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

Citation

Aghion, Philippe, Antoine Dechezleprêtre, David Hémous, Ralf Martin, and John Van Reenen. 2016. "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry." *Journal of Political Economy* 124 (1) (February): 1–51. doi:10.1086/684581.

Published Version

doi:10.1086/684581

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:27759048>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

NBER WORKING PAPER SERIES

CARBON TAXES, PATH DEPENDENCY AND DIRECTED TECHNICAL CHANGE:
EVIDENCE FROM THE AUTO INDUSTRY

Philippe Aghion
Antoine Dechezleprêtre
David Hemous
Ralf Martin
John Van Reenen

Working Paper 18596
<http://www.nber.org/papers/w18596>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2012

We would like to thank Daron Acemoglu, Robin Burgess, Michael Greenstone, John Hassler, Ken Judd, Rebecca Henderson, Per Krusell, Ed Lazear, Torsten Persson, Estaban Rossi-Hansberg, Nick Stern and Scott Stern for many helpful comments. Participants at seminars in the AEA, Arizona, Birmingham, Cambridge, CIAR, Imperial, INSEAD, LSE, Manchester, Mannheim, NBER, Paris, Stanford, Stirling, Stockholm and Venice have all improved the paper. Financial support has come from the British Academy and the ESRC through the Centre for Economic Performance and the Centre for Climate Change Economics and Policy. Antoine Dechezleprêtre gratefully acknowledges the support of the ESRC under the ESRC Postdoctoral Fellowship Scheme (award no: PTA-026-27- 2756). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2012 by Philippe Aghion, Antoine Dechezleprêtre, David Hemous, Ralf Martin, and John Van Reenen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

Philippe Aghion, Antoine Dechezleprêtre, David Hemous, Ralf Martin, and John Van Reenen
NBER Working Paper No. 18596

December 2012

JEL No. L62,O13,O3

ABSTRACT

Can directed technical change be used to combat climate change? We construct new firm-level panel data on auto industry innovation distinguishing between "dirty" (internal combustion engine) and "clean" (e.g. electric and hybrid) patents across 80 countries over several decades. We show that firms tend to innovate relatively more in clean technologies when they face higher tax-inclusive fuel prices. Furthermore, there is path dependence in the type of innovation both from aggregate spillovers and from the firm's own innovation history. Using our model we simulate the increases in carbon taxes needed to allow clean to overtake dirty technologies.

Philippe Aghion
Department of Economics
Harvard University
1805 Cambridge St
Cambridge, MA 02138
and NBER
paghion@fas.harvard.edu

Antoine Dechezleprêtre
Grantham Research Institute on Climate Change and
the Environment & Centre for Economic Performance
London School of Economics
Houghton Street
London WC2A 2AE
United Kingdom
A.Dechezlepretre@lse.ac.uk

David Hemous
INSEAD
Boulevard de Constance
77305 Fontainebleau, France
david.hemous@insead.edu

Ralf Martin
Imperial College Business School
London SW7 2AZ, UK
and Centre for Economic Performance,
London School of Economics
R.Martin@lse.ac.uk

John Van Reenen
Department of Economics
London School of Economics
Centre for Economic Performance
Houghton Street
London WC2A 2AE
United Kingdom
and NBER
j.vanreenen@lse.ac.uk

1 Introduction

With increasing concern over climate change there is much interest in how new technologies can help reduce greenhouse gas emissions like carbon dioxide (CO_2). Many climate change models assume exogenous technological change (e.g. Stern, 2006), but dealing with the challenge of global warming almost certainly requires more climate change related innovation (e.g. Henderson and Newell, 2011). But what policies can be used to achieve this?

Recent models of climate change with endogenous technical progress suggest that the market will generate insufficient innovation to reduce climate change and too much R&D investment directed at “dirty” technologies. For example, in Acemoglu et al. (2012a) there is path-dependence in the *direction* of technical change. Firms in economies that have innovated a lot in dirty technologies in the past will find it more profitable to innovate in dirty technologies in the future. This path dependency feature when 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 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. For path dependence to occur, however, the associated clean products need to be perceived as substitutes for dirty products by consumers. Autos are a good example of this as electric and hybrid vehicles are potential substitutes for vehicles based on internal combustion engines (another example would be renewable energy plants vs. fossil fuel plants). Furthermore, vehicles are major contributors to greenhouse gas emission. According to the International Energy Agency (IEA) in 2009 road transport accounted for 4.88 GT of CO_2 , which represented 16.5% of global CO_2 emissions (transport as a whole was responsible for 22.1%). Consequently, the automobile sector is the focus of our paper. Note that even with such substitutable products as autos, however, there may be decreasing returns to dirty innovation in which case the market would do part of the job of redirecting technical change towards clean technologies. Thus, estimating the effects of

a firm or country’s history of innovation on future innovation trajectories has to be part of the dynamic evaluation of different climate change policies.

In this paper, we construct a new panel dataset on auto innovations to examine whether firms redirect technical change in response to fuel prices (our proxy for a carbon tax) in the context of path dependent innovation. Our main data are drawn from the European Patent Office’s (EPO) World Patent Statistical database (PATSTAT). These data cover close to the population of all worldwide patents since the mid 1960s. To mitigate the well-known problem that many patents are of very low value, our outcome measure focuses on “triadic” patents which are those that have been taken out in all three of the world’s major patents offices: the European Patent Office, the Japanese Patent Office (JPO) and the United States Patents and Trademark Office (USPTO). We treat multiple filings of an invention (a patent family) as one innovation.¹ We show that our results are robust to alternative ways to controlling for heterogeneous patent values such as using wider categories of patents like “biadics” (US and EPO) or weighting by future citations. Since 1978 when the EPO began there have been around 6,500 triadic patents filed in “clean” auto technologies (electric vehicles, hybrid vehicles, fuel cells for hydrogen vehicles, etc.) compared to about 18,500 triadic patents in “dirty” auto technologies, namely those that affect internal combustion engines. 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.² We also know the geographical location of the inventors listed on the patent so we can examine location based knowledge spillovers.

Our main results can be summarized as follows. First, higher fuel prices induce firms to redirect technical change towards clean innovation and away from dirty innovation. Fuel prices vary across countries and time primarily due to changes in fuel taxes and will have a differential impact across auto firms because of firm-specific profile of market shares due

¹Triadic patents have been used extensively as a way to focus on high-value patents (Grupp et al., 1996; Grupp, 1998; Dernis, Guellec and van Pottelsberghe, 2001; Dernis and Khan, 2004; Guellec and van Pottelsberghe, 2004).

²We do not consider radical innovations in upstream industries such as biofuels, for instance. To explore this is beyond the scope of the current paper which takes the more positive approach of exploring the determinants of clean innovation in vehicles.

to national tastes for home brands, etc. Second, a firm’s propensity to innovate in clean technologies appears to be stimulated by its exposure to past clean technologies (and vice versa for dirty technologies) consistent with the path-dependency hypothesis. We measure this past exposure by (i) aggregate spillovers based on the initial (pre-regression sample) location of the firm’s inventors across different countries combined with the changes over time in the country-specific stocks of clean and dirty auto patents; and (ii) the firm’s own lagged stocks of clean and dirty innovations. Using econometric models that include the fuel price and both individual and aggregate spillover variables we find evidence in favor of firm-specific and aggregate path dependency.

Our paper focuses primarily on radically clean innovations (electric, hydrogen and hybrid cars) rather than on innovations in fuel efficiency (which we label “grey” innovation). While the latter can also contribute to curbing climate change, our focus is motivated by several considerations. First, innovations in fuel efficiency may be partly offset by the resulting increase in driving (the “rebound effect” discussed in West, 2004, for example). Second, there are technological limits on how energy efficient an internal combustion car can be. Third, perfectly separating purely dirty innovations from these energy efficient grey inventions is difficult. Nevertheless, we present regressions where we separate out the grey innovations from the clean and purely dirty and find results that support our model. The effect of higher fuel prices is positive on clean, negative on the purely dirty with the effect on grey innovation in the middle.

Our paper relates to three main strands of the literature. First, in the climate change literature Nordhaus (1994) developed a dynamic Ramsey based model of climate change (the DICE model), which added equations linking production to emissions. Subsequent contributions examined the implications of risk and discounting for the optimal design of environmental policy³ or have looked at the choice between taxes and quotas, building on Weitzman (1974). Recently Golosov, Hassler, Krusell and Tsyvinski (2011) have extended this literature by solving for the optimal policy in a full dynamic stochastic general equilibrium framework.

³In particular see Stern (2006), Weitzman (2007, 2009), Dasgupta (2008), Nordhaus (2007), von Below and Persson (2008), Mendelsohn et al (2008), and Yohe, Tol and Anthoff (2009).

Second, our paper relates to the literature on directed technical change, in particular Acemoglu (1998, 2002; 2008) which itself was inspired by early contributions by Hicks (1932) and Habakkuk (1962).⁴ An application of this idea is the empirical literature linking environmental policy and directed technical change (surveyed in Popp, Newell and Jaffe, 2009). 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. He finds a significant impact from both energy prices and past knowledge stocks on directed innovation. However, since Popp uses aggregate data a concern is that his regressions also capture macro-economic shocks correlated with both innovation and the energy price. Our work uses international firm-level panel data which allows us to exploit differences in the extent to which firms in different countries are affected by policy-induced shocks to the energy price (e.g. from fuel taxes). We control for global macro shocks with a full set of time dummies. Further evidence of directed technical change in the context of energy-saving can be found in Newell, Jaffe and Stavins (1999) which focuses on the air-conditioning industry, or in Crabb and Johnson (2010) who also look at energy-efficient automotive technology. However, neither of these papers use multi-country data nor analyze whether there is path dependency in the direction of technical change. Hassler, Krussell and Olovsson (2011) find evidence for a trend increase in energy saving technologies following oil price shocks. They measure the energy-saving bias of technology as essentially a residual which is attractive as it side-steps the need to classify patents into distinct classes. The advantage of our technology variables is that they are more directly related to the innovation we want to measure. Finally, Acemoglu, Akcigit, Hanley and Kerr (2012) calibrate a microeconomic model of directed technical change to derive quantitative estimates of the optimal climate change policy.

Third, we draw on the extensive literature in industrial organization that estimates the

⁴The theoretical literature on directed technical change is well developed. See for example Messner (1997), Grubler and Messner, (1998), Goulder and Schneider (1999), Manne and Richels (2002), Nordhaus (2002), Van der Zwaan et al. (2002); Buonanno et al (2003), Nordhaus (2002), Sue Wing (2003), Gerlagh (2008) and Gerlagh et al (2009). In contrast, directed technical change has rarely been empirically tested. However, see Acemoglu and Linn (2004), Acemoglu and Finkelstein (2008), Hascic et al (2008) and more recently Hanlon (2012) who shows how relative price changes induced by the Union blockade during the American Civil War induced British textile innovation in machines using Indian cotton.

demand for vehicles (energy efficient and otherwise) as a function of fuel prices and other factors. For example, using around 86 million transactions Alcott and Wozny (2011) find that fuel prices reduce demand for autos, but by less than an equivalent increase in the vehicle price.⁵ Busse, Knittel and Zettelmeyer (2011) use similar data in a more reduced form approach but, by contrast, find a much larger impact of fuel price on auto demand. Although the magnitude of the fuel price effect on demand differs between studies, it is generally accepted that there is an important effect of fuel prices on vehicle demand.

The paper is organized as follows. Section 2 develops a simple model to guide our empirical analysis and Section 3 presents the econometric methodology. The data is presented in Section 4 with some descriptive statistics. Section 5 presents the results, discusses their robustness and some extensions. In Section 6 we explore the implications of our results by simulating the aggregate evolution of future clean and dirty knowledge stocks and analyzing how this evolution would be affected by changes in carbon taxes. Section 7 concludes.

2 Theoretical Model

In this section we present a simple model to guide our empirical analysis. This model rationalizes path dependence in firms' own knowledge stock as well as the impact of a change in the price of fuel on clean and dirty innovation. We then show how one can add knowledge spillovers and energy efficiency innovations ("grey innovations") to our framework.

2.1 Basic framework

We consider a one-period model of an economy where consumers derive utility from an outside good and from motor vehicle services. Utility is quasi-linear with respect to the outside good C_0 (chosen as the numeraire) and β is the elasticity of consumption of motor vehicle services with respect to its index price. This parameter measures the degree of substitutability between motor vehicle services and the outside good.

⁵They argue that this is a behavioral bias causing consumers to undervalue fuel price changes. Readers are referred to this paper for an extensive review of the literature on fuel prices and the demand for autos.

To consume motor vehicle services, consumers need to buy cars and fuel (call it a “dirty car bundle”) or cars and electricity (call it a “clean car bundle”). Utility is then given by:

$$U = C_0 + \frac{\beta}{\beta - 1} \left(\left(\int_0^1 Y_{ci}^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1} \frac{\varepsilon-1}{\varepsilon}} + \left(\int_0^1 Y_{di}^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1} \frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1} \frac{\beta-1}{\beta}},$$

where the consumption of variety i of clean cars together with the corresponding clean energy (electricity) is:⁶

$$Y_{ci} = \min(y_{ci}, \xi_c e_{ci});$$

the consumption of variety i of dirty cars together with the corresponding dirty energy (fuel) is:

$$Y_{di} = \min(y_{di}, \xi_d e_{di});$$

where e_{zi} is the amount of energy consumed for variety i of type- z car where $z = c, d$, i.e. $z = \text{Clean}, \text{Dirty}$; ε is the elasticity of substitution between the clean and dirty cars and σ is the elasticity of substitution among varieties within each type of car. Finally, ξ_c (respectively ξ_d) is the energy efficiency of clean (respectively dirty) cars.

We assume that $1 < \varepsilon \leq \sigma$, so that clean cars are more substitutable with each other than with dirty cars, and that $\varepsilon > \beta$: the elasticity of substitution between clean and dirty cars is larger than the price elasticity for motor vehicle services.

We denote by f_c the price of electricity and f_d the price of fuel. In the first part of the analysis innovation will be cost saving for producers, but later we investigate energy saving innovation. Cars are produced by local monopolists, and each monopolistic firm i produces a fixed number of clean and/or dirty varieties. The monopoly producer of variety i of a type- z car produces A_{zi} cars using one unit of outside good as input.

To complete the description of the model, we need to specify the innovation technology. We assume that at the beginning of the period, by incurring total R&D cost $\frac{1}{2}\psi x_{zi}^2$ in the

⁶The analysis can be easily extended to the case where we allow a less stark form of complementarity. For example, consider a CES instead of Leontief form:

$$Y_{zi} = (y_{zi}^{\frac{\rho-1}{\rho}} + (\xi_z e_{zi})^{\frac{\rho-1}{\rho}})^{\frac{\rho}{\rho-1}}, \rho < 1.$$

outside good the producer of variety i of type- z car can increase productivity according to:

$$A_{zi} = (1 + x_{zi}) A_{zi0} \text{ for } z \in \{c, d\}, \quad (1)$$

where A_{zi0} is the initial productivity for producing that type of cars.⁷

The timing is very simple: at the beginning of the period producers invest in R&D and innovate; at the end of the period, given their productivities resulting from R&D activities, producers make production decisions to maximize profits.

2.2 Solving the model

2.2.1 Equilibrium profits

Define the price indexes for dirty and clean car bundles as:

$$P_z = \left(\int_0^1 \left(p_{zi} + \frac{f_z}{\xi_z} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}, \text{ for } z \in \{c, d\}.$$

The inverse demand curves for clean and dirty cars, which simply result from utility maximization subject to budget constraint, are given by:

$$y_{zi} = \left(p_{zi} + \frac{f_z}{\xi_z} \right)^{-\sigma} P_z^{\sigma-\varepsilon} (P_c^{1-\varepsilon} + P_d^{1-\varepsilon})^{\frac{\varepsilon-\beta}{1-\varepsilon}} \text{ for } z \in \{c, d\}. \quad (2)$$

For given (end-of-period) productivity A_{zi} , the producer of variety i of type- z car solves:

$$\pi_{zi} = \max_{y_{zi}} \left\{ p_{zi} y_{zi} - \frac{1}{A_{zi}} y_{zi} \right\}$$

where y_{zi} for $z \in \{c, d\}$ is given by (2).

This yields the following expression for the equilibrium profit of the corresponding car producer:

$$\pi_{zi} = \frac{(\sigma-1)^{\sigma-1}}{\sigma^\sigma} \left(\frac{1}{A_{zi}} + \frac{f_z}{\xi_z} \right)^{1-\sigma} P_z^{\sigma-\varepsilon} (P_c^{1-\varepsilon} + P_d^{1-\varepsilon})^{\frac{\varepsilon-\beta}{1-\varepsilon}} \text{ for } z \in \{c, d\}. \quad (3)$$

In particular an increase in the firm's productivity A_{zi} or a reduction in the price of energy f_z increases the equilibrium profit π_{zi} since we assumed $\sigma > 1$.

⁷Innovation in dirty technologies as modeled here is not directed at improving energy efficiency. Efficiency-enhancing innovations will be introduced in subsection 2.3 below.

2.2.2 Equilibrium innovation efforts and path-dependence

Moving back to the beginning of period the equilibrium innovation intensities x_{zi} solve

$$\max_{x_{zi}} \left\{ \pi_{zi} - \frac{1}{2} \psi x_{zi}^2 \right\}$$

where π_{zi} is given by equation (3) and A_{zi} satisfies the growth equation (1). For ψ sufficiently large the solution to this maximization problem is uniquely given by the first order condition. The resulting x_{zi} satisfies:

$$x_{zi} \propto \frac{1}{A_{zi0} (1 + x_{zi})^2} \left(\frac{1}{A_{zi0} (1 + x_{zi})} + \frac{f_z}{\xi_z} \right)^{-\sigma}. \quad (4)$$

In particular, the equilibrium innovation intensities x_{zi} increases in the firm's corresponding technology stocks A_{zi0} as long as the elasticity of substitution σ is sufficiently large or if the price of fuel f_z represents a sufficiently small share of the total costs of a car.⁸ Then, there is *path dependence with respect to the firm's own innovation history*.

The ambiguity of the effect of the firm's own technology stock A_{zi0} on the firm's innovation incentives reflects two counteracting forces of a higher A_{zi0} . On the one hand, the term $\frac{1}{A_{zi0}(1+x_{zi})^2}$ reflects the negative marginal effect this has on the price charged by the producer; on the other hand, the term $\left(\frac{1}{A_{zi0}(1+x_{zi})} + \frac{f_z}{\xi_z} \right)^{-\sigma}$ reflects the positive marginal effect of a higher A_{zi0} on the demand for bundle (i, z) . This latter effect is strongest when the price of fuel f_z represents a sufficiently small share of the total costs of a car or when the elasticity of substitution σ between variety (i, z) and other varieties of the z -type cars is sufficiently large, which in turn implies that a lower price allows producer (i, z) to capture more market share from other (j, z) producers.

