., Jaffe, . B., and Stavins, R. N. (1999). The Induced Innovation Hypothesis and Energy-Saving Technological Change. *Quarterly Journal of Economics*.

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

Peri, G. (2005). Determinants of Knowledge Flows and their Effects on Innovation. *Review of Economics and Statistics*, 87 (2):308–322.

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

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

Stern, N. (2006). *The Stern Review on the Economics of Climate Change*. HM Treasury, London, UK.

Sue Wing, I. (2003). Induced Technical Change and the Cost of Climate Policy. Technical report, MIT Joint Program on the Science and Policy of Global Change.

van Zeebroeck, N. (2010). The puzzle of patent value indicators. *Economics of Innovation and New Technology*.

A Appendix 1: Construction of the Tax-adjusted Fuel Price

The firm-specific tax-adjusted fuel price is

$$P_{it} = \sum_c w_{ic} \ln P_{ct}$$

where P_{ct} is the fuel price (including taxes) in country c at time t and w_{ic} is a firm and country specific weight constructed from patterns of patenting. To construct the weights we use patent data in the pre-sample period (i.e. 1978-1985 when we estimate regressions from 1985 and 1978-1990 when we estimate from 1990). Suppose firm i holds PAT_{ic} patents in country c. We then compute w_{ic} as

$$w_{ic} = \frac{GDP_c \times P\tilde{A}T_{ic}}{\sum_c GDP_c \times P\tilde{A}T_{ic}} \tag{14}$$

where $P\tilde{A}T_{ic} = 1+PAT_{ic}$; where we also take into account the market size of a country by weighting with its GDP. Note that we add a 1 to every country's patent count to deal with firms that had no patents in the 1978 to 1985 period. Such a firm would then be allocated the GDP weighted fuel price across countries. Adding a 1 even to firms that have patents in 1978 to 85 ensures that there are no discrete jumps in our index as we move from firms with no to firms with some patents.

B Appendix 2: Patent categories

Table 17 presents the patent classification codes used to construct the data sets.

Table 17: Definition of IPC Patent classes

Description	IPC code
Electric vehicles	
Electric propulsion with power supplied within the vehicle	B60L 11
Electric devices on electrically-propelled vehicles for safety purposes; Monitoring operating variables, e.g. speed, deceleration, power consumption	B60L 3
Methods, circuits, or devices for controlling the traction- motor speed of electrically-propelled vehicles	B60L 15
Arrangement or mounting of electrical propulsion units	B60K 1
Conjoint control of vehicle sub-units of different type or different function / including control of electric propulsion units, e.g. motors or generators / including control of energy storage means / for electrical energy, e.g. batteries or capacitors	B60W 10/08, 24, 26
Hybrid vehicles	
Arrangement or mounting of plural diverse prime-movers for mutual or common propulsion, e.g. hybrid propulsion systems comprising electric motors and internal combustion engines	B60K 6
Control systems specially adapted for hybrid vehicles, i.e. vehicles having two or more prime movers of more than one type, e.g. electrical and internal combustion motors, all used for propulsion of the vehicle	B60W 20
Regenerative braking	
Dynamic electric regenerative braking	B60L 7/1
Braking by supplying regenerated power to the prime mover of vehicles comprising engine -driven generators	B60L 7/20
Fuel cells	
Conjoint control of vehicle sub-units of different type or different function; including control of fuel cells	B60W 10/28
Electric propulsion with power supplied within the vehicle - using power supplied from primary cells, secondary cells, or fuel cells	B60L 11/18
Fuel cells; Manufacture thereof	H01M 8
Combustion engines	
Combustion engines	F02 (excl. C/G/ K)

C Appendix 4: Disaggregating the non-linear effects

Table 18: Regressions on the level of clean patenting with quadratic term

	(1)	(2)	(3)
Dep. Variable	Ratio between Clean and Dirty Patent applications		
Fuel Price (including tax)	0.609***	0.514**	-0.156
$\ln P_{it-1}$	(0.155)	(0.203)	(0.202)
Fuel Price squared	0.234***	0.138	0.120
$(\ln P_{it-1})^2$	(0.045)	(0.086)	(0.085)
Stock of clean patents	0.216***	0.216***	0.212***
$\ln(1+KPATC_{it-1})$	(0.014)	(0.014)	(0.015)
Stock of dirty patents	0.040***	0.040***	0.073***
$\ln(1+KPATD_{it-1})$	(0.006)	(0.006)	(0.010)
Stock of clean patents X Fuel Price			-0.265***
$\ln(1+KPATC_{it-1}) \times \ln P_{it-1}$			(0.066)
Stock of dirty patents X Fuel Price			0.134***
$\ln(1+KPATD_{it-1}) \times \ln P_{it-1}$			(0.036)
Stock of clean patents X Fuel Price^2			-0.613***
$\ln(1+KPATC_{it-1}) \times (\ln P_{it-1})^2$			(0.159)
Stock of dirty patents X Fuel Price^2			0.013
$\ln(1+KPATD_{it-1}) \times (\ln P_{it-1})^2$			(0.061)
Controls for population & GDP	no	yes	yes
Firm Fixed Effects	yes	yes	yes
Year Fixed Effects	yes	yes	yes
Observations	141284	141284	141284
Firms	6422	6422	6422

*Notes:**=significant at the 10% level, **=significant at the 5% level, ***significant at the 1% level. The dependent variable in all columns is the log of the number of clean patents. All columns are estimated by OLS with standard errors in parentheses (clustered by firms). All columns include a full set of country dummies and year dummies and a set of company fixed effects. The tax adjusted fuel price variable is a weighted average of 26 country-specific fuel prices where the firm-specific weights are constructed according to the firm's patent portfolio in these countries (see Appendix 1). To construct the weights we use patent data 1978-1985. Regressions estimated 1986-2007. ln(GDP) and ln(GDP per capita) are constructed using the same time-invariant firm-specific weights.

Table 19: Regressions on the level of dirty patenting with quadratic term

	(1)	(2)	(3)
Dep. Variable	Ratio between Clean and Dirty Patent applications		
Fuel Price (including tax)	-0.668***	-0.171	-0.239
$\ln P_{it-1}$	(0.225)	(0.298)	(0.267)
Fuel Price squared	0.012	0.240*	0.331**
$(\ln P_{it-1})^2$	(0.069)	(0.144)	(0.133)
Stock of clean patents	0.057***	0.057***	0.062***
$\ln(1+KPATC_{it-1})$	(0.011)	(0.011)	(0.012)
Stock of dirty patents	0.120***	0.120***	0.120***
$\ln(1+KPATD_{it-1})$	(0.016)	(0.016)	(0.020)
Stock of clean patents X Fuel Price			-0.084**
$\ln(1+KPATC_{it-1}) \times \ln P_{it-1}$			(0.041)
Stock of dirty patents X Fuel Price			-0.057
$\ln(1+KPATD_{it-1}) \times \ln P_{it-1}$			(0.050)
Stock of clean patents X Fuel Price^2			-0.219**
$\ln(1+KPATC_{it-1}) \times (\ln P_{it-1})^2$			(0.091)
Stock of dirty patents X Fuel Price^2			-0.105
$\ln(1+KPATD_{it-1}) \times (\ln P_{it-1})^2$			(0.089)
Controls for population & GDP	no	yes	yes
Firm Fixed Effects	yes	yes	yes
Year Fixed Effects	yes	yes	yes
Observations	141284	141284	141284
Firms	6422	6422	6422

*Notes:**=significant at the 10% level, **=significant at the 5% level, ***significant at the 1% level. The dependent variable in all columns is the log of the number of dirty patents. All columns are estimated by OLS with standard errors in parentheses (clustered by firms). All columns include a full set of country dummies and year dummies and a set of company fixed effects. The tax adjusted fuel price variable is a weighted average of 26 country-specific fuel prices where the firm-specific weights are constructed according to the firm's patent portfolio in these countries (see Appendix 1). To construct the weights we use patent data 1978-1985. Regressions estimated 1986-2007. ln(GDP) and ln(GDP per capita) are constructed using the same time-invariant firm-specific weights.