⁸The precise condition is $\frac{(\sigma-1)}{A_{zi0}(1+x_{zi})} > \frac{f_z}{\xi_z}$. A back of the envelope calculation shows that this condition is likely to be satisfied in practice. In our set-up the price of a new car is given by $p_{zi} = \frac{\sigma}{\sigma-1} \frac{1}{A_{zi}} + \frac{1}{\sigma-1} \frac{f_z}{\xi_z}$, so that the price ratio between fuel expenditure and a new car is equal to $x = \frac{f_z}{\xi_z} / \left(\frac{\sigma}{\sigma-1} \frac{1}{A_{zi}} + \frac{1}{\sigma-1} \frac{f_z}{\xi_z} \right)$. The average price of a new car in the US is \$26,850 with a fuel efficiency of 33.8 miles per gallon and Americans drive on average 14,500 miles per year (source: http://www.bts.gov/publications/national_transportation_statistics/). Assuming that fuel costs \$4 a gallon (with a price increasing at the Hotelling rate), and that a car lasts for 10 years, we obtain a price ratio of $x = 0.64$. Simple algebra gives that $\frac{f_z}{\xi_z} / \left(\frac{\sigma-1}{A_{zi}} \right) = \frac{\sigma}{(\sigma-1)(\frac{\sigma-1}{x}-1)}$ which is equal to 0.70, so that $\frac{f_z}{\xi_z} < \frac{\sigma-1}{A_{zi}}$, when the elasticity of substitution σ is equal to 3. An elasticity of substitution of 3 seems to be a low value: in the model the elasticity of substitution is the same as the price elasticity for car varieties, which Berry, Levinsohn and Pakes (1995) estimated to lie between 5.05 and 37.49 (depending on the car variety).

2.2.3 Redirecting innovation through changes in the fuel price

We now investigate the impact of a change in the fuel price on clean and dirty innovations. Totally differentiating equation (4) for $z = c$ with respect to the fuel price, and then using the notation $\widehat{X} = \frac{dX}{X}$, we obtain:

$$(1 - \omega) \widehat{x}_{ci} = \left(\sigma - \varepsilon + (\varepsilon - \beta) \frac{P_c^{1-\varepsilon}}{P_c^{1-\varepsilon} + P_d^{1-\varepsilon}} \right) \widehat{P}_c + (\varepsilon - \beta) \frac{P_d^{1-\varepsilon}}{P_c^{1-\varepsilon} + P_d^{1-\varepsilon}} \widehat{P}_d, \quad (5)$$

where $\omega \equiv \frac{d}{dx_{zi}} \ln \left(\frac{1}{(1+x_{zi})^2} \left(\frac{1}{A_{ci0}(1+x_{ci})} + \frac{f_z}{\xi_z} \right)^{-\sigma} \right) < 1$.⁹

For sufficiently small innovation intensities, one can neglect the indirect impact of an increase in fuel price via the innovation response of other firms, so that $\widehat{P}_c \approx 0$.¹⁰ Then we approximately have:

$$\widehat{x}_{ci} \propto (\varepsilon - \beta) \frac{P_d^{1-\varepsilon}}{P_c^{1-\varepsilon} + P_d^{1-\varepsilon}} \widehat{P}_d.$$

This in turn implies that the equilibrium intensity of clean innovation increases with fuel price since we assumed $\varepsilon \geq \beta$. Intuitively, a higher fuel price makes dirty car bundles more expensive; this in turn might favor the demand either for clean car bundles or for the outside good. It will boost the demand for clean car bundles if the elasticity of substitution between dirty and clean car bundles is higher than the price elasticity of motor-vehicle services as a whole, as we assumed.

Similarly, we get:

$$(1 - \omega) \widehat{x}_{di} \approx \sigma \left(\widehat{P}_d - \frac{\frac{f_d}{\xi_d}}{\frac{1}{A_{di}} + \frac{f_d}{\xi_d}} \widehat{f}_d \right) - \left(\varepsilon \frac{P_c^{1-\varepsilon}}{P_c^{1-\varepsilon} + P_d^{1-\varepsilon}} + \beta \frac{P_d^{1-\varepsilon}}{P_c^{1-\varepsilon} + P_d^{1-\varepsilon}} \right) \widehat{P}_d \quad (6)$$

once we neglect the indirect impact of a change in fuel price on the price indexes working through the innovation response.

The first term captures a reallocation effect *among* varieties of dirty cars, from most to least productive dirty car producers. This term is indeed equal to zero if all firms had the same dirty technologies, otherwise it has the sign of $\widehat{f}_d$ for the least productive dirty firms and the opposite sign for the most productive dirty firms.

⁹That ω be less than 1 follows from the fact that at the equilibrium the left-hand side of (4) crosses the right-hand side from below.

¹⁰Another reason to neglect this indirect impact is that firms typically operate in several markets, with different exposures to each market for each firm. Therefore the allocation of innovation of the competitors does not depend only on the fuel price in a given country but also on the fuel price in other countries.

The second term captures a substitution effect *between* clean and dirty car producers. This term has the opposite sign from that of $\widehat{f_d}$: namely, an increase in fuel price reduces the benefit of dirty innovation both because it induces substitution towards clean cars and because it reduces the overall consumption of cars.

Overall, an increase in fuel price should decrease dirty innovation intensity (as long as the firm's initial productivity in dirty technologies is not too low compared to that of other firms) and it should increase clean innovation intensity.

2.3 Extensions

Knowledge spillovers. In the empirical part of the paper we investigate not only the effects of firms' own past knowledge but also the effects of aggregate knowledge spillovers across firms in the country where innovation occurs. To introduce the possibility of such aggregate spillovers in our model, suppose the existence for each firm i of knowledge spillovers from a set Ω_i of neighboring varieties of cars of the same type $z = c, d$, so that abstracting from innovation firm i 's initial productivity is:

$$\widetilde{A}_{zi0} = A_{zi0}(1 + \eta(\overline{A}_{zi0}))$$

where

$$\overline{A}_{zi0} = \int_{\Omega_i} A_{zj0} dj$$

and η is an increasing function.¹¹

Equilibrium innovation x_{zi} is now given by:

$$x_{zi} \propto \frac{1}{\widetilde{A}_{zi0}(1 + x_{zi})^2} \left(\frac{1}{\widetilde{A}_{zi0}(1 + x_{zi})} + \frac{f_z}{\xi_z} \right)^{-\sigma} \quad (7)$$

which again is increasing in the aggregate initial knowledge variable $\overline{A}_{zi0}$ if σ is sufficiently large.

“Grey” innovation. Innovations in the car industry may also involve improvements in energy efficiency. Innovations that increase energy efficiency for clean cars would not react differently to an increase in fuel price than cost reducing innovations.

¹¹Our modeling of knowledge spillovers is dictated by our empirical estimation strategy, whereby for each patenting firms we compute the stock of patents generated by scientists who are geographically close.

Now, let us analyze innovations that improve *dirty* energy efficiency. More specifically, suppose that at the beginning of every period, a firm can increase its energy efficiency from ξ_{di0} to

$$\xi_{di} = (1 + x_{\xi i}) \xi_{di0}$$

if it spends $\frac{1}{2}\psi x_{\xi i}^2$ units of the outside good in R&D. We refer to this type of innovations as “grey” innovations.

For sufficiently high ψ , i.e for small innovation intensity the solution to the maximization problem is uniquely defined by the first order condition, totally differentiating the latter with respect to the fuel price leads to:

$$\widehat{x_{\xi i}} \propto \sigma \left(\widehat{P_d} - \frac{\frac{f_d}{\xi_{di}}}{\frac{1}{A_{di}} + \frac{f_d}{\xi_{di}}} \widehat{f_d} \right) - \left(\varepsilon \frac{P_c^{1-\varepsilon}}{P_c^{1-\varepsilon} + P_d^{1-\varepsilon}} + \beta \frac{P_d^{1-\varepsilon}}{P_c^{1-\varepsilon} + P_d^{1-\varepsilon}} \right) \widehat{P_d} + \widehat{f_d}.$$

This expression is similar to equation (6) except for the last term which captures a direct positive effect of an increase in the fuel price on energy efficiency innovation. The overall impact of an increase in the fuel price on grey innovation is therefore ambiguous, whereas we saw that it is unambiguously positive on clean innovation. The reason for the ambiguous effect is that on the one hand an increase in fuel price both, reduces the demand for dirty cars and therefore the profitability of producing (and therefore innovating) in the dirty sector altogether, but on the other hand it induces dirty firms to save more on fuel energy by improving dirty energy efficiency.

We use the expression “grey” as the impact of these innovations on the environment is also ambiguous.¹² On the one hand these innovations increase energy efficiency and therefore reduce the amount of fuel consumption per car; on the other hand these innovations make fossil fuel cars cheaper, thereby increasing total consumption of these cars.¹³

¹²Formally, one obtains

$$e_{di} = \frac{y_{di}}{\xi_{di}} = \frac{\sigma}{\sigma - 1} \frac{1}{\xi_{di}} \left(\frac{1}{A_{di}} + \frac{f_d}{\xi_{di}} \right)^{-\sigma} P_d^{\sigma-\varepsilon} (P_c^{1-\varepsilon} + P_d^{1-\varepsilon})^{\frac{\varepsilon-\beta}{1-\varepsilon}},$$

so that replacing g_i with e_{di}

$$\frac{dg_i}{d\xi_{di}} = \left((\sigma - 1) \frac{f_d}{\xi_{di}} - \frac{1}{A_{di}} \right) \frac{1}{\xi_{di}^2} \left(\frac{1}{A_{di}} + \frac{f_d}{\xi_{di}} \right)^{-\sigma-1} \frac{\sigma}{\sigma - 1} P_d^{\sigma-\varepsilon} (P_c^{1-\varepsilon} + P_d^{1-\varepsilon})^{\frac{\varepsilon-\beta}{1-\varepsilon}},$$

which is ambiguously signed. The expression is negative if the price of fuel is sufficiently low relative to other costs, but it is positive if the elasticity of substitution across cars is sufficiently large.

¹³Empirically, this latter “rebound effect” is estimated at around 20 - 25% (see for instance Small and

2.4 Summarizing our main predictions

We have presented a simple model of clean versus dirty innovation, in which we predict that:

1. producers have a higher propensity to innovate in clean relative to dirty technologies the larger the fuel price f_d .
2. for σ sufficiently high or the price of fuel f_d sufficiently small, producers have a higher propensity to innovate in clean technologies:
 - (i) the higher A_c^0 and the lower A_d^0 , i.e. the higher their initial stock of clean compared to dirty technologies (lagged own innovation stocks);
 - (ii) the higher the aggregate level of clean (or the lower the aggregate level of dirty) technology in neighboring varieties (spillovers).

We now confront these three predictions with our panel data on clean and dirty innovation in the auto industry.

3 Econometrics

3.1 General approach

Consider the following Poisson specification for the determination of firm innovation in clean technologies: ¹⁴

$$PAT_{Clean,it} = \exp(\beta_{C,P} \ln FP_{it-1} + A_{C,it-1}) + u_{C,it} \quad (8)$$

where $PAT_{Clean,it}$ is the number of patents applied for in clean technologies by firm i in year t ; $A_{C,it}$ is the firm's knowledge stock relevant for clean innovation, which depends both upon its own stocks of past clean and dirty innovation and the aggregate spillovers from other firms (discussed below); $u_{C,it}$ is an error term; $\exp(.)$ is the exponential operator; and FP_{it} is fuel price (the empirical equivalent of f_d in the theory section). We lag prices

Van Dender, 2007), however some studies have estimated much larger rebound effects (87% in West, 2004).

¹⁴In our regressions we use an equivalent equation for dirty technologies. We initially discuss only one of these equations to simplify the notation.

and knowledge by a year to reflect delayed response and also mitigate contemporaneous feedback effects.¹⁵ In the robustness section we show this form is reasonable comparing it to alternative dynamic representations using other lag structures and the Popp (2002) approach.

The fuel price has independent variation across time and countries primarily because of country-specific taxes and we show the robustness of our results to using fuel taxes instead of prices. The profile of car sales across countries differs across auto firms. For example, GM has some “home bias” towards the US market whereas Toyota has a “home bias” towards the Japanese market, i.e. they sell more in these countries than one would expect from country observables alone. Thus, different firms are likely to be differently exposed to tax changes in different countries and the fuel price has a firm-specific component. This firm-specific difference in market shares could be because of product differentiation and heterogeneous tastes or it may be because of government policies to promote domestic firms (e.g. trade barriers). To take this heterogeneity into account 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 fuel prices for the corresponding market. Simply put, an unexpected increase in Japanese fuel taxes will have a more salient impact on car makers with a bigger market share in Japan than those with a smaller market share. We discuss this in more detail in Section 4.

We parameterize the firm’s total knowledge stock as:

$$A_{Cit} = \beta_{C,1} \ln SPILL_{C,it} + \beta_{C,2} \ln SPILL_{D,it} + \beta_{C,3} \ln K_{C,it} + \beta_{C,4} \ln K_{D,it} \quad (9)$$

The firm’s knowledge will likely depend on its own past history of innovation and we denote this as $K_{Clean,it}$ (own stock of clean innovation) and $K_{Dirty,it}$ (own stock of dirty innovation).¹⁶ In addition to building on its own past innovations firms will also “stand on

¹⁵In principle, the price should be the firm’s expectation of the future evolution of the fuel price based on the information set at time of making the innovation investment decision. Fuel prices appear to be well approximated by a random walk process (e.g. Anderson et al, 2011a,b), so given our assumption that decisions are made on $t - 1$ information, lagged prices should be a sufficient statistic for this expectation.

¹⁶We construct stocks using the perpetual inventory method, but show robustness to using patent flows and to considering alternative assumptions over knowledge depreciation rates. Some firms have zero lagged knowledge stock in some periods, so we also add in three dummy indicator variables for when lagged clean stock is zero, lagged dirty stock is zero or both are zero.

the shoulders of giants”, so we allow their knowledge stock to depend on spillovers from other firms both in clean ($SPILL_{C,it}$) and dirty technologies ($SPILL_{D,it}$). We use stocks of economy wide patents to construct these country-specific spillover measures. Drawing on the evidence that knowledge has a local component (e.g. Jaffe, Trajtenberg, and Henderson, 1993) we use the firm’s distribution of inventors across countries to weight the country spillover stocks. In other words, if the firm has many inventors in the US regardless of whether the headquarters of the firm is in Tokyo or Detroit, then the knowledge stock in the US is given a higher weight (see Section 4).

There are of course other factors that may influence innovation in addition to fuel prices and the past history of innovation. These include policy measures such as government R&D subsidies for clean innovation and regulations over emissions and the size and wealth of the countries a firm is expecting to sell to (proxied by GDP per capita and GDP) which we denote by the vector $w_{C,it}$. We also allow for unobservable factors by introducing a firm fixed effect ($\eta_{C,i}$), a full set of time dummies (T_{Ct}) and an error term ($u_{C,it}$) that is uncorrelated with the right hand side variables. Adding these extra terms and substituting equation (9) into (8) gives us our main empirical equation for clean innovation:

$$\begin{aligned} PAT_{Clean,it} = & \exp(\beta_{C,P} \ln FP_{it-1} + \beta_{C,1} \ln SPILL_{C,it-1} + \beta_{C,2} \ln SPILL_{D,it-1} \\ & + \beta_{C,3} \ln K_{C,it-1} + \beta_{C,4} \ln K_{D,it-1} \\ & + \beta_{C,w} w_{it} + T_{C,t}) \eta_{C,i} + u_{C,it} \end{aligned} \quad (10)$$

Symmetrically, we can derive an equation for dirty innovation:

$$\begin{aligned} PAT_{Dirty,it} = & \exp(\beta_{D,P} \ln FP_{it-1} + \beta_{D,1} \ln SPILL_{C,it-1} + \beta_{D,2} \ln SPILL_{D,it-1} \\ & + \beta_{D,3} \ln K_{C,it-1} + \beta_{D,4} \ln K_{D,it-1} \\ & + \beta_{D,w} w_{it} + T_{D,t}) \eta_{D,i} + u_{D,it} \end{aligned} \quad (11)$$

The theory yields predictions on the signs of the coefficients in these two equations. If higher fuel prices induces more clean than dirty innovation then the marginal effect of the fuel price must be larger on clean innovation than on dirty innovation: $\beta_{C,P} > \beta_{D,P}$ and we would expect that $\beta_{C,P} > 0$ and $\beta_{D,P} < 0$.¹⁷ Next, for there to be path dependence in the

¹⁷Note that these two stronger second conditions are not necessary for induced (redirected) technical

direction of innovation it should be the case that (*ceteris paribus*) firms that are exposed to more dirty spillovers become more prone to conduct dirty innovation in the future: i.e. $\beta_{D,2} > 0$ and $\beta_{D,2} > \beta_{C,2}$. In the clean innovation equation we have $\beta_{C,1} > 0$ and $\beta_{C,1} > \beta_{D,1}$. And path dependence should involve similar effects working through a firm's own accumulated knowledge: $\beta_{D,4} > 0$ and $\beta_{D,4} > \beta_{C,4}$. ($\beta_{C,3} > 0$ and $\beta_{C,3} > \beta_{D,3}$.) Also, we would expect that the positive effect of dirty spillovers and dirty knowledge stocks on dirty innovation be larger than the effects of clean spillovers and clean knowledge stocks: $\beta_{D,2} > \beta_{D,1}$ and $\beta_{D,4} > \beta_{D,3}$. The reverse predictions should all apply for the clean equation: $\beta_{C,2} < \beta_{C,1}$ and $\beta_{C,3} > \beta_{C,4}$.

3.2 Dynamic count data models with fixed effects

To estimate equation (10) and (11) we use:

$$PAT_{zit} = \exp(x_{it}\beta_z)\eta_{zi} + u_{zit} \quad (12)$$

where $z \in \{Dirty, Clean\}$ and x_{it} is the vector of regressors. We compare a number of econometric techniques to account for firm level fixed effects η_{zi} in these Poisson models. Our baseline is an econometric model we label CFX, the Control Function Fixed Effect estimator. It builds on the pre-sample mean scaling estimator proposed in Blundell, Griffith and Van Reenen (1999), henceforth BGVR.¹⁸ BGVR suggest conditioning on the pre-sample average of the dependent variable to proxy out the fixed effect. Like BGVR, CFX uses a control function approach to deal with the fixed effect but rather than using information from the pre-sample period in the control function, we simultaneously estimate the main regression equation and a second equation allowing us to identify the control function from *future* data (similar to the idea of taking orthogonal deviations in the linear panel data literature, see Arellano, 2003). The full details on this are provided in Appendix A, but in a nutshell, we use CFX to deal with a potential concern with the BGVR approach, namely that it requires long pre-sample history of realizations of the dependant variable. However, in our data - particularly for clean - patenting is

change as the absolute sign of the price effects will depend on the elasticity of substitution between cars and other goods.

¹⁸See also Blundell, Griffith and Windmeijer (2002) and Blundell, Griffith and Van Reenen (1995).

concentrated towards the end of our sample period. Below, we provide results using both the CFX and BGVR method as well as two other common approaches. First, we use the Hausman, Hall and Griliches (1984) method (the count data equivalent to the within groups estimator) even though this requires strict exogeneity, which is inconsistent with models including functions of the lagged dependent variable as we have in equations (10) and (11). Second, we implement some simple linear within groups models adding an arbitrary constant to the dependent variable before taking logarithms. We show that all these approaches deliver similar qualitative results, although CFX provides the best overall fit.

4 Data

4.1 Main dataset

Our main data are drawn from the World Patent Statistical Database (PATSTAT) maintained by the European Patent Office.¹⁹ Patent documents are categorized using the International Patent Classification (IPC) and national classification systems. We extract all the patents relating to “clean” and “dirty” technologies in the automotive industry.

Clean is identified by patents whose technology class is specifically related to electric, hybrid and hydrogen vehicles. Our selection of relevant IPC codes for clean technologies relies heavily on previous work by the OECD.²⁰ The precise description of the IPC codes used to identify relevant patents can be found in Table 1 Panel A. Some typical IPC classification codes included in the clean category are B60L11 (“Electric propulsion with power supplied within the vehicle”) and B60K6 (“Arrangement or mounting of hybrid propulsion systems comprising electric motors and internal combustion engines”). US patent #6456041²¹ is an example of a clean patent from our dataset: it describes a “Power supply system for electric vehicle”. It was first filed by Yamaha Motor in Japan in 1998 and was then filed at the European Patent Office and at the USPTO in 1999. The front page and technical diagram of the patent is shown in Appendix Figure A1.

¹⁹PATSTAT can be ordered from EPO at <http://www.epo.org/searching/subscription/raw/product-14-24.html>

²⁰See www.oecd.org/environment/innovation, Vollebergh (2010) and Hasčič et al (2008).

²¹We use the publication numbers in this and the following patent examples.

“Dirty” includes patents with an IPC that is related to the internal combustion engine. These can be found in various sub-categories of the F02 group, for example F02B (“Combustion engines in general”), F02F (“Cylinders, pistons or casings for combustion engines”) or F02N (“Starting of combustion engines”). The full list of IPC codes used to identify dirty patents is in Table 1 Panel B. Each of these groups includes several dozen sub-classes and an example of the full list of sub-classes for the F02F group is shown in Appendix Figure A2. The dirty category typically includes patents covering the various parts that make up an internal combustion engine. For example, EPO patent #0967381 protects a “Cylinder head of an internal-combustion engine” and USPTO patent #5844336 protects a “Starter for an internal combustion engine”.

As noted above, an important feature of the dirty category is that some patents included in this group aim at improving the fuel efficiency of internal combustion engines, making the dirty technology less dirty. We refer to these fuel-efficiency patents as “Grey” patents. In our baseline results, grey patents are included in the dirty category, but we also disaggregate the category to estimate models separately for grey and “Pure Dirty” innovations separately (as well as splitting up the knowledge stocks along these lines on the right hand side of the regressions). To select IPC codes for grey technologies, we use recent work at the European Patent Office related to the new climate change mitigation patent classification (see Veefkind et al. 2012). We complement this with information from interviews with engineers working in the automobile industry.²² The list of IPC codes is shown in Panel C of Table 1. An example of a grey patent is EPO patent #0979940, which protects a “Method and device for controlling fuel injection into an internal combustion engine” Electronic fuel injection technologies constantly monitor and control the amount of fuel burnt in the engine, with a view to reduce the amount of fuel unnecessarily burnt, thus optimizing fuel consumption. Appendix Figure A3 has the front page and technical diagram of this patent.

Alongside the grey fuel efficiency innovations there are many purely dirty patents, such as EPO patent #0402091, which covers a “Four-cycle twelve cylinder engine” (see

²²We are especially indebted to Christian Hue de la Colombe for many extremely helpful discussions.

Appendix Figure A4). Fuel consumption is proportional to the number and the volume of cylinders: the average car sold in Europe has four cylinders whereas it is six in the US. A twelve cylinder engine is much more powerful than a four or six cylinders engine, but this comes at the cost of increased fuel consumption. Twelve cylinder engines are used by many car makers for their top-range models, including Aston Martin, Audi, BMW and Rolls Royce. These cars typically run about 15 miles/gallon, while the average new car sold in the US in 2011 obtains 33.8 miles/gallon.²³

To measure innovation, we use a count of patents by application/filing date. The advantages and limitations of patenting as a measure of innovation have been extensively discussed.²⁴ For our purposes, there are three advantages of using patents. First, they are available at a highly technologically disaggregated level. We can distinguish innovations in the auto industry according to specific technologies whereas R&D investment cannot be easily disaggregated. Second, R&D is not reported for small and medium sized firms in Europe nor for privately listed firms in the US (they are exempt from the accounting requirement to report R&D). Third, the auto sector is innovation intensive where patents are perceived as an effective means of protection against imitation, something which is not true in all sectors (Cohen et al., 2000).²⁵ These considerations make patents a reasonably good indicator of innovative activity in the auto sector.

Patents do suffer from a number of limitations. 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 is that patent values are highly heterogeneous with most patents having a very low valuation. Finally, the number of patents that are granted for a given innovation varies significantly across patent

²³See <http://www.fueleconomy.gov> for details on car consumption and http://www.bts.gov/publications/national_transportation_statistics/html/table_04_23.html for US average. Note grey technical change appears to account for a minority of innovation in the internal combustion engines. For example, between 1980 and 2004 the fuel efficiency of passenger cars increased by only 6.5%, while horsepower increased by 80% (Knittel, 2011).

²⁴See Griliches (1990) and for a recent overview, OECD (2009). Dating by application is conventional in the empirical innovation literature as it is much more closely timed with when the R&D was performed than the grant date.

²⁵Cohen et al. (2000) conducted a survey questionnaire administered to 1,478 R&D labs in the U.S. manufacturing sector. They rank sectors according to how effective patents are considered as a means of protection against imitation, and find that the top three industries according to this criterion are medical equipment and drugs, special purpose machinery and automobile.

offices with concerns over increasing laxity in recent years particularly in the USPTO (e.g. Jaffe and Lerner, 2004).

To mitigate these problems, we focus on “triadic” patents²⁶ which are those patents that have been taken out in all three of the world’s major patents offices in the US, Europe and Japan (USPTO, EPO and JPO).²⁷ Focusing on triadic patents has a number of advantages. First, triadic patents provide us with a common measure of innovation worldwide, which is robust to administrative idiosyncracies of the various patent offices. For example, if the same invention is covered by one patent in the US and by two patents in Japan, all of which are part of the same triadic patent family, we will count it as one single invention. Secondly, triadic patents cover only the most valuable inventions which explains why they have been used so extensively to capture high-quality patents.²⁸ Thirdly, triadic patents typically protect inventions that have a potential worldwide application so these patents are thus relatively independent of the countries in which they are filed.

Our data set includes 6,419 “clean” and 18,652 “dirty” triadic patents.²⁹ Since the EPO was created in 1978 our triadic patent data only starts in that year. The last year of fully comprehensive triadic data is 2005, so this is our end year.³⁰ Our basic dataset consists of all those applicants (both firms and individuals) who applied for at least one of these clean or dirty auto patents. We identify 3,423 distinct patent holders, which breaks down into 2,427 companies and 996 individuals. For every patent holder we subsequently identify all the patents they filed. We also extract other pieces of information based on

²⁶To identify triadic patents we use the INPADOC dataset in PATSTAT. For details on the construction of patent families see Martinez, 2010.

²⁷Following standard practice we use all patents filed at the EPO and JPO and granted by the USPTO. This is because the USPTO only published patent applications that were subsequently granted until 2001 (when they changed policy in line with the other major patent offices). For consistency we thus consider only triadic patents granted by the USPTO both before and after 2001. For the official definition of triadic patents and how triadic patent families are constructed, see Dernis and Kahn, 2004, and Martinez, 2010.

²⁸It 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). Grupp et al. (1996); Grupp (1998); Dernis, Guellec and van Pottelsberghe (2001); Dernis and Khan (2004); Guellec and van Pottelsberghe, (2004)

²⁹In total, the PATSTAT data set includes 213,668 “clean” and 762,708 “dirty” patent applications across all 80 patents offices. Thus by using triadic patents we focus on the high end of the quality distribution.

³⁰The number of triadic patent in all technologies (i.e. including patents that are neither clean nor dirty) starts falling in 2006. This is because of time lags between application and grant date at the USPTO.

this sample which we use to construct weights for prices and spillovers. For example, we identify all the other patents filed by holders of at least one clean or dirty triadic patent, which represents a total of 4,467,362 patent applications.

4.2 Tax-inclusive fuel prices

To estimate the impact of a carbon tax on innovation in clean and dirty technologies, we use information on fuel prices (FP_{ct}) and fuel taxes. Data on tax-inclusive fuel prices are available from the International Energy Agency (IEA) for 25 major countries from 1978 onwards.³¹ We construct a time-varying country-level fuel price defined as the average of diesel and gasoline prices. The average fuel price across countries for our regression sample period 1986-2005 is shown in Panel A of Figure 1. Prices fell from the mid to late 1980s then rose peaking just before the Dot-Com bust of 2000-01. Prices then fell before recovering after 2003. Average fuel taxes have followed a broadly similar pattern falling in late 1980s, rising throughout the 1990s and falling back in the 2000s (Panel B of Figure 1). What is more striking, however, is the high variability across countries of changes in the fuel price over time, much of it being driven by cross-country differences in tax policies (see Figures 2 and 3).

Fuel prices are available only at the country-year level, whereas our dependent variable has firm level variation that we would like to exploit. A related issue is that the auto market is global and government policies abroad might be at least as important for firms' innovation decisions as domestic policies in the company's headquarter country. We allow fuel prices to have a different effect across firms by noting that some geographical markets matter more than others for reasons that are idiosyncratic to an auto firm. Firstly, auto manufacturers have different styles of vehicles reflecting their heterogeneous capabilities and branding that are differentially popular depending on local tastes (e.g. Berry et al, 1995; Goldberg, 1995; Verboven, 1999). Second, there is typically some home bias towards "national champions" auto manufacturers in national tastes and government policies. For example, the recent auto bailouts in Detroit were paid for by US taxpayers whereas the bailout of Peugeot has been shouldered by the French. The upshot of this is that auto

³¹The IEA reports some incomplete data for an additional 13 countries.

firms display heterogeneous current and expected market shares across nations and their R&D decisions will be more influenced by prices and policies in some countries than in others.

To operationalize this idea we construct a fuel price variable for each firm as a weighted average of fuel prices across countries based on a proxy of where the firm expects its future market to be. Our price index for firm i at time t is defined as:

$$\ln FP_{it} = \sum_c w_{ic0}^{FP} \ln FP_{ct} \quad (13)$$

where FP_{ct} is the tax-inclusive fuel price discussed above and w_{ic0}^{FP} is a firm-specific weight (this is time invariant and uses information only prior to the regression sample period). The weight is determined by the importance of county c as a market outlet for firm i , so we define w_{ic0}^{FP} as the fraction of firm i 's patents taken out in country c . The rationale for doing this is that a firm will seek intellectual property protection in jurisdictions where it believes it will need to sell in the future (even if it licenses the technology, the value of license will depend on whether it has obtained intellectual property protection in relevant growth markets). For every patent applied for, 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 EU 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 expectation of where its markets will be in the future, we can use this distribution to construct a firm-specific fuel price, FP_{it} , whose value is computed as the weighted mean of the $\ln(\text{fuel prices})$ in the relevant markets, with weights w_{ic0}^{FP} equal to the shares of the corresponding countries in the firm's patent portfolio. For example, if a firm had filed 30 patents, 20 in the US and 10 in Japan, the price changes in the US would get a weight of two-thirds and the Japanese price changes a weight of one third. In order to reflect the greater importance of larger countries, we also weight by each country's average GDP.

We calculate the weights using the patent portfolio of each company averaged over the 1965-1985 "pre-sample" period, whereas we run regressions over the period 1986-2005.

This is to make sure that the weights are weakly exogenous as patent location could be influenced by shocks to innovation. Choosing 1985 as the cut-off is to ensure there is enough time pre-sample to construct the weights. We perform robustness tests using different pre-sample periods to check that nothing is driven by the precise year of cut-off (e.g. use 1965-1990 as the pre-sample period and estimate the regressions from 1991 onwards).³²

Why not use an alternative weighting scheme which simply reflects where firms currently sell their products (e.g. as in Bloom, Schankerman and Van Reenen, 2012)? First, we believe that the information on where firms choose to take patent protection is a potentially better measure because it reflects their *expectations* of where their future markets will be. Second, there is a data constraint: although sales distributions by geographic area are available for larger firms they are not available for smaller firms - and there are many patents from these smaller firms. We show our weights compared to sales weights for some of the largest car firms in Appendix Table A1 - Toyota, VW, Ford, Honda and Peugeot. The correlation is generally high suggesting that the weights we choose do a reasonable job at reflecting market shares.³³

4.3 Firm lagged Patent stocks and spillovers

Firm patent stocks are calculated in a straightforward manner using the patent flows ($PAT_{z,it}$) described above. Following Cockburn and Griliches (1988) and Peri (2005), the patent stock is calculated using the perpetual inventory method:

$$K_{z,it} = PAT_{z,it} + (1 - \delta)K_{z,it-1} \quad (14)$$

where $z \in \{Dirty, Clean\}$. We take δ , the depreciation of R&D capital, to be 20%, as is often assumed in the literature, but we check the robustness of our results to other values of the depreciation parameter.

To construct aggregate spillovers for a firm, we use information on the geographical

³²Our approach to weighting is conservative and may cause us to underestimate the importance of fuel prices for innovation decisions. If firms reduce patent protection in countries where fuel taxes rise, this will generate a measurement error biasing the estimated coefficients towards zero.

³³One exception is that VW appears to have a much higher patent share in Germany (its home country) than its sales would suggest.

location of the various inventors in that firm. These are geographically located regardless of nationality of the firm's headquarters or the location of the office where the patent was filed (e.g. the patents of Toyota's scientists working in US research labs are part of this US spillover pool). Implicit in our approach is the view that the geographical location of an inventor is likely to be a key determinant of knowledge spillovers rather than the jurisdiction over which the patent is taken out (which matters more as a signal of where the market for sales is likely to be). Many papers have documented the importance of the geographical component of knowledge spillovers in patents and other indicators (e.g. Henderson, Jaffe and Trajtenberg, 1993, 2005 and Griffith, Lee and Van Reenen, 2011).

To construct a *firm*-specific spillover pool we use an analogous empirical strategy to that for the fuel price. The spillover weight w_{ic0}^S is the share of all firm i 's inventors (i.e. where the inventor worked) in country c between 1965 and 1985. This weight is distinct from w_{ic0}^P in equation (13) as it is based on the location of inventors who are more likely to benefit from research conducted locally. Importantly, the distribution of the patent portfolio across countries and the distribution of inventors vary considerably across firms. This is illustrated for the US in Appendix Figure A5.

The spillover for firm i is:

$$SPILL_{z,it} = \sum_c w_{ic0}^S SPILL_{z,ct} \quad (15)$$

where $SPILL_{z,ct}$ is the spillover pool in country c at time t . This is defined as:

$$SPILL_{z,ct} = \sum_{j \neq i} w_{jc0}^S K_{z,jt} \quad (16)$$

The spillover pool of a country is the sum of all other firms' patent stocks with a weight that depends on how many inventors the other firm has in that country.³⁴

As noted above, a common problem with patent data is that the value of patents is highly heterogeneous. We mitigate this problem by conditioning on triadic patents, which

³⁴An alternative approach would be to define the country level spillover as

$$SPILL_{z,ct} = \sum_j K_{z,jct} \quad (17)$$

where $K_{z,jct} = PAT_{z,jct} + (1 - \delta)K_{z,jct-1}$ and $PAT_{z,jct}$ is the number of patents filed by inventors of company j located in country c at year t . Empirically these two methods give very similar results.

screen out the very low value patents. But we also perform two other checks. First, we weight patents by the number of future citations. Second, we use “biadic” patents filed at the EPO and at the USPTO, following Henderson and Cockburn (1993) who argued that patents were important if they had been applied in at least two of the three major economic regions. Our results are robust to these two variants.

4.4 Descriptive statistics

Figure 4 shows that aggregate triadic clean and dirty patents have been rising over time. Dirty patents increased steadily between 1978 and 1988 and then again between 1992 and 2000, but have been decreasing during the last five years of our dataset. The number of clean patents was low for a decade until 1992, then began rising particularly after 1995 (at an average annual growth rate of 23%), peaking at 724 in 2002 alone, before falling back slightly. Consequently, while the number of clean patents represented only 10% of the number of dirty patents filed annually during the 1980s this reached 60% by 2005. Descriptive statistics for our dataset used in the regressions are shown in Table 2. In any given year, the average number of dirty patents is 0.22 and the average number of clean patents is 0.08.

Appendix B discusses more descriptives showing more of the cross-country distribution of patent filing (Table A2) and citation patterns (Table A3) which are consistent with spillovers being much stronger within the two categories (clean or dirty) than between them.

Table A4 displays the top 10 patenters in clean technologies between 1978 and 2005 and Table A5 shows the equivalent for dirty technologies. Japanese and German companies predominate although most top companies’ portfolios include both clean and dirty (the only exception is Samsung SDI, a battery specialist). Recall that this is based on triadic patents and US companies tend to file disproportionately more patents in the US than in Europe and Japan. Tables A6 to A9 report top clean and dirty patenters at the EPO and at the USPTO separately. General Motors is the third largest patenter of clean technologies at the USPTO whereas it is not even in the top 10 at the EPO. ³⁵

³⁵While it is clear that there a number of big companies active in both clean and dirty automotive

5 Results

5.1 Main Results

Our main results are shown in Table 3. The first three columns use the number of clean patents (a flow) in a firm as the dependent variable and the last three columns uses the flow of dirty patents. All estimates include firm fixed effects using the Control Function Fixed Effect (CFX) approach (described in Section 3), year dummies and GDP per capita. Column (1) shows that the coefficient on the (tax inclusive) fuel price is positive and significant. The elasticity of 0.97 implies that a 10% higher fuel price is associated with about 10% more clean patents. The coefficients on spillovers and lagged patent stocks take signs consistent with the path dependency hypothesis. Firms who are more exposed to larger stocks of clean innovation by other firms' ("clean spillovers", $SPILL_{C,it-1}$) are significantly more likely to produce clean patents, whereas those benefiting more from dirty spillovers ($SPILL_{D,it-1}$) are significantly less likely to innovate in clean technologies. If the coefficients can be interpreted causally, increasing the clean spillover stock by 10% increases a firm's clean innovation by 2.7%. By contrast, an increase in the exposure to dirty spillovers by 10% reduces clean innovation by 1.7%.

In addition to path dependency at the economy level through spillovers, there is also path dependency at the firm level. Column (1) of Table 3 suggests that firms which have innovated in clean innovation in the past ($K_{C,it-1}$) are much more likely to continue to innovate in clean technologies in the future, with a significant elasticity of 0.306. Interestingly, a firm's own history of dirty innovation ($K_{D,it-1}$) is also associated with more clean innovation with an elasticity of 0.139. This coefficient is, however, much smaller than the corresponding coefficient on past dirty innovation stocks in the dirty innovation equation (column (4)) which is four times as large (0.557). In other words, firms with a history of dirty innovation are more likely to innovate in the future in either clean or dirty (compared to those with little innovation), but this effect is much stronger for dirty

patenting, computing a Herfindahl Index for patenting over 1978 to 2005 for clean innovation we find a Herfindahl of 0.023 and for dirty we find a HHI of 0.038, implying low concentration. The top 10 patent holders in clean account for 35.6% of patents over 1978 to 2005 whereas the corresponding figure is 46.6% for dirty.

innovations than for clean innovation leading to path dependence. Moreover, note that in column (1) the coefficient on a firm's past dirty innovation stock on future clean innovation (0.139) is much smaller than the effect of past clean innovations on future clean innovation (0.306).³⁶

Columns (2) and (3) of Table 3 include a measure of R&D subsidies for clean technologies and a control for emission regulations. R&D subsidies are from the IEA's Energy Technology Research Database and the emissions regulations index are from Dechezlepretre, Perkins and Neumayer (2010) with details in Appendix B. In contrast to the proxy for carbon taxes (fuel prices) neither of these additional policy variables is statistically significant and the coefficients on the other variables do not change much. The absence of an R&D subsidy effect is surprising, but we suggest why in Table 4.

Columns (4)-(6) of Table 3 repeat the specification in the first three columns but use dirty patents as the dependent variable instead of clean. The coefficient on fuel prices is negative and significant in all columns. In column (4) a 10% increase in fuel prices is associated with about a 6% decrease in dirty innovation. The estimates on spillovers and knowledge stocks are symmetric to those in the clean equation. Exposure to dirty spillovers fosters future dirty innovation, whereas clean spillovers reduces dirty patenting. The coefficients suggest that a firm's own history of either dirty or clean patenting has a positive effect on further dirty patenting, but the effect of past dirty patenting is stronger on future dirty innovation than past clean.

In summary, Table 3 offers considerable support for our model. First, higher fuel prices significantly encourage clean innovation and significantly discourage dirty innovation. Second there is path dependency in the direction of technical change: countries and firms that have a history of relatively more clean (dirty) innovation are more likely to innovate in clean (dirty) technologies in the future.

³⁶This effect is not predicted by the theory but could arise from cross-technology knowledge spillovers (the cost coefficient ψ could depend on the amount of accumulated knowledge in clean and in dirty).

5.2 Grey Innovations

Our dirty category includes innovations relating to improvements in the energy efficiency of internal combustion engines. We labeled these “grey” innovations and consider disaggregating the dirty category into these grey and purely dirty innovations. As noted in the theory Section the effect of fuel prices are more ambiguous in this middle grey category. On the one hand, there are incentives to substitute research away from purely dirty into grey innovation when the fuel price rises. On the other hand, there is also an incentive to switch away from the internal combustion engines completely (including grey) towards alternative clean vehicles.

Table 4 presents the results and shows that, as expected, the coefficient on the fuel price for grey innovation in the second column (0.282) lies between the coefficients on clean (positive at 0.848 in column (1)) and purely dirty (very negative at -0.832 in column (3)). This is consistent with fuel prices having a positive effect on energy efficient innovation, although smaller and insignificant when compared to the effect of fuel prices on purely clean innovations. Another interesting feature of the results is that the coefficient on R&D subsidies is positive and significant in the grey innovation equation whereas it continues to be insignificant in the clean and purely dirty equations. This is consistent with the fact that the majority of these government subsidies are for energy efficiency (see Appendix B) rather than more radical clean technologies.

Since we have also disaggregated the spillover stocks and the firm’s own past innovation stocks into the three categories now we have six variables reflecting path dependency on the right hand side of the regression. The coefficients on these variables take a broadly sensible pattern, but precision has fallen as there is likely to be some collinearity issues with a large number of highly correlated variables.

Given how demanding this specification is we find the overall results from Table 4 encouraging and consistent with the theory.

5.3 Extensions and robustness

Table 5 considers the alternative econometric approaches for dynamic count data models with firm fixed effects discussed in Section 3. First, we follow Hausman, Hall and Griliches (“HHG”) in column (1) for clean patents and column (3) for dirty patents. The signs of coefficients are generally the same as in our baseline model of Table 3, but the marginal effect of fuel price is much greater in absolute magnitude for dirty innovation and smaller (and insignificant) for clean. Indeed the magnitude of the estimated elasticity for dirty patents seems unreasonably large (-2.457). We suspect that the assumption of strict exogeneity underlying HHG is problematic in our context, as we have a highly dynamic specification. Columns (2) and (4) implement the Blundell et al (1999, BGVR) estimator. The pattern of the spillover effects and dynamics remain similar to the baseline regression, and we still obtain a positive and significant effect of fuel prices on clean innovation and a negative and significant effect on dirty innovation. The fuel price coefficients are comparable to the baseline case.³⁷

The final two columns of Table 5 uses relative patenting $\ln(1 + PAT_{Clean,it}) - \ln(1 + PAT_{Dirty,it})$ as the dependent variable in an OLS regression with firm dummies (i.e. the linear within groups estimator). Column (5) shows that there is a significant and positive effect of fuel prices on relative innovation. Column (6) shows that this result is robust to including a full set of country by year fixed effects to absorb any potential country specific time varying policy variables.³⁸

Could the results somehow be driven by firms who were not patenting prior to 1986? Table 6 repeats the baseline regressions for our three count data models (BGVR, HHG and CFX) restricting the sample to firms with at least one patent before 1986. This leads to only small changes in the coefficients and no change in the overall qualitative patterns.

Oil prices are broadly global, so most of the country-specific variation over time in

³⁷However, notice that we find larger values for the effects of clean knowledge stocks on clean patenting and dirty knowledge stocks on dirty patenting than in both the baseline CFX and the HHG specification. This could mean that the BGVR approach is not fully controlling for all the fixed effects by relying on pre-sample patenting only.

³⁸The country here is based on the headquarters whereas the previous country variables like fuel price with weighted averages using patent weights. Note that it is computationally infeasible to include the full set of country by time dummies in the non-linear count data models.

fuel prices comes from differential taxation. Consequently, Table 7 substitutes fuel taxes for fuel prices showing again a similar pattern of results. One difference is that the point estimates of the fuel price response are smaller in absolute terms for both types of innovation. This is to be expected as demand is driven by the final price the consumer pays rather than the fuel tax itself.

We experimented with changing the cut-off year to 1990 and ran the regressions 1991-2005 using data from 1965-1990 to construct the weights. The results in Table 8 are quite comparable to our baseline, although standard errors are a little larger as we would expect from using a smaller sample for the regressions.

Table 9 reports alternative dynamic specifications for fuel prices. The first five columns are for clean and use fuel prices dated in the current year in column (1), lagged one year in our (baseline) of column (2), lagged two years in column (3) and lagged three years in column (4). In column (5) we construct a geometrically weighted average of past fuel price levels as proposed by Popp(2002).³⁹ We repeat these specifications in the last five columns but use dirty patents instead. With all these approaches we find price coefficients that are very similar to our earlier estimates with a positive elasticity of clean patents with respect to fuel price of around unity and a negative elasticity of dirty patents of around -0.6.⁴⁰

We conducted other robustness tests reported in Appendix A. First, our outcome variable is Triadic patents, those filed at all three main patent offices in the world (USPTO, EPO and JPO). A concern is that this screens out too many of the lower value patents. To address this we ran our regressions based on biadic rather than triadic patents; i.e. we included all patents into the construction of the innovation and knowledge stock variables that are filed at the EPO and the USPTO but not necessarily the JPO. Table A10 shows that the results are robust to this experiment. Second, we constructed the patent

³⁹Popp (2002) uses an adaptive expectations model of prices, in which the expected future price of energy is a weighted average of past prices: $P_{it}^* = \sum_{k=0}^n \lambda^k P_{i,t-k}$. The parameter λ captures the speed at which agents adjust their expectations based on the gap between the predicted and the realized values. For comparison purposes we use the same adjustment factor of $\lambda = 0.83$ as in Popp (2002).

⁴⁰We tried to pin down more precisely the dynamic response structure by including multiple lags of price simultaneously but autocorrelation in prices made it difficult as all coefficients tended to be zero, as in Popp (2002).

stock variables - including the spillover variables - using citation weighted counts from all worldwide patents (Table A11). This led to qualitatively similar results, e.g. the fuel price response is larger for clean patents than for dirty patents.⁴¹ Third, we experimented with a wide range of other country specific variables and report that the results are robust to these additional covariates.⁴² For example, in Table A12 we included total GDP in addition to GDP per capita. The coefficient on GDP is insignificant and the basic pattern of our results is robust to this extra control (see Table A12).

6 Simulation results

To obtain a better sense of the aggregate magnitude of the results we report a number of counterfactual experiments. We explore the implications of our econometric models for the evolution of future clean and dirty knowledge stocks and how this is affected by an increase in the fuel price (generated, for example, by an international carbon tax). We recursively compute values of expected patenting under different policy scenarios, use those to update the knowledge stock variables (including the spillover variables) and feed these into the next iteration. Hence if we split the right hand side variables x_{it} into variables that are functions of the lagged knowledge stock (k_{it}) and other variables such as the fuel price (p_{it}), we can write $x_{it} = [k_{it}, p_{it}]$ and a particular iteration in period T greater than t as defined by:

$$\begin{aligned} \widehat{PAT}_{z,it+T} &= \exp(k_{it+T-1}\beta_{kz} + p_{i,t+T}^{CF}\beta_{pz}) \eta_{z,i} \\ k_{i,t+T} &= f(k_{i,t+T-1}, \widehat{PAT}_{Clean,it+T}, \widehat{PAT}_{Dirty,it+T}) \end{aligned} \quad (18)$$

where $\widehat{PAT}_{Clean,it+T}$ and $\widehat{PAT}_{Dirty,it+T}$ are vectors of predicted patent flows for firms in the sample and p_{it+T}^{CF} are potentially counterfactual values of the policy and other control variables. Our results imply that there is path dependence in the type of innovation pursued, both through internal knowledge stock effects as well as external spillovers. Here we explore how important this path dependence is in quantitative terms by studying the evolution of both clean and dirty knowledge stock implied by our fitted models into

⁴¹If anything, the results are generally stronger with elasticities that are larger in magnitude.

⁴²Other country-level variables matched in across all the locations a firm expects to sell in like population did not appear to be robustly significant in our empirical models.

the future. We do this for every firm in the dataset and then aggregate across the world economy in each period.

More specifically, we are looking for conditions under which *the clean knowledge stock for the aggregate economy exceeds the dirty knowledge stock*. In line with Acemoglu et al. (2012a) this would be a requirement for clean technologies to be able to compete with dirty ones, even without policy intervention. Our projections should be considered as a rough exploration into the importance of carbon taxes and path dependency rather than precise forecasts of future innovation.

We focus on the period up to 2030 with 2020 as a focal point. This is somewhat arbitrary but in line with scenarios of the International Energy Agency (IEA)⁴³ suggesting that globally fossil fuel use must peak by 2020 to avoid highly risky climate change. It is also consistent with the European Commission’s 2020 targets.⁴⁴

We first check the within sample performance of the model by implementing simulation runs providing recursively generated knowledge stocks over the regression sample period (1986-2005) in Appendix Figure A6.⁴⁵ Clean and dirty patent stocks are reported on the y-axis. Comparing predicted aggregate patents to the actual values suggests that our preferred CFX model does a reasonably good job at tracking the aggregate changes in clean and dirty patenting (Panel A). The alternative BGVR and HHG estimates are not terrible but do less well in later years (Panels B and C).

Figure 5 reports simulations based on the regressions from Table 6 columns (1) and (4) for years through to 2030. In Panel A we report the baseline case keeping fuel prices (and time dummies) at their 2005 values.⁴⁶ The regressions imply a strong enough path dependency for the gap between dirty and clean knowledge stocks to remain far apart for a considerable period of time. Clean innovation catches up with dirty only well after 2030.⁴⁷ To what extent can carbon taxes speed up this convergence process? We examine

⁴³<http://blogs.ft.com/energy-source/2009/11/10/fossil-fuel-use-must-peak-by-2020-warns-iea/#axzz1tQmZyLoy>

⁴⁴See <http://ec.europa.eu/news/economy/100303.en.htm>

⁴⁵For the simulations we restrict the sample to the firms where we have pre sample information. In this way we do not have to make further assumptions as to how changes in the spillover and policy variables would affect firms where these variables are essentially missing.

⁴⁶We assume per capita GDP grows at 1.5% p.a., but report alternative assumptions in Figure 6.

⁴⁷Catch up occurs because of delayed reaction to fuel price hikes leading up to 2005 and GDP per

the effects of a permanent worldwide increase in fuel prices in 2006 (and fixed at this level thereafter) of 10%, 20%, 30%, 40% and 50% in Panels B through F respectively. In Panel B we see that the gap between clean and dirty becomes smaller with a fuel price increase of 10% both because there is more clean innovation and because there is less dirty innovation. However, parity is achieved between clean and dirty only after 2030. It would take an increase of 40% in fuel prices in order to achieve parity in 2020 according to our model (Panel E). This is a pretty large increase - comparable with the increase that took place in the 1990s in Figure 1.

One criticism of the simulation is that we would expect such a large increase in the fuel prices to have a negative effect on GDP per capita due to deadweight costs of taxation, adjustment costs and so on. This in turn could slow down the growth of clean innovation (e.g. Gans, 2012). To obtain some insight into the magnitude of these effects, Figure 6 considers the 40% fuel tax hike scenario coupled with a negative effect on GDP per capita growth. Panel A reproduces the baseline case where there is no effect on GDP (as in Figure 5 Panel E). Panel B considers a fall in the growth rate by 0.25 percentage point (e.g. from 1.5% to 1.25%). This postpones the crossover year because income growth has a stronger positive effect on clean innovation than dirty innovation in our estimates. But the effect is rather small, moving the crossover year from 2020 to 2022, only two years. Larger tax-driven falls in GDP per capita growth postpone things further, but it would take a full one percentage point a year fall in the growth rate to postpone the crossover year beyond 2030. We view it as unlikely that fuel taxes would knock a percentage point off annual growth for 15 years or more and this also ignores the damaging effects of global warming itself on economic growth over the medium run. We therefore take some comfort from Figure 6 that incorporating output effects would not dramatically change the conclusions from Figure 5.

In Figure 7 we explore the importance of path dependence for the simulations. First we repeat the baseline specifications allowing for all dynamic adjustments in the cases of no fuel price change (Panel A) and of a 40% increase (Panel B). In panels C and D

capita growth which relatively favors clean innovation.

we repeat this exercise while fixing all innovation stock variables - i.e. both spillovers and own knowledge stocks - at their 2005 levels. As a consequence both clean and dirty innovation and thus the growth rate of knowledge stocks reduces markedly as firms no longer benefit from standing on the shoulders of either their own or others' past innovation success. Also note that in Panel C where we keep prices fixed the gap between clean and dirty is now much narrower than in the equivalent Panel A. Despite this, the 40% increase in fuel prices in Panel D is much less effective than in Panel B where the dynamic effects from knowledge stocks are switched on. This illustrates that path dependency is a double edged sword as pointed out by Acemoglu et al (2012a). In the absence of effective policies it creates a kind of lock-in for dirty innovation. But if effective policies are introduced like a carbon tax or R&D subsidy, path dependency can help reinforce the growth of clean innovation as the economy accumulates clean knowledge more rapidly. Hence, if we switch off the two path dependency channel, innovation trends become less responsive to tax policy.

7 Conclusion

In this paper we have combined several patent datasets to analyze directed technical change in the auto sector, which is a key industry of concern for climate change. We use patenting data from 3,423 firms and individuals between 1965 and 2005 across 80 patent offices. Consistent with what theory predicts we find that “clean” innovation is stimulated by increases in the tax-inclusive fuel prices (our proxy for a carbon tax) whereas dirty innovation is depressed. We exploit the fact that prices evolve differentially across time in countries in our dataset and that firms are differentially exposed to these price changes because of their heterogeneous expected sales in geographic markets.

Our second key result is that there is strong evidence for “path dependency” in the sense that firms more exposed to clean innovation are more likely to direct their research energies to clean innovation in the future (a directed knowledge spillover effect). Similarly, firms with a history of dirty innovation in the past are more likely to focus on dirty innovation in the future. The fact that such path dependency holds for clean (as well as

dirty) innovation highlights the desirability of acting sooner to shift incentives for climate change innovation. Since the stock of dirty innovation is greater than clean, the path dependency effect will tend to lock economies into high carbon emissions, even after the introduction of a mild carbon tax or R&D subsidies for clean. So this may make the case for stronger action now, which could be relaxed in the future as the economy's stock of knowledge shifts in more of a clean direction. Increases to carbon prices can bring about a change in direction. For example, our baseline results suggest an increase of 40% of fuel prices with respect to the 2005 fuel price will allow clean innovation stocks to overtake dirty stocks after fifteen years.

Our analysis could be extended in several ways. First, we could extend the analysis of output effects beyond the macro adjustments in the simulations of Table 6 to examine the firm-level effects. This would require a large extension in terms of using data on sales, however. Second, we could use our framework to simulate other policies, such as country specific changes in carbon taxes (or R&D subsidies) to see how this would affect the innovation profiles in specific countries rather than just globally. Third, the same basic approach could be taken to look at other sectors than automobiles such as the energy sector as in Acemoglu et al (2012b). Finally, we could 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. As argued in Acemoglu et al (2012a), these parameters play as important a role as the discount rate in characterizing the optimal environmental policy.

References

- Acemoglu, D. (1998). Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. *Quarterly Journal of Economics*, 113:1055–1089.
- Acemoglu, D. (2002). Directed Technical Change. *Review of Economic Studies*, 69(4):781–810.
- Acemoglu, D. (2007). Equilibrium Bias of Technology. *Econometrica*, 75(5):1371–1410.
- Acemoglu, D., Aghion, P., Bursztyn, L., and Hémous, D. (2012a). The Environment and Directed Technical Change. *The American Economic Review*, 102 (1):131–166.
- Acemoglu, D., Akcigit, U., Hanley, D., and Kerr, W. (2012b). The Transition to Clean Technology.

- Acemoglu, D. and Finkelstein, A. (2008). Input and Technology Choices in Regulated Industries: Evidence from the Health Care Sector. *Journal of Political Economy*, 116(5):837–880.
- Acemoglu, D. and Linn, J. (2004). Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry. *Quarterly Journal of Economics*, 119(3):1049–1090.
- Alcott, H. and Wozny, N. (2011). Gasoline Prices, Fuel Economy, and the Energy Paradox. *mimeo NYU*.
- Anderson, S. T., Kellogg, R., and Sallee, J. M. (2011a). What Do Consumers Believe About Future Gasoline Prices? NBER Working Papers 16974, National Bureau of Economic Research, Inc.
- Anderson, S. T., Kellogg, R., Sallee, J. M., and Curtin, R. T. (2011b). Forecasting Gasoline Prices Using Consumer Surveys. *American Economic Review*, 101(3):110–14.
- Arellano, M. (2003). *Panel Data Econometrics*. Number 9780199245291 in OUP Catalogue. Oxford University Press.
- Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–90.
- Bloom, N., Schankerman, M., and Reenen, J. V. (2012). Identifying Technology Spillovers and Product Market Rivalry. Technical Report dp0675, Centre for Economic Performance, LSE.
- 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.
- Blundell, R., Griffith, R., and Windmeijer, F. (2002). Individual effects and dynamics in count data models. *Journal of Econometrics*, 108(1):113–131.
- Buonanno, P., Carraro, C., and Galeotti, M. (2003). Endogenous Induced Technical Change and the Costs of Kyoto. *Resource and Energy Economics*, 25(1):11–34.
- Busse, M. R., Knittel, C. R., and Zettelmeyer, F. (2011). Pain at the Pump: The Effect of Gasoline Prices on New and Used Automobile Markets. *mimeo Northwestern University*.
- 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.
- Cohen, W. M., Nelson, R. R., and Walsh, J. P. (2000). Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not). Working Paper 7552, National Bureau of Economic Research.
- Crabb, J. M. and Johnson, D. K. N. (2010). Fueling Innovation: The Impact of Oil Prices and CAFE Standards on Energy-Efficient Automotive Technology. *The Energy Journal*, Volume 31(Number 1):199–216.

- Dasgupta, P. (2008). Discounting climate change. *Journal of Risk and Uncertainty*, 37(2).
- Dechezleprêtre, A., Perkins, R., and Neumayer, E. (2012). Regulatory Distance and the Transfer of New Environmentally Sound Technologies: Evidence from the Automobile Sector. Working Papers 2012.33, Fondazione Eni Enrico Mattei.
- 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.
- Gans, J. (2012). Innovation and Climate Change Policy. *American Economic Journal: Policy (forthcoming)*.
- Gerlagh, R. (2008). A climate-change policy induced shift from innovations in carbon-energy production to carbon-energy savings. *Energy Economics*, 30(2):425–448.
- Gerlagh, R., Kverndokk, S., and Rosendahl, K. E. (2009). Optimal Timing of Climate Change Policy: Interaction Between Carbon Taxes and Innovation Externalities. *Environmental and Resource Economics*, 43(3):369–390.
- Goldberg, P. K. (1995). Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry. *Econometrica*, 63(4):891–951.
- Golosov, M., Hassler, J., Krusell, P., and Tsyvinski, A. (2011). Optimal Taxes on Fossil Fuel in General Equilibrium. Technical Report 17348, National Bureau of Economic Research.
- Goulder, L. H. and Schneider, S. H. (1999). Induced Technological Change and the Attractiveness of C\$O_2\$ Abatement Policies. *Resource and Energy Economics*, 21(3-4):211–253.
- Griffith, R., Lee, S., and Van Reenen, J. (2011). Is distance dying at last? Falling home bias in fixedeffects models of patent citations. *Quantitative Economics*, 2(2):211–249.
- Griliches, Z. (1990). Patent Statistics as Economic Indicators: A Survey. *Journal of Economic Literature*, 28(4):1661–1707.
- Grübler, A. and Messner, S. (1998). Technological Change and the Timing of Mitigation Measures. *Energy Economics*, 20(5-6):495–512.
- Grupp, H. (1998). *Foundations of the Economics of Innovation: Theory, Measurement and Practice*. Elgar, Cheltenham.
- Grupp, H., Mnt, G., and Schmoch, U. (1996). Assessing Different Types of Patent Data for Describing High-Technology Export Performance.
- Guellec, D. and Van Pottelsberghe de la Potterie, B. (2004). From R&D to Productivity Growth: Do the Institutional Settings and the Source of Funds of R&D Matter? *Oxford Bulletin of Economics and Statistics*, 66(3):353–378.
- Habakkuk, H. J. (1962). *American and British technology in the nineteenth century: The search for labour-saving inventions*. University Press.

- Hanlon, W. (2011). Necessity is the mother of invention: Input supplies and Directed Technical Change. http://www.columbia.edu/~wwh2104/hanlon_dtc.pdf.
- 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.
- Haščič, I., De Vries, F., Johnstone, N., and Medhi, N. (2008). Effects of Environmental Policy on the Type of Innovation: The Case of Automotive Emissions Control Technologies. *OECD Journal: Economic Studies*, 2009(1).
- Hassler, J., Krussell, J., and Olovsson, C. (2011). Energy saving technical Change. IIES, Stockholm University.
- 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.
- Henderson, R., Jaffe, A., and Trajtenberg, M. (2005). Patent Citations and the Geography of Knowledge Spillovers: A Reassessment: Comment. *The American Economic Review*, 95(1).
- Henderson, R. and Newell, R. G., editors (2011). *Accelerating Energy Innovation: Insights from Multiple Sectors*. University of Chicago Press.
- Hicks, J. R. (1932). The Theory of Wages.
- Jaffe, A. B. and Lerner, J. (2004). *Innovation and its discontents: How our broken patent system is endangering innovation and progress, and what to do about it*. Princeton University Press.
- Jaffe, A. B. and Trajtenberg, M. (2002). Patents, citations, and innovations a window on the knowledge economy.
- Jaffe, A. B., Trajtenberg, M., and Henderson, R. (1993). Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *The Quarterly Journal of Economics*, 108(3):577–98.
- Knittel, C. R. (2011). Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector. *American Economic Review*, 101(7):3368–99.
- Knittel, C. R. (2012). Reducing Petroleum Consumption from Transportation. *The Journal of Economic Perspectives*, pages 93–118.
- Lanjouw, J. O., Pakes, A., and Putnam, J. (1998). How to Count Patents and Value Intellectual Property: The Uses of Patent Renewal and Application Data. *The Journal of Industrial Economics*, 46(4).
- Manne, A. and Richels, R. (2004). The impact of learning-by-doing on the timing and costs of CO₂ abatement. *Energy Economics*, 26(4):603–619.
- Martinez, C. (2010). Insight into Different Types of Patent Families. OECD Science, Technology and Industry Working Papers 2010/2, OECD Publishing.

- Mendelsohn, R., Sterner, T., Persson, U. M., and Weyant, J. P. (2008). Comments on Simon Dietz and Nicholas Stern's Why Economic Analysis Supports Strong Action on Climate Change: A Response to the Stern Review's Critics. *Review of Environmental Economics and Policy*, 2(2):309–313.
- Messner, S. (1997). Endogenized technological learning in an energy systems model. *Journal of Evolutionary Economics*, 7(3):291–313.
- Newell, R. G., Jaffe, A. B., and Stavins, R. N. (1999). The Induced Innovation Hypothesis and Energy-Saving Technological Change. *The Quarterly Journal of Economics*, 114(3):941–975.
- Nordhaus, W. D. (1994). *Managing the Global Commons: The Economics of Climate Change*. MIT Press, Cambridge, MA.
- Nordhaus, W. D. (2002). Modeling Induced Innovation in Climate Change Policy. In *Technological Change and the Environment*. Resources for the Future Press, Washington, DC.
- Nordhaus, W. D. (2007). A Review of the Stern Review on the Economics of Climate Change. *Journal of Economic Literature*, 45(3):686–702.
- OECD (2009). OECD patent statistics manual. Technical report, OECD.
- OECD (2011). Patent Search Strategies for the identification of selected environment related technologies.
- Peri, G. (2005). Determinants of Knowledge Flows and their Effects on Innovation. *Review of Economics and Statistics*, 87 (2):308–322.
- Perkins, R. and Neumayer, E. (2010). Does the 'California Effect' Operate Across Borders? Trading- and Investing-Up in Automobile Emission Standards. *Social Science Research Network Working Paper Series*.
- Popp, D. (2002). Induced Innovation and Energy Prices. *The American Economic Review*, 92(1):160–180.
- Popp, D., Newell, R. G., and Jaffe, A. B. (2009). Energy, the Environment, and Technological Change. NBER Working Papers 14832, National Bureau of Economic Research, Inc.
- Small, K. A. and Van Dender, K. (2007). Long Run Trends in Transport Demand, Fuel Price Elasticities and Implications of the Oil Outlook for Transport Policy. Technical Report 2007/16, OECD Publishing.
- 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 der Zwaan, B. C. C., Gerlagh, R., Klaassen, G., and Schrattenholzer, L. (2002). Endogenous Technological Change in Climate Change Modelling. *Energy Economics*, 24(1):1–19.

- Veefkind, V., Hurtado-Albir, J., Angelucci, S., Karachalios, K., and Thumm, N. (2012). A new EPO classification scheme for climate change mitigation technologies. *World Patent Information*, 34(2):106–111.
- Verboven, F. (1999). The Markets for Gasoline and Diesel Cars in Europe. *CEPR discussion paper*, 2069.
- Vollebergh, H. (2010). *Fuel Taxes, Motor Vehicle Emission Standards and Patents Related to the Fuel Efficiency and Emissions of Motor Vehicles*.
- von Below, D. and Persson, T. (2008). Uncertainty, Climate Change and the Global Economy. Working Paper 14426, National Bureau of Economic Research.
- Weitzman, M. L. (1974). Prices vs. Quantities. *The Review of Economic Studies*, 41(4):477–491.
- Weitzman, M. L. (2007). A Review of the Stern Review on the Economics of Climate Change. *Journal of Economic Literature*, 45(3):703–724.
- Weitzman, M. L. (2009). On Modelling and Interpreting the Economics of Catastrophic Climate Change. *Review of Economics and Statistics*, 91(1):1–19.
- West, S. E. (2004). Distributional effects of alternative vehicle pollution control policies. *Journal of Public Economics*, 88(3& 4):735–757.
- Yohe, G. W., Tol, R. S. J., and Anthoff, D. (2009). Discounting for Climate Change. *Economics - The Open-Access, Open-Assessment E-Journal*, 3(24):1–22.

Web Appendix (not intended for publication unless requested)

A Econometric Models

We separately examine clean and dirty patent counts using a standard Poisson model

$$PAT_{zit} = \exp(x_{it}\beta_z) \eta_{zi} + u_{zit} \quad (19)$$

where $z \in \{Dirty, Clean\}$ and x_{it} is a vector of regressors including functions of the lagged dependent variable. For identification we assume $E(u_{zit}|x_{it}) = 0$.⁴⁸ We consider four alternative estimation techniques that allow for the possibility of firm level fixed effects η_{zit} in the propensity to patent. The standard approach is Hausman, Hall and Griliches (1984, HHG) who suggest a transformation akin to the within groups estimator in the linear panel data context. In GMM terms, their estimator can be expressed as relying on the following moment condition for identification (e.g. Blundell, Griffith, Windmeijer, 2002):

$$E \left\{ \left(PAT_{zit} - \mu_{zit} \frac{P\bar{A}T_{zi}}{\bar{\mu}_{zi}} \right) x_{kit} \right\} = 0$$

for all variable in x_{it} where $\mu_{zit} = \exp(x_{it}\beta_z)$ and a bar represents the average of a variable over time for a specific firm. Note that

$$PAT_{zit} - \mu_{zit} \frac{P\bar{A}T_{zi}}{\bar{\mu}_{zi}} = u_{zit} - \frac{\mu_{zit}}{\bar{\mu}_{zi}} \bar{u}_{zit}$$

implying that we require strict exogeneity, i.e. the shock u_{zit} must be uncorrelated with x_{it} not only contemporaneously, but in all periods; i.e. $E\{u_{zit}|x_{i\tau}\} = 0$ for all t and τ . When using regressors that depend on past realizations of the dependent variable such as the knowledge capital stocks, this assumption is violated.

Blundell, Griffith and Van Reenen (1999, BGVR) proposed an alternative estimator which is robust to relaxing the strict exogeneity assumption. It relies on introducing a control function term for the fixed effects, which is identified from realizations of the dependent variable in a pre-sample period. Hence, the idea is to think of the fixed effect as the combination of a control term $\phi(\cdot)$ and an error, ω_i .

$$\eta_{zi} = \phi(\ln P\bar{A}T_{zi0}, I\{P\bar{A}T_{zi0} = 0\}) + \omega_i$$

where $P\bar{A}T_{zi0}$ is the average amount of patenting by firm i in the pre-sample period. BGVR show that with $\phi(\cdot) = \exp(\phi_{z1} \ln P\bar{A}T_{zi0} + \phi_{z2} I\{P\bar{A}T_{zi0} = 0\})$, pre-determined x_{it} ⁴⁹ and stationarity in the dynamic system implied by equation (19) estimates of β_z are

⁴⁸Note that we can equivalently represent the model in terms of a multiplicative shock ν_{zit} with $E(\nu_{zit}|x_{it}) = 1$. We would have

$$\nu_{zit} = 1 + \frac{u_{zit}}{\exp(x_{it}\beta_z) \eta_{zi}}$$

and our assumptions concerning u_{zit} imply $E(\nu_{zit}|x_{it}) = 1$.

⁴⁹i.e. $E\{u_{i\tau}|x_{it}\} = 0$ for $\tau \geq t$.

unbiased as the duration of the pre-sample period becomes large. Thus, effectively we estimate the following model:

$$PAT_{zit} = \exp(x_{it}\beta + \phi_{z1} \ln \bar{PAT}_{zi0} + \phi_{z2} I\{ \bar{PAT}_{zi0} = 0 \}) + u_{zit}$$

The BGVR approach requires the realizations of the dependent variable in the pre-sample period to be representative of a firm's behavior over the sample period. Formally, the series must be mean stationary (conditionally on the time dummies). It is easy to see why this might be violated in particular for clean patents, whose realizations are concentrated towards the end of our sample period. Consequently, for many firms we do not observe any clean patenting in the pre-sample period which could inform us about variations in their fixed propensity to patent in clean.

To address this problem we propose a new estimator in the same spirit of using a control function as in BGVR. However, rather than using information from the pre-sample period to calibrate the control function, we simultaneously exploit future data. We estimate the main regression equation as well as a second equation allowing us to identify the control function from future data. The key idea is the following. In general, a control term $\check{\phi}_{zit}(\cdot)$ will lead to consistent estimates, if the resulting error term $\check{\omega}_{zit} = \eta_{zit} - \check{\phi}_{zit}(\cdot)$ is orthogonal to x_{it} ; i.e. $E\{\check{\omega}_{zit} | x_{it}\} = 0$. Note, that given a parameter vector β we can obtain such an estimate by regressing⁵⁰

$$\frac{PAT_{ziT}}{\mu_{ziT}} = \eta_{zi} + \frac{u_{ziT}}{\mu_{ziT}} = \check{\phi}_z(x_{it}) + \check{\omega}_{zit} \quad (20)$$

with $T > t$, provided that the variables in x_{it} are pre-determined because then

$$E\left\{\frac{u_{ziT}}{\mu_{ziT}} | x_{it}\right\} = 0 \quad (21)$$

and we can interpret $\check{\phi}_{zit}(x_{it})$ as the expectation of the fixed effects given x_{it} :

$$\check{\phi}_z(x_{it}) = E\{\eta_i | x_{it}\}$$

As in the standard case we parameterize $\check{\phi}_z(x_{it})$ as an exponential function,⁵¹

$$\check{\phi}_z(x_{it}) = \exp(x_{zit}\gamma)$$

Notice, that given this control function we can transform our main regression equation as

$$\frac{PAT_{zit}}{\check{\phi}_z(x_{it})} = \exp(x_{it}\beta_z) + \exp(x_{it}\beta_z) \frac{\check{\omega}_{zit}}{\check{\phi}_z(x_{it})} + \frac{u_{zit}}{\check{\phi}_z(x_{it})} \quad (22)$$

where we replaced η_i by $\check{\phi}_{zit}(x_{it}) + \check{\omega}_{zit}$ and divided by $\check{\phi}_{zit}(x_{it})$. Because the x_{it} are pre-determined, given the definition of $\check{\omega}_{zit}$ and recalling the definition $\mu_{zit} = \exp(x_{it}\beta_z)$ we have that

⁵⁰For notational simplicity we write the following equation with just one future term. In practice we can improve efficiency by regressing on an average of future values $\frac{1}{T-t+1} \sum_{\tau=t}^T \frac{PAT_{zi\tau}}{\mu_{zi\tau}}$. In our regressions reported above we identify the control function from averages over the current and one future period; i.e. $T = t + 1$

⁵¹In theory we can even allow a more flexible specification where the conditional expectation varies over time; i.e. $\check{\phi}_{zt}(x_{it}) = E_t\{\eta_i | x_{it}\}$. This could reflect firms learning more about their fixed effect over time for instance. In practice this increases the number of parameters to be estimated greatly and becomes computationally very burdensome. In our baseline results we therefore fix $\check{\phi}_z(\cdot)$ over time.

$$E \left\{ \left(\mu_{zit} \frac{\check{\omega}_{zit}}{\check{\phi}_z(x_{it})} + \frac{u_{zit}}{\check{\phi}_z(x_{it})} \right) | x_{it} \right\} = 0 \quad (23)$$

Hence, we have two equations that depend on each other as well as two sets of moment conditions. We can consequently estimate equations (20) and (22) as a system of two simultaneous equations using the sample analog of the following moments

$$E \left\{ \left(\begin{array}{c} \frac{PAT_{zit}}{\check{\phi}_z(x_{it})} - \mu_{zit} \\ \frac{PAT_{ziT}}{\mu_{ziT}} - \check{\phi}_z(x_{it}) \end{array} \right) | x_{it} \right\} = 0$$

We refer to this approach below as the control function fixed effects estimator (CFX).

In addition to these three dynamic count data approaches we also explore the common practice of implementing equation (19) as a linear panel data estimator by taking logs of the dependent variable after simply adding the value of unity (an arbitrary constant); i.e. the regression equation becomes:

$$\ln(1 + PAT_{zit}) = x_{it}\beta_z + \alpha_{zi} + \varepsilon_{zit}$$

Although this model has undesirable features like generating negative predicted values of patenting it is attractive because it is straightforward to estimate a relative clean vs. dirty regression; i.e.

$$\ln(1 + PAT_{Clean,it}) - \ln(1 + PAT_{Dirty,it}) = x_{it}(\beta_{Clean} - \beta_{Dirty}) \quad (24)$$

$$+ (\alpha_{Clean,i} - \alpha_{Dirty,i}) + (\varepsilon_{Clean,it} - \varepsilon_{Dirty,it}) \quad (25)$$

We show in the results section that the results are qualitatively similar no matter which precise estimation technique we use.

B Data Appendix

B.1 Basic dataset

As described in the main text we draw from PATSTAT data all patent filings relating to IPC classes over clean and dirty auto innovation as defined in Table 1 and illustrated in Appendix Figure A1-A4.⁵² Our patent data is drawn from the World Patent Statistical Database (PATSTAT) maintained by the European Patent Office. We use the September 2009 version of PATSTAT. The innovation outcomes we use as the dependent variable and in the construction of the spillovers and own lagged innovation stocks are triadic patents (filed in all three of USPTO, EPO or JPO). For the weights (see below) we use a wider definition to patenting.

If a single patent filing has multiple IPC codes we include it so long as at least one of the IPC codes relates to clean or dirty innovation.⁵³ Patents are coded by whichever firm

⁵²To identify clean and dirty innovations filed at the US patent office we use the same IPC codes as the ones used for EPO and JPO patents. However, the USPTO has only recently adopted the IPC classification so a few older US patents do not have IPC codes. We therefore complement IPC codes with their US equivalents using the IPC/US concordance table available on the USPTO website

⁵³In the small number of cases where a patent had both a dirty and a clean IPC code we coded the patent to be clean, but nothing hinges on this.

first applied so we ignore traded patents, but these are rare: less than 3% of triadic patents are traded. As is standard in the literature, patents are dated by their application/filing date as this is close to the time when the R&D was performed.

B.2 Identifying unique patent holders

The PATSTAT database reports the name of patent applicants, but a common problem with patent data is that the name of patentees often varies, because of spelling mistakes, typographical errors and name variants. To identify unique patent holders we use the ECOOM-EUROSTAT-EPO PATSTAT Person Augmented Table (EEE-PPAT) database, available at <http://www.ecoom.be/nl/eee-ppat>, which provides a dictionary of harmonized patent applicants' names produced through a computer algorithm followed by visual inspection. We then manually check the name match, which allows us to put together companies that a typical computer algorithm would consider distinct. For example we match Ford Motor Company with Ford Werke, its German subsidiary. As a result, we are able to reduce the number of distinct patent holders of clean and dirty patents from 20,916 to 3,423; 2,427 of which are companies and 996 are individuals.

B.3 Firm-level weights

B.3.1 Weights based on patent portfolios

As explained above in the main text, the firm-specific fuel price is computed as the weighted geometric mean of the fuel prices across countries with weights reflecting the shares of the corresponding countries in the firm's patent portfolio. Our price variable is thus defined as:

$$\ln FP_{it} = \sum_c w_{ic0}^{FP} \ln FP_{ct} \quad (26)$$

where FP_{ct} is the tax-inclusive fuel price in country c at time t and w_{ic0}^{FP} is the firm-specific weight for country c . In order to make sure that the computed exposures are an exogenous source of variation across firms, the weights are calculated using the patent portfolio of each company over the 1965-1985 "pre-sample" period (with the regressions performed on the 1986-2005 period). We cross check the 1985 cut-off in the robustness section using 1965-1990 as the pre-sample period for weights and 1991-2005 for the regression sample.

To make matters concrete consider the example of Hitachi, a large Japanese car parts manufacturer, who filed 90,381 patents between 1965 and 1985. 63,175 of these filings were in Japan, 8,315 in the US and 3,498 in Germany. The rest were in a large number of other patent offices. Note that there are a larger number of filings than there are patents, as one invention can be filed in multiple patent offices. For example, Hitachi's patent 11464997 (this is the DOCDB family number) was developed by a Japanese inventor and filed in 1980 both in Japan and in the US. This patent enters twice in the patent-portfolio weight: once for Japan and once for US, since it indicates that both the US and Japan matter for Hitachi. Hitachi's 90,381 patents filed between 1965 and 1985 correspond to only 70,526 distinct inventions (or patent families), some of which were patented in several countries even though almost all of Hitachi's R&D activities are conducted in Japan (we use inventor location below for spillovers - see next section). In order to reflect the greater importance of larger countries when constructing fuel price weights, we take each country's average GDP over 1965-1985 into account (although nothing hinges on this for the results). The firm-specific weight for country c is thus equal to:

$$w_{ic0}^{FP} = \frac{s_{ic0}^{FP} GDP_c}{\sum_c s_{ic0}^{FP} GDP_c} \quad (27)$$

where s_{ic0}^{FP} is the share of country c in Hitachi's patent portfolio between 1965 and 1985 and GDP_c is the share of country c in the world's GDP over 1965-1985. The weights used for Hitachi are 68.8% for Japan, 23.9% for US and 2.7% for Germany. The weights summed across all other countries was 4.6% so the total weights sum to 100

We use the patent-portfolio weights, w_{ic}^{FP} , to construct the fuel price, fuel tax, GDP per capita and emission regulations variables. Note that in constructing the weights we use all patent filings from applicant firms who have filed at least one auto-related patent. These are all applicants who have filed a dirty or clean patent as defined by Table 1 from the OECD or in an IPC class defined as autos according to the OECD's cross walk. We could have also included patent filings by applicants who were part of the auto-related firms who had never filed for a clean or dirty auto patent according to our definitions. This would have increased our sample of patent filings from 4.5m to about 16m. We chose not to do this as many of these patents are only distantly related to autos and so would not be relevant for tracking the demand for cars. Going in the other direction, we could narrow our definition to include only patents in IPC classes we deem as clean or dirty and exclude all other patents by the same applicants. Building weights from this narrower pool led to similar results to those presented in the main text.

Although we have filings in 80 patent offices, the 25 countries we use for the fuel price data are Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Japan, Luxembourg, Mexico, Netherlands, New Zealand, Norway, Portugal, South Korea, Spain, Sweden, Switzerland, UK and the USA.

B.3.2 Weights based on location of inventors

To construct the firm-specific spillover pools in clean and dirty knowledge we use an analogous empirical strategy to that for the fuel price. The firm-specific spillover pool is computed as the weighted geometric mean of the knowledge pools across countries with weights reflecting the shares of the corresponding countries in the firm's pool of inventors. The spillover pool for firm i is calculated as:

$$SPILL_{zit} = \sum_c w_{ic0}^S SPILL_{zct} \quad (28)$$

where $z \in \{Dirty, Clean\}$ and $SPILL_{zct}$ is the spillover pool in country c at time t , which can be firm specific (see below). The spillover weight w_{ic0}^S is the share of all firm i 's inventors (i.e. where the inventors worked when they discovered the invention) in country c between 1965 and 1985.

This weight differs from the patent-portfolio weight w_{ic0}^{FP} described above in two ways. First, instead of using information on where each patent was filed (for example, the USPTO) we use the location of the patent inventors (who are more likely to benefit from other research conducted locally). Inventor countries are counted fractionally, so if a patent is filed by two inventors, one from Germany and one from the US, each country will receive one half.⁵⁴ Note that we use information on the *country of residence* of the inventor, not on his nationality. This seems natural because the geographical location of the inventor is likely to be the critical issue for knowledge spillovers.

⁵⁴We do this in order to avoid giving an artificially higher weight to a patent with multiple inventors compared to one with just a single named inventor

The second difference with respect to patent-portfolio weight is that each invention is only counted once, no matter in how many patent offices it has been filed. This is to avoid double counting. Returning to Hitachi’s patent 11464997 filed in 1980 both in Japan and in the US, this patent enters twice in the patent-portfolio weight but only once in the inventor location weight, as a Japan-developed invention. So although $w_{Hitachi, Japan}^{FP} = 0.688$ as above, $w_{Hitachi, Japan}^S = 0.99$. This indicates that although almost all Hitachi’s R&D is based in Japan, it sells car parts to a much wider geographical market.

The spillover pool $SPILL_{z,ct}$ is defined as:

$$SPILL_{z,ct} = \sum_{j \neq i} w_{jc}^S K_{z,jt} \quad (29)$$

i.e. the spillover pool of a country is the sum of all other firms’ patent stocks with a weight that depends on how many inventors the other firm has in that country. The aggregate stocks in equation (29) are thus entirely based on firm level stocks. This allows us to make out of sample simulations of aggregate stocks below using firm level equations only.

As an alternative strategy we constructed country level spillover stocks by aggregating over all patents of inventors based in that country:

$$SPILL_{z,ct} = \sum_{j \in \text{Inventors based in } c} K_{z,cjt} \quad (30)$$

where $K_{z,jct} = PAT_{z,it} + (1 - \delta)K_{z,jt-1}$ and $PAT_{z,jt}$ are the patents filed that associated with inventor j in year t . Empirically, both methods give very similar results. For consistency with our simulation results we use the first method (Equation (29)) throughout the paper.

We also use the inventor weights to construct the amount of R&D in energy-efficient transportation in country c at time t .

B.4 More descriptive statistics on patents filing and citations

For every patent in our data set, we know whether the invention has also been filed (prior to or following the first filing of the patent at USPTO, EPO or JPO) at any other patent office included in PATSTAT (over 80 offices). Table A2 provides information on the geographical coverage of clean and dirty innovations for some of the main patent offices. Interestingly, 31% of clean inventions are also patented in China. This is almost twice the rate for dirty inventions (18%). Germany’s specialization in traditional combustion engines is apparent from this table, with 61% of dirty patents protected in Germany but only 41% of clean patents.

When a patent is filed, it must include citations to earlier patents that are related to the new invention. Citations to earlier patents are indicative of the accumulated knowledge used by the inventor to develop the new invention (e.g. Jaffe and Trajtenberg, 2002). There are 181,151 citations for all clean and dirty triadic patents included in our data set (13.1 citations for the average patent). Among patents cited by clean patents, 47% are clean, whereas 5% are dirty. The remaining 48% refer to other, neither clean nor dirty (Table A3). If citations were not technology specific we would expect that the likelihood of a citation to a dirty patent would be three times higher than towards a clean patent. The likelihood of a clean on clean citations (47%) is almost as high as the likelihood of dirty on dirty citations (59%). This suggests that within category spillovers are much higher than between category spillovers. This is consistent with path-dependent innovation as the theory suggests.

B.4.1 Other data

Fuel price and fuel tax come from the International Energy Agency's Energy Prices and Taxes database, available online at <http://data.iea.org>. We use Households End-Use Prices in USD PPP/unit. Since data are available for both diesel and gasoline fuels, we define fuel price as the average of diesel and gasoline prices.

Data on public R&D expenditures comes from the IEA's Energy Technology Research and Development database, available online at <http://data.iea.org>. We use Total R&D in Million USD (2010 prices and exchange rates). We use the data on public R&D expenditures in "Energy efficiency - transportation" (Flow 13). This includes: electric cars, hybrid cars and stirling motors; analysis and optimization of energy consumption in the transport sector; efficiency improvements in light-duty vehicles, heavy-duty vehicles, non-road vehicles; public transport systems; engine-fuel optimization; use of alternative fuels (liquid, gaseous); fuel additives; diesel engines. Note that the IEA also reports R&D on "Hydrogen and fuel cells", in particular "fuel cells for mobile applications" but the data only start to be available in 2004, at the very end of our sample.

Data for environmental standards governing maximum permissible levels of tailpipe emissions for pollutants from new automobiles were sourced from a dataset originally constructed by Perkins and Neumayer (2012). Countries' regulatory stringency is coded on a scale of 0 to 5. The basis of the classification scheme is the European Union's (EU) Euro emission standards which were originally implemented across member states in 1992 and have subsequently been tightened in a series of incremental steps. Countries are coded 0 if they had no national emissions standards in place for new vehicles, or if standards were less stringent than the equivalent of Euro 1, during the year in question. Countries where Euro 1 or its equivalent was legally enforceable are coded 1, and so on, with 5 for countries having implemented the equivalent of the Euro 5 standard.

Data on GDP, GDP per capita and population are taken from the World Bank's World Development Indicators, available at <http://data.worldbank.org/>. GDP and GDP per capita are PPP and constant 2005 USD.

Sales data used to compare the patent weights with sales distribution are from company accounts (see the URLs in notes to Table A1).

Figure 1: Average fuel price and fuel tax 1986-2005
Panel A: Fuel price **Panel B: Fuel tax**

Note: this graph shows the average annual price of fuel for all countries available in the IEA database. The fuel price is the average between diesel and gasoline price. Prices are in 2005 USD PPP.
 Source: IEA.

Note: this graph shows the average annual tax on fuel for all countries available in the IEA database. The fuel tax is the average between diesel and gasoline tax. Tax is in 2005 USD PPP.
 Source: IEA.

Figure 2: Fuel price time series

Note: this graph shows the annual price of fuel for all countries available in the IEA database. The fuel price is the average between diesel and gasoline price. Prices are in 2005 USD PPP. Source: IEA.

Figure 3: Fuel Tax time series

Note: this graph shows the average annual tax on fuel for all countries available in the IEA database. The fuel tax is the average between diesel and gasoline tax. Taxes are in 2005 USD PPP.
Source: IEA.

Figure 4: Number of clean and dirty triadic patents 1978-2005

Note: this graph shows the number of annual triadic patents filed worldwide between 1978 and 2005 in clean and dirty technologies.
Source: authors' calculations based on the PATSTAT database.

Figure 5: Simulations over time of the effects of increases in fuel price
A: No change in fuel price **B: 10% increase in fuel prices**

Notes: these graphs show the simulated evolution of the aggregate clean and dirty knowledge stocks between 2005 and 2030 depending on the variation in fuel prices. The knowledge stock is the discounted sum of past patents. Fuel prices are assumed to increase at once in 2005 and remain constant thereafter. Simulations are based on CFX estimations presented in Table 6 columns (1) and (4).

Figure 6: Simulations over time of the effect of a 40% increase in fuel prices allowing for a negative effect of the carbon tax on GDP per capita growth
A: Baseline case : No effect of carbon tax on GDP per capita growth

B: Tax reduces GDP per capita growth by 0.25 percentage points

C: Tax reduces GDP per capita growth by 0.50 percentage points

D: Tax reduces GDP per capita growth by 0.75 percentage points

E: Tax reduces GDP per capita growth by 1.0 percentage points

Notes: These graphs show the simulated evolution of the aggregate clean and dirty knowledge stocks between 2005 and 2030 after a fuel price increase of 40% using the model in Table 6 columns (1) and (4). We consider a negative effect on per capita GDP growth of the carbon tax of between zero as in the baseline case (Panel A replicates Panel E of Figure 5) and one percentage point (in Panel E).

Figure 7: Simulations over time based on partial updating of innovation stock variables

Notes: these graphs show the simulated evolution of the aggregate clean and dirty knowledge stocks between 2005 and 2030. The knowledge stock is the discounted sum of past patents. Fuel prices are assumed to increase at once in 2005 and remain constant thereafter. In Panels A and B knowledge stocks and spillover stocks are recursively updated using the estimates from Table 6 columns (1) and (4). In panels C and D we switch off the effects of past innovation stocks by the firm itself and of spillovers. In all figures we assume a 1.5% growth rate of per capita GDP.

Table 1: Definition of IPC patent classes for clean and dirty patents

Panel A- Clean patents

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
Hydrogen vehicles / 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

Panel B- Dirty patents

Description	IPC code
Internal combustion engine	
Internal-combustion piston engines; combustion engines in general	F02B
Controlling combustion engines	F02D
Cylinders, pistons, or casings for combustion engines; arrangement of sealings in combustion engines	F02F
Supplying combustion engines with combustible mixtures or constituents thereof	F02M
Starting of combustion engines	F02N
Ignition (other than compression ignition) for internal-combustion engines	F02P

Panel C- Grey patents

Description	IPC code
Fuel efficiency of internal combustion engines	
Fuel injection apparatus	F02M39-71
Idling devices for carburettors preventing flow of idling fuel	F02M3/02-05
Apparatus for adding secondary air to fuel-air mixture	F02M23
Engine-pertinent apparatus for adding non-fuel substances or small quantities of secondary fuel to combustion-air, main fuel, or fuel-air mixture	F02M25
Electrical control of supply of combustible mixture or its constituents	F02D41
Methods of operating engines involving adding non-fuel substances or anti-knock agents to combustion air, fuel, or fuel-air mixtures of engines, the substances including non-airborne oxygen	F02B47/06

Table 2: Descriptive statistics

Variable	Mean	SD	Min	Max
Clean Patents (PAT_{Clean})	0.081	1.231	0	125
Dirty Patents (PAT_{Dirty})	0.227	3.424	0	355
Fuel Price ($\ln FP$)	-0.276	0.251	-1.053	0.438
Government R&D subsidies ($\ln R\&D$)	3.885	1.447	0	5.725
Emission Regulations Index	1.573	1.334	0	5
Clean Spillover ($\ln SPILL_C$)	3.774	1.258	-9.864	7.071
Dirty Spillover ($\ln SPILL_D$)	5.401	0.991	-5.509	7.677
Own Stock Clean innovation ($\ln K_C$)	-0.174	0.790	-6.718	5.740
Own Stock Dirty innovation ($\ln K_D$)	-0.910	1.618	-7.593	6.958

Notes: These are the values from our regression sample of 68,240 observations across 3,412 firms between 1986 and 2005. Emission Regulations for maximum level of tailpipe emissions for pollutants for new automobiles are coded between 0 and 5 following Dechezlepretre, Perkins and Neumayer (2012). Government R&D subsidies on clean transportation is from the IEA. See Appendix B for exact definitions

Table 3: Regressions of clean and dirty patents

Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)
	Clean Patents			Dirty Patents		
Fuel Price	0.970***	0.962**	0.886**	-0.565***	-0.553***	-0.644***
$\ln(FP)$	(0.374)	(0.379)	(0.362)	(0.146)	(0.205)	(0.143)
R&D subsidies		-0.005	-0.001		-0.006	-0.014
$\ln(R\&D)$		(0.025)	(0.024)		(0.021)	(0.021)
Emission Regulation			0.055			0.046
			(0.276)			(0.197)
Clean Spillover	0.268***	0.301***	0.266***	-0.093*	-0.078	-0.058
	(0.076)	(0.087)	(0.087)	(0.048)	(0.067)	(0.066)
Dirty Spillover	-0.168**	-0.207**	-0.160*	0.151**	0.132	0.114
	(0.085)	(0.098)	(0.097)	(0.064)	(0.082)	(0.081)
Own Stock Clean	0.306***	0.320***	0.303***	-0.002	-0.004	0.016
	(0.026)	(0.027)	(0.026)	(0.022)	(0.022)	(0.026)
Own Stock Dirty	0.139***	0.135***	0.139***	0.557***	0.549***	0.542***
	(0.017)	(0.017)	(0.017)	(0.031)	(0.022)	(0.020)
Observations	68240	68240	68240	68240	68240	68240
Firms	3412	3412	3412	3412	3412	3412

Notes: *, **, *** = significant at 10%, 5%, 1%. Standard errors are clustered at the firm level. Estimation is by the CFX method. All regressions include controls for GDP per capita, year dummies, fixed effects and three dummies for no clean knowledge, no dirty knowledge and no dirty or clean knowledge (in the previous year). Fuel price is the tax-fuel price faced. . R&D subsidies are public R&D expenditures in energy efficient transportation. Emissions Regulations are maximum levels of tailpipe emissions for pollutants from new automobiles. The dependent variable is the flow of clean patents in columns (1)-(3) and the flow of dirty patents in columns (4)-(6).

Table 4: Disaggregating dirty patents into fuel efficiency for internal combustion engine (grey) and purely dirty

Dependent variable:	(1) Clean Patents	(2) Grey Patents	(3) Purely Dirty Patents
Fuel Price	0.848* (0.461)	0.282 (0.398)	-0.832*** (0.214)
R&D subsidies	0.031 (0.047)	0.081** (0.034)	-0.02 (0.030)
Clean Spillover	0.333** (0.165)	-0.171* (0.098)	-0.014 (0.094)
Grey Spillover	0.215 (0.228)	0.173 (0.112)	0.235** (0.102)
Purely Dirty Spillover	-0.509 (0.377)	0.045 (0.136)	-0.208 (0.161)
Own Stock Clean	0.379*** (0.090)	-0.005 (0.035)	0.047 (0.035)
Own Stock Grey	0.185* (0.106)	0.418*** (0.035)	-0.141*** (0.025)
Own Stock Purely Dirty	-0.011 (0.066)	0.192*** (0.038)	0.544*** (0.026)
Observations	68240	68240	68240
Firms	3412	3412	3412

Notes: *, **, *** = significant at 10%, 5%, 1%. Standard errors are clustered at the firm level. Estimation is by the CFX method. This table disaggregates the dirty patents into those that are “grey” (related to fuel efficiency) and those that are not (“purely dirty”). We construct all spillovers and own past stocks based on this disaggregation and include on the right hand side (hence two extra terms compared to Table 3). We estimate two dirty equations, one where grey innovations are the dependent variable (in column (2)) and one for the purely dirty in column (3). All regressions include controls for GDP per capita, year dummies, fixed effects and 4 dummies for no own innovations in (i) clean, (ii) grey (iii) dirty and (iv) no clean, grey nor purely dirty in the previous year. Fuel price is the tax-inclusive fuel price faced. R&D subsidies are public R&D expenditures in energy efficient transportation.

Table 5: Alternative Econometric Models

Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)
	Clean Patents		Dirty Patents		Difference between Clean and Dirty $\ln(1 + PAT_{Clean}) - \ln(1 + PAT_{Dirty})$	
Model	HHG	BGVR	HHG	BGVR	Quasi Linear Within Groups	
Fuel Price	0.295 (1.062)	0.672** (0.332)	-2.457*** (0.897)	-0.614*** (0.192)	0.141** (0.061)	0.143** (0.061)
Clean Spillover	0.495** (0.236)	0.294*** (0.077)	0.393** (0.197)	-0.136** (0.054)	-0.007 (0.007)	-0.009 (0.007)
Dirty Spillover	-0.409 (0.484)	-0.277*** (0.084)	0.254 (0.300)	0.198*** (0.065)	0.015 (0.014)	0.010 (0.014)
Own Stock Clean	0.424*** (0.051)	0.883*** (0.031)	0.042 (0.036)	-0.003 (0.021)	0.048*** (0.007)	0.059*** (0.011)
Own Stock Dirty	0.133 (0.087)	0.091*** (0.029)	0.648*** (0.042)	1.069*** (0.022)	-0.016*** (0.004)	-0.010 (0.008)
Country X year effects	no	no	no	no	no	Yes
Firm fixed effects	yes	yes	yes	yes	yes	Yes
Observations	22420	68240	42300	68240	68240	68240
Firms	1121	3412	2115	3412	3412	3412

Notes: *,**,***= significant at 10%, 5%, 1%. Standard errors are clustered at the firm level. Regressions are same specifications as Table 3, i.e. column (3) for clean and column (6) for dirty. Fuel price is the tax- inclusive fuel price faced by the firm. The dependent variable is the flow of clean patents in columns (1)-(2), the flow of dirty patents in columns (3)-(4) and the log-ratio of clean to dirty patents in columns (5) and (6). Different columns control for fixed effects in different ways: HHG is the Hausman, Hall and Griliches (1984), BGVR is Blundell, Griffith and Van Reenen (1999) and last two columns are Within Groups (i.e. adding a dummy variable for each firm).

Table 6: Regressions for sample of firms with at least one pre-sample clean or dirty patent

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
	Clean Patents			Dirty Patents		
Model	CFX	HHG	BGVR	CFX	HHG	BGVR
Fuel Price	0.632** (0.296)	-0.293 (1.091)	0.825** (0.331)	-0.580*** (0.147)	-2.194*** (0.738)	-0.488*** (0.171)
Clean Spillover	0.240*** (0.068)	0.451* (0.247)	0.317*** (0.076)	-0.07 (0.051)	0.358 (0.230)	-0.126** (0.057)
Dirty Spillover	-0.152** (0.074)	-0.223 (0.473)	-0.281*** (0.085)	0.139** (0.068)	0.395 (0.280)	0.197*** (0.069)
Own Stock Clean	0.300*** (0.025)	0.403*** (0.060)	0.834*** (0.038)	-0.001 (0.027)	0.126*** (0.037)	0.002 (0.021)
Own Stock Dirty	0.142*** (0.017)	0.13 (0.089)	0.098*** (0.032)	0.523*** (0.018)	0.467*** (0.045)	1.040*** (0.022)
Observations	25400	7900	25400	25400	13340	25400
Firms	1270	395	1270	1270	667	1270

Notes: *,**,***= significant at 10%, 5%, 1%. Standard errors are clustered at the firm level. This is a sub-sample of the data in Table 3 where we condition on firms having at least one patent in the pre-sample period. All regressions include controls for GDP per capita, fixed effects, year dummies, three dummies for no clean knowledge, no dirty knowledge and no dirty or clean knowledge (in the previous year). Fuel price is the tax- inclusive fuel price faced by the firm. The dependent variable is the flow of clean patents in columns (1)-(3) and the flow of dirty patents in columns (4)-(6). HHG is the Hausman et al (1984) method; BGVR is the Blundell et al (1999) method and CFX is the Control Function Fixed Effect method.

Table 7: Regressions with fuel taxes instead of fuel price

Dependent variable Model	(1)	(2)	(3)	(4)	(5)	(6)
	Clean Patents			Dirty Patents		
	CFX	HHG	BGVR	CFX	HHG	BGVR
Fuel Tax	0.400** (0.167)	-0.969 (0.901)	0.227 (0.203)	-0.229*** (0.069)	-2.643*** (0.850)	-0.301*** (0.091)
Clean Spillover	0.284*** (0.075)	0.442* (0.228)	0.286*** (0.077)	-0.085* (0.047)	0.394 (0.257)	-0.142*** (0.049)
Dirty Spillover	-0.193** (0.084)	-0.433 (0.487)	-0.275*** (0.077)	0.141** (0.061)	0.093 (0.288)	0.204*** (0.063)
Own Stock Clean	0.327*** (0.027)	0.430*** (0.052)	0.884*** (0.032)	-0.008 (0.021)	0.051 (0.036)	-0.005 (0.021)
Own Stock Dirty	0.134*** (0.017)	0.126 (0.087)	0.091*** (0.029)	0.546*** (0.028)	0.645*** (0.041)	1.071*** (0.022)
Observations	68240	22420	68240	68240	42300	68240
Firms	3412	1121	3412	3412	2115	3412

Notes: *, **, *** = significant at 10%, 5%, 1%. Standard errors are clustered at the firm level. All regressions include controls for GDP per capita, year dummies, and three dummies for no clean knowledge, no dirty knowledge and no dirty or clean knowledge in the previous year. The dependent variable is the flow of clean patents in columns (1)-(3) and is the flow of dirty patents in columns (4)-(6). HHG is the Hausman et al (1984) method and BGVR is the Blundell et al (1999) method.

**Table 8: Alternative sample period
(pre-sample period for weights from 1990 and before, regressions run on data 1991-2005)**

Dependent variable Model	(1)	(2)	(3)	(4)	(5)	(6)
	Clean Patents			Dirty Patents		
	CFX	HHG	BGVR	CFX	HHG	BGVR
Fuel Price	0.806** (0.341)	-0.742 (1.110)	-0.038 (0.315)	-0.235 (0.233)	-2.547*** (0.904)	-0.602** (0.273)
Clean Spillover	0.177** (0.077)	0.684* (0.381)	0.390*** (0.111)	-0.05 (0.066)	0.763* (0.397)	-0.066 (0.093)
Dirty Spillover	-0.106 (0.084)	-0.31 (0.549)	-0.367*** (0.138)	0.136* (0.075)	0.024 (0.334)	0.134 (0.094)
Own Stock Clean	0.349*** (0.023)	0.258*** (0.069)	0.892*** (0.035)	0.009 (0.032)	0.128** (0.051)	0.024 (0.022)
Own Stock Dirty	0.136*** (0.018)	0.153 (0.097)	0.138*** (0.042)	0.519*** (0.053)	0.318*** (0.060)	1.098*** (0.032)
Observations	50820	15105	50820	50820	23985	50820
Firms	3388	1007	3388	3388	1599	3388

Notes: *, **, *** = significant at 10%, 5%, 1%. Standard errors are clustered at the firm level. All regressions include controls for GDP per capita, year dummies, and three dummies for no clean knowledge, no dirty knowledge and no dirty or clean knowledge in the previous year. Fuel price is the tax- inclusive fuel price faced by the firm (using pre-sample patent portfolios as weights). The dependent variable is the flow of clean patents in columns (1)-(3) and the flow of dirty patents in columns (4)-(6). HHG is the Hausman et al (1984) method and BGVR is the Blundell et al (1999) method.

Table 9: Alternative dynamic specifications on fuel price

Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Clean Patents					Dirty Patents				
Fuel Price in t $\ln(\text{FP}_t)$	0.882** (0.354)					-0.550*** (0.152)				
Fuel Price in $t-1$ $\ln(\text{FP}_{t-1})$		0.970*** (0.374)					-0.565*** (0.146)			
Fuel Price in $t-2$ $\ln(\text{FP}_{t-2})$			1.102*** (0.390)					-0.568*** (0.140)		
Fuel Price in $t-3$ $\ln(\text{FP}_{t-3})$				1.081*** (0.401)					-0.571*** (0.138)	
Fuel Price (Popp, 2002) $\ln(\text{FP}_{\text{Popp}})$					1.047*** (0.403)					-0.591*** (0.157)
Observations	68240	68240	68240	68240	68240	68240	68240	68240	68240	68240
Firms	3412	3412	3412	3412	3412	3412	3412	3412	3412	3412

Notes: **, *** = significant at 10%, 5%, 1%. Standard errors are clustered at the firm level. Estimation is by the CFX method. All regressions include controls for GDP per capita, year dummies, fixed effects (BGV method) and three dummies for no clean knowledge, no dirty knowledge and no dirty or clean knowledge in the previous year. Fuel price is the tax-inclusive fuel price. FP_{Popp} is the geometrically weighted average fuel price from 1978 until current year with a discount factor of 0.829 (following Popp, 2002). The dependent variable is the flow of clean patents in columns (1)-(5) and is the flow of dirty patents in columns (6)-(10). HHG is the Hausman et al (1984) method and BGVR is the Blundell et al (1999) method.

Figure A1: Front page of patent US 6456041 B1

US006456041B1

(12) **United States Patent**
Terada et al.

(10) **Patent No.:** **US 6,456,041 B1**
(45) **Date of Patent:** **Sep. 24, 2002**

(54) **POWER SUPPLY SYSTEM FOR ELECTRIC VEHICLE**
(75) Inventors: **Junji Terada; Satoshi Yamamoto**, both of Iwata (JP)
(73) Assignee: **Yamaha Hatsudoki Kanushiki Kaisha, Iwata (JP)**
(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

5,789,898 A * 8/1998 Suzuki et al. 320/104
5,798,702 A * 8/1998 Okamoto et al. 320/106
5,834,922 A * 11/1998 Ikawa et al. 320/136
5,864,224 A * 1/1999 Takechi et al. 320/152
5,896,024 A * 4/1999 Bradus et al. 320/125
5,942,878 A * 8/1999 Ito 320/131
5,963,140 A * 10/1999 Kawaguchi et al. 320/127
5,982,148 A * 11/1999 Mercer 320/134
6,034,507 A * 3/2000 Ikawa et al. 320/136
6,114,836 A * 9/2000 Hagiwara et al. 320/132

* cited by examiner

(21) Appl. No.: **09/601,391**
(22) PCT Filed: **Dec. 28, 1999**
(86) PCT No.: **PCT/JP99/07418**
§ 371 (c)(1),
(2), (4) Date: **Oct. 20, 2000**

Primary Examiner—Edward H. Tso
Assistant Examiner—Pia Tibbits
(74) Attorney, Agent, or Firm—Ernest A. Beutler

(87) PCT Pub. No.: **WO00/38944**
PCT Pub. Date: **Jul. 6, 2000**

(30) **Foreign Application Priority Data**
Dec. 28, 1998 (JP) 10-374096
Oct. 25, 1999 (JP) 11-302929
(51) Int. Cl.⁷ **H02J 7/00**
(52) U.S. Cl. **320/132**
(58) Field of Search 320/131, 132,
320/DIG. 19, DIG. 21

(57) **ABSTRACT**

A power source system 402 for electric motor-operated vehicles adapted to determine an actual capacity of a chargeable battery 400, namely the maximum capacity learning value at the current time point after it is used in cycles, comprises; a discharge means 403 for performing refreshment discharge of the chargeable battery 400, a charge control means 404 for controlling the discharge means 403 to perform the refreshment discharge with a discharge current including a pulse waveform portion, and an actual capacity learning means 405 for learning the actual capacity of the chargeable battery 400 on the basis of discharge capacity including the discharge capacity at the time of refreshment with the pulse waveform current.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,680,031 A * 10/1997 Pavlovic et al. 320/145

19 Claims, 15 Drawing Sheets

Figure A2: Disaggregation of F02F group

F02F CYLINDERS, PISTONS, OR CASINGS FOR COMBUSTION ENGINES; ARRANGEMENTS OF SEALINGS IN COMBUSTION ENGINES (specially adapted for rotary-piston or oscillating-piston internal-combustion engines F02B; specially adapted for gas-turbine plants F02C; specially adapted for jet-propulsion plants F02K) [2]

- (1) Attention is drawn to the Notes preceding class F01.
(2) Class F16 takes precedence over this subclass, except for subject matter specific to combustion engines.

1/00	Cylinders; Cylinder heads (in general F16J)	1/40	 cylinder heads with means for directing, guiding, or distributing liquid stream (F02F 1/38 takes precedence)
1/02	. . having cooling means (cylinder heads F02F 1/26)		
1/04	. . . for air cooling		
1/06	. . . Shape or arrangement of cooling fins; Finned cylinders	1/42	. . Shape or arrangement of intake or exhaust channels in cylinder heads
1/08	 running-liner and cooling-part of cylinder being different parts or of different material		
1/10	. . . for liquid cooling	3/00	Pistons (in general F16J)
1/12	. . . Preventing corrosion of liquid-swept surfaces	3/02	. . having means for accommodating or controlling heat expansion
1/14	. . . Cylinders with means for directing, guiding, or distributing liquid stream	3/04	. . having expansion-controlling inserts
1/16	. . . Cylinder liners of wet type	3/06	. . . the inserts having bimetallic effect
1/18	. Other cylinders	3/08	. . . the inserts being ring-shaped
1/20	. . characterised by constructional features providing for lubrication	3/10	. having surface coverings (F02F 3/02 takes precedence)
1/22	. . characterised by having ports in cylinder wall for scavenging or charging	3/12	. . on piston heads
1/24	. Cylinder heads	3/14	. . . within combustion chambers
1/26	. . having cooling means	3/16	. having cooling means
1/28	. . . for air cooling	3/18	. . the means being a liquid or solid coolant, e.g. sodium, in a closed chamber in piston
1/30	 Finned cylinder heads	3/20	. . the means being a fluid flowing through or along piston
1/32	 the cylinder heads being of overhead-valve type	3/22	. . . the fluid being liquid
1/34	 with means for directing or distributing cooling medium (F02F 1/32 takes precedence)	3/24	. having means for guiding gases in cylinders, e.g. for guiding scavenging charge in two-stroke engines
1/36	 for liquid cooling	3/26	. having combustion chamber in piston head (the surface thereof being covered F02F 3/14)
1/38	 the cylinder heads being of overhead-valve type	3/28	. Other pistons with specially-shaped head
		5/00	Piston rings, e.g. associated with piston crown
		7/00	Casings, e.g. crankcases (engine casings in general F16M)
		11/00	Arrangements of sealings in combustion engines (piston rings F02F 5/00; sealings <u>per se</u> F16J)

Figure A3: Front page and diagram for patent EP 0979940 B1

Europäisches Patentamt European Patent Office Office européen des brevets		(11) EP 0 979 940 B1
	(12) EUROPEAN PATENT SPECIFICATION		
(45) Date of publication and mention of the grant of the patent: 17.11.2004 Bulletin 2004/47	(51) Int Cl.7: F02M 59/44 , F02M 59/36, F02M 63/02, F02M 39/00, F02D 41/06, F02D 41/38, F02M 69/34, F02M 39/02		
(21) Application number: 99115753.8			
(22) Date of filing: 10.08.1999			
(54) Method and device for controlling fuel injection into an internal combustion engine Verfahren und Vorrichtung zum Steuern der Kraftstoffeinspritzung in einer Brennkraftmaschine Procédé et dispositif de commande de l'injection de carburant dans un moteur à combustion interne			
(84) Designated Contracting States: DE ES FR GB IT SE	• Suzui, Kosuke, c/o Toyota Jidosha K. K. Toyota-shi, Aichi-ken, 471-8571 (JP)		
(30) Priority: 11.08.1998 JP 22690898			
(43) Date of publication of application: 16.02.2000 Bulletin 2000/07	(74) Representative: Leson, Thomas Johannes Alois, Dipl.-Ing. et al TBK-Patent, P.O. Box 20 19 18 80019 München (DE)		
(73) Proprietor: TOYOTA JIDOSHA KABUSHIKI KAISHA Aichi-ken 471-8571 (JP)			
(72) Inventors: • Koga, Nobuhiko, c/o Toyota Jidosha K. K. Toyota-shi, Aichi-ken, 471-8571 (JP) • Kojima, Susumu, c/o Toyota Jidosha K. K. Toyota-shi, Aichi-ken, 471-8571 (JP) • Takeda, Keiso, c/o Toyota Jidosha K. K. Toyota-shi, Aichi-ken, 471-8571 (JP)	(56) References cited: EP-A- 0 481 964 EP-A- 0 677 655 US-A- 5 063 900 US-A- 5 605 133		
Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).			

EP 0 979 940 B1

Fig.1

Figure A4: Front page and diagram for patent EP 0402091 B1

Europäisches Patentamt European Patent Office Office européen des brevets		Publication number: 0 402 091 B1
---	--	--	---

⑫ **EUROPEAN PATENT SPECIFICATION**

④⑤ Date of publication of patent specification: **21.09.94**

⑤① Int. Cl.⁵: **F02B 27/00**

②① Application number: **90306087.9**

②② Date of filing: **05.06.90**

⑤④ **A four-cycle twelve cylinder engine.**

<p>③① Priority: 06.06.89 JP 143392/89 06.06.89 JP 143393/89</p> <p>④③ Date of publication of application: 12.12.90 Bulletin 90/50</p> <p>④⑤ Publication of the grant of the patent: 21.09.94 Bulletin 94/38</p> <p>⑤④ Designated Contracting States: DE FR GB</p> <p>⑤⑥ References cited:</p> <table style="width: 100%;"> <tr> <td>EP-A- 0 158 008</td> <td>FR-A- 2 591 665</td> </tr> <tr> <td>GB-A- 369 784</td> <td>JP-A- 6 341 621</td> </tr> <tr> <td>JP-A-61 237 823</td> <td>US-A- 2 034 368</td> </tr> <tr> <td>US-A- 4 698 876</td> <td>US-A- 4 708 097</td> </tr> <tr> <td>US-A- 4 829 941</td> <td></td> </tr> </table>	EP-A- 0 158 008	FR-A- 2 591 665	GB-A- 369 784	JP-A- 6 341 621	JP-A-61 237 823	US-A- 2 034 368	US-A- 4 698 876	US-A- 4 708 097	US-A- 4 829 941		<p>⑦③ Proprietor: Mazda Motor Corporation No. 3-1, Shinch Fuchu-cho Aki-gun Hiroshima-ken (JP)</p> <p>⑦② Inventor: Kurokawa, Toshikazu 327 Komatsubara, Akitsu-cho Toyota-gun, Hiroshima-ken (JP) Inventor: Tokushima, Takashige Seto-Haimu 4-11-20, Fuchu-cho Aki-gun, Hiroshima-ken (JP) Inventor: Sumimoto, Hiroshi 2-3-44 Arikiyo, Ondo-cho Aki-gun, Hiroshima-ken (JP) Inventor: Suemori, Tadao 2252-16 Iida, Hachihonmatsu-cho Higashi-Hiroshima-shi, Hiroshima-ken (JP)</p> <p>⑦④ Representative: Brunner, Michael John et al GILL JENNINGS & EVERY Broadgate House 7 Eldon Street London EC2M 7LH (GB)</p>
EP-A- 0 158 008	FR-A- 2 591 665										
GB-A- 369 784	JP-A- 6 341 621										
JP-A-61 237 823	US-A- 2 034 368										
US-A- 4 698 876	US-A- 4 708 097										
US-A- 4 829 941											

EP 0 402 091 B1

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid (Art. 99(1) European patent convention).

Rank Xerox (UK) Business Services
(3.10/3.09/3.3.3)

Figure A5: Patent weights vs inventor weights for US

Notes: this graph shows the share of companies' patent portfolio at the USPTO (on the y-axis) together with the share of inventors located in the US for the same companies (on the x-axis). The patent weight is used to calculate the firm-level fuel price and the inventor weight is used to calculate the firm-level spillover variables. Each point corresponds to a combination of patent weight and inventor weight for the US for a given company. We see (along the y-axis) that many companies file patent in the US but do not carry out R&D in this country. There are also a few companies (along the x-axis), which have R&D labs in the US but file their patents only in foreign countries.

Figure A6: Simulations over sample period
A: CFX

B: HHG

C: BGVR

Notes: These graphs show the simulated evolution of annual clean and dirty patents stocks from 1986 onwards. Prices and other exogenous variables (including time dummies, controls for GDP per capita) are set at their actual values. Own knowledge stocks as well as spillover effects are simulated. The simulated series are therefore directly comparable to the actual knowledge stocks over time, which are also plotted in the graphs. Estimation taken from the relevant columns of Table 6.

Table A1: Car sales and patent portfolios across countries for selected large auto firms

Company and markets	Car sales	Patent weights
TOYOTA (2003-2005)		
Japan	0.34	0.42
North America	0.31	0.34
Europe	0.13	0.23
VW (2002-2005)		
Germany	0.19	0.52
UK	0.07	0.07
Spain	0.06	0.03
Italy	0.05	0.05
France	0.05	0.08
USA	0.07	0.14
Mexico	0.03	0.00
Canada	0.02	0.00
Japan	0.01	0.02
FORD (1992-2002)		
USA	0.59	0.59
Canada	0.04	0.01
Mexico	0.02	0.00
Britain	0.08	0.08
Germany	0.06	0.15
Italy	0.03	0.03
Spain	0.02	0.02
France	0.02	0.04
Australia	0.02	0.00
Japan	0.01	0.05
Peugeot (2001-2005)		
Western Europe	0.75	0.83
France	0.25	0.31
Other countries	0.50	0.52
The Americas	0.04	0.13
Asia-Pacific	0.12	0.04
Honda (2004-2005)		
Japan	0.23	0.31
North America	0.50	0.48
Europe	0.08	0.20

Notes: Car sales are taken from company annual reports from the years as noted. Patent weights are constructed from filings in each country across patent offices for the same years as noted. Sources for sales data are the following (last accessed 25th November 2012):

TOYOTA: http://www.toyota-global.com/investors/ir_library/annual/pdf/2005/pdf/04.pdf

VW (Volkswagen):

<http://www.volkswagen.co.uk/assets/common/content/volkswagen-world/annual-report-2003.pdf>;

<http://www.volkswagen.co.uk/assets/common/content/volkswagen-world/annual-report-2004.pdf>;

<http://www.volkswagen.co.uk/assets/common/content/volkswagen-world/annual-report-2005.pdf>;

FORD: http://corporate.ford.com/doc/2002_full.pdf

PEUGEOT: <http://www.psa-peugeot-citroen.com/en/publications>

HONDA: http://world.honda.com/investors/library/annual_report/2006/ar2006.pdf

Table A2: Geographical coverage of patent protection

Share of inventions also patented in:	Type of technology:	
	Clean	Dirty
Germany	40.9%	61.0%
China	31.1%	18.3%
Canada	30.4%	12.6%
Korea	16.6%	11.2%
Australia	15.8%	11.0%
Brazil	7.3%	10.7%
Spain	7.0%	10.6%
Austria	9.6%	9.0%
France	3.8%	3.9%
UK	3.4%	3.8%

Note: the patents in our data set are triadic patents, filed in USA, Japan and at the European Patent Office. The table reports the share of patents that are also filed in Germany, China, Canada, Korea Australia, Brazil, Spain, Austria, France and the UK for each category.

Source: authors' calculations based on the PATSTAT database.

Table A3: Citation patterns

Citing patent	Cited patent		
	Clean	Dirty	Other
Clean	46.8%	5.2%	48.0%
Dirty	1.5%	59.6%	38.9%
Other	0.2%	0.5%	99.3%

Note: the table shows the type of patents cited by triadic patents in clean, dirty and other (ie, neither clean nor dirty) technologies. For example, 46.8% of patents cited by clean patents are clean, 5.2% are dirty and 48.0% pertain to other technologies (ie, neither clean nor dirty).

Source: authors' calculations based on the PATSTAT database.

Table A4: Main clean patent holders 1978-2005 – Triadic patents

Company	Clean patents	Dirty patents	Other patents	Total patents
Toyota	568	1238	3500	5306
Nissan	472	811	1735	3018
Honda	374	904	1871	3149
Hitachi	169	746	6987	7902
Robert Bosch	111	2734	4534	7379
Siemens	105	426	6786	7317
Mitsubishi	95	445	8138	8678
Daimler-Benz	87	295	1421	1803
Samsung	75	3	5123	5201
NGK Spark Plug	74	195	1611	1880

Note: the table reports the top 10 clean triadic patent holders between 1978 and 2005. We also report the number of dirty patents and the number of total patents (including clean, dirty and other patents) held by these applicants.

Source: authors' calculations based on the PATSTAT database.

Table A5: Main dirty patent holders 1978-2005 – Triadic patents

Company	Dirty patents	Clean patents	Other patents	Total patents
Robert Bosch	2734	111	4534	7379
Toyota	1238	568	3500	5306
Honda	904	374	1871	3149
Nissan	811	472	1735	3018
Hitachi	746	169	6987	7902
Denso	454	38	947	1439
Mitsubishi	445	95	8138	8678
Siemens	426	105	6786	7317
Isuzu	336	28	236	600
Yamaha	312	48	869	1229

Note: the table reports the top 10 dirty triadic patent holders between 1978 and 2005. We also report the number of clean patents and the number of total patents (including clean, dirty and other patents) held by these applicants.

Source: authors' calculations based on the PATSTAT database.

Table A6: Main clean patent holders 1978-2005 – EPO

Company	Clean patents	Dirty patents	Other patents	Total patents
Toyota	473	1280	4272	6025
Nissan	423	730	2465	3618
Honda	378	886	2726	3990
Siemens	313	1612	32454	34379
Daimler-Benz	201	844	4910	5955
Hitachi	162	784	9838	10784
Ballard Power Systems	155	0	46	201
International Fuel Cells	153	31	1957	2141
Panasonic	135	2	6078	6215
Robert Bosch	132	4109	11627	15868

Note: the table reports the top 10 clean patent holders at the EPO between 1978 and 2005. We also report the number of dirty patents and the number of total patents (including clean, dirty and other patents) held by these applicants.

Source: authors' calculations based on the PATSTAT database.

Table A7: Main dirty patent holders 1978-2005 – EPO

Company	Dirty patents	Clean patents	Other patents	Total patents
Robert Bosch	4109	132	11627	15868
Siemens	1612	313	32454	34379
Toyota	1280	473	4272	6025
Honda	886	378	2726	3990
Daimler-Benz	844	201	4910	5955
Ford	825	72	2849	3746
Hitachi	784	162	9838	10784
Nissan	730	423	2465	3618
Audi	697	103	2821	3621
BMW	542	86	2626	3254

Note: the table reports the top 10 dirty patent holders at the EPO between 1978 and 2005. We also report the number of clean patents and the number of total patents (including clean, dirty and other patents) held by these applicants.

Source: authors' calculations based on the PATSTAT database.

Table A8: Main clean patent holders 1978-2005 – USPTO

Company	Clean patents	Dirty patents	Other patents	Total patents
Honda	909	3107	7767	11783
Toyota	735	2832	8753	12320
General Motors	532	1587	7923	10042
Nissan	474	2180	5508	8162
International Fuel Cells	429	75	4556	5060
Hitachi	360	1819	31719	33898
Ford	325	2112	5862	8299
Ballard Power Systems	255	2	84	341
Daimler-Benz	249	1571	6134	7954
Mitsubishi	228	2138	27985	30351

Note: the table reports the top 10 clean patent holders at the USPTO between 1978 and 2005. We also report the number of dirty patents and the number of total patents (including clean, dirty and other patents) held by these applicants.

Source: authors' calculations based on the PATSTAT database.

Table A9: Main dirty patent holders 1978-2005 – USPTO

Company	Dirty patents	Clean patents	Other patents	Total patents
Robert Bosch	4476	165	7774	12415
Honda	3107	909	7767	11783
Toyota	2832	735	8753	12320
Nissan	2180	474	5508	8162
Mitsubishi	2138	228	27985	30351
Ford	2112	325	5862	8299
Denso	1954	143	6149	8246
Hitachi	1819	360	31719	33898
General Motors	1587	532	7923	10042
Daimler-Benz	1571	249	6134	7954

Note: the table reports the top 10 dirty patent holders at the USPTO between 1978 and 2005. We also report the number of clean patents and the number of total patents (including clean, dirty and other patents) held by these applicants.

Source: authors' calculations based on the PATSTAT database.

Table A10: Using the biadic patents

Dependent variable Model	(1)	(2)	(3)	(4)	(5)	(6)
	Clean Patents			Dirty Patents		
	CFX	HHG	BGVR	CFX	HHG	BGVR
Fuel Price	0.980** (0.395)	0.425 (0.999)	0.845** (0.332)	-0.516*** (0.187)	-1.960** (0.791)	-0.368* (0.207)
Clean Spillover	0.233*** (0.079)	0.657*** (0.185)	0.314*** (0.077)	-0.058 (0.047)	0.277* (0.168)	-0.044 (0.059)
Dirty Spillover	-0.162* (0.090)	-0.955** (0.467)	-0.334*** (0.080)	0.101 (0.067)	-0.042 (0.265)	0.088 (0.073)
Own Stock Clean	0.385*** (0.035)	0.424*** (0.046)	0.888*** (0.030)	0.024 (0.022)	0.017 (0.034)	-0.017 (0.021)
Own Stock Dirty	0.124*** (0.019)	0.107 (0.078)	0.106*** (0.024)	0.517*** (0.013)	0.683*** (0.041)	1.090*** (0.023)
Observations	92700	29480	92700	92700	57500	92700
Firms	4635	1474	4635	4635	2875	4635

Notes: *, **, *** = significant at 10%, 5%, 1%. Standard errors are clustered at the firm level. Sample includes all patents taken out at both EPO and USPTO (Triadic patents used in the main paper are a sub-sample of these who were also filed in the JPO). All regressions include controls for GDP per capita, year dummies, three dummies for no clean knowledge, no dirty knowledge and no dirty or clean knowledge in the previous year. Fuel price is the tax-inclusive fuel price faced. The dependent variable is the flow of clean patents in columns (1)-(3) and is the flow of dirty patents in columns (4)-(6). HHG is the Hausman et al (1984) method, BGVR is the Blundell et al (1999) method, CFX is Control Function Fixed Effect method.

Table A11: Using citation weighted knowledge stocks

Dependent variable Model	(1)	(2)	(3)	(4)	(5)	(6)
	Clean Patents			Dirty Patents		
	CFX	HHG	BGVR	CFX	HHG	BGVR
Fuel Price	1.972*** (0.585)	1.59 (1.318)	2.909*** (0.714)	-0.817*** (0.167)	-1.004 (1.258)	0.960** (0.462)
Clean Spillover	0.957*** (0.203)	0.859** (0.360)	0.628*** (0.179)	-0.066 (0.088)	0.914*** (0.281)	-0.103 (0.120)
Dirty Spillover	-0.674*** (0.211)	-0.67 (0.549)	-0.333 (0.210)	0.186* (0.098)	0.168 (0.599)	0.232* (0.133)
Own Stock Clean	0.330*** (0.047)	0.412*** (0.086)	0.657*** (0.046)	0.073** (0.037)	0.225*** (0.056)	0.061 (0.045)
Own Stock Dirty	0.064 (0.043)	0.163* (0.088)	0.086* (0.052)	0.300*** (0.020)	0.259*** (0.060)	0.697*** (0.051)
Observations	68240	22420	68240	68240	42300	68240
Firms	3412	1121	3412	3412	2115	3412

Notes: *, **, *** = significant at 10%, 5%, 1%. Standard errors are clustered at the firm level. All regressions include controls for GDP per capita, year dummies, and three dummies for no clean knowledge, no dirty knowledge and no dirty or clean knowledge in the previous year. Fuel price is the tax-inclusive fuel price faced. The dependent variable is the flow of clean patents in columns (1)-(3) and is the flow of dirty patents in columns (4)-(6). HHG is the Hausman et al (1984) method, BGVR is the Blundell et al (1999) method, and CFX is Control Function Fixed Effect method.

Table A12: Controlling for GDP

Dependent variable Model	(1)	(2)	(3)	(4)	(5)	(6)
	Clean Patents			Dirty Patents		
	CFX	HHG	BGVR	CFX	HHG	BGVR
Fuel Price	0.857** (0.334)	0.203 (1.162)	0.618* (0.321)	-0.408 (0.356)	-2.521*** (0.878)	-0.410* (0.214)
GDP	0.219 (0.201)	0.138 (1.809)	0.021 (0.190)	0.248 (0.217)	0.896 (2.802)	0.212 (0.163)
GDP per capita	1.650 (1.500)	-2.174 (3.260)	2.342** (1.071)	-0.765 (0.578)	-3.178 (2.326)	-0.826* (0.455)
Clean Spillover	0.308*** (0.078)	0.478** (0.231)	0.296*** (0.074)	-0.107 (0.080)	0.405* (0.224)	-0.124** (0.052)
Dirty Spillover	-0.201** (0.086)	-0.438 (0.488)	-0.280*** (0.090)	0.164* (0.085)	0.238 (0.284)	0.188*** (0.064)
Own Stock Clean	0.302*** (0.032)	0.426*** (0.052)	0.883*** (0.032)	0.039 (0.033)	0.044 (0.037)	0.000 (0.023)
Own Stock Dirty	0.134*** (0.019)	0.131 (0.087)	0.091*** (0.029)	0.519*** (0.023)	0.648*** (0.042)	1.065*** (0.023)
Observations	68240	22420	68240	68240	42300	68240
Firms	3412	1121	3412	3412	2115	3412

Notes: *, **, *** = significant at 10%, 5%, 1%. Standard errors are clustered at the firm level. Estimation is by the various methods described in the Econometrics Section. All regressions include controls for GDP per capita, GDP, year dummies, three dummies for no clean knowledge, no dirty knowledge and no dirty or clean knowledge in the previous year. The dependent variable is the flow of clean patents in columns (1)-(3) and is the flow of dirty patents in columns (4)-(6). HHG is the Hausman et al (1984) method, BGVR is the Blundell et al (1999) method, and CFX is Control Function Fixed Effect method.