

**Policy-Induced Innovation in
Clean Technologies:
Evidence from the Car Market**

Rik L. Rozendaal, Herman Vollebergh

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Policy-Induced Innovation in Clean Technologies: Evidence from the Car Market

Abstract

This article tests the effects of fuel economy and greenhouse gas emission standards on the direction of innovation, in particular on breakthrough technologies in the automotive industry. We develop an intuitive measure of standard stringency that captures the policy's most important features for the decision as to whether or not to innovate. To test the role of these standards relative to prices and taxes, we construct a firm-level panel of patents in clean and dirty automotive technologies for the years 2000-2016. Our results indicate that standards are a very robust driver inducing clean innovation, whereas taxes also seem to play a role but prices (net of taxes) do not. This effect is driven by patenting for breakthrough technologies, in particular electric vehicle and hydrogen fuel cell technologies. We find no evidence that these policies negatively impact dirty innovation.

JEL-Codes: O300, Q550, Q580.

Keywords: environmental policy instruments, regulatory stringency, innovation, directed technical change.

*Rik L. Rozendaal**
Tilburg University / The Netherlands
r.l.rozendaal@tilburguniversity.edu

Herman Vollebergh
Tilburg University & PBL Netherlands
Environmental Assessment Agency
herman.vollebergh@pbl.nl

*corresponding author

November 11, 2021

This project originally started as Rik Rozendaal's Research Master Thesis. We would like to thank Malik Çürük, Arik Levinson, Sjak Smulders and participants at the GSS Seminar (Tilburg), the EAERE Summer School (Graz), the Macro Study Group (Tilburg) and the Dutch Economists Day (The Hague) for many helpful comments.

1 Introduction

Many countries regulate greenhouse gas (GHG) emissions from road transport in two ways: excise taxes on car fuels and standards that set targets for vehicles' fuel economy or emissions at the automobile manufacturer level. Taxes in most countries have been relatively stable in real terms over the past two decades, whereas the stringency of standards has increased rapidly in the EU and the US since new targets were announced in 2009 and 2010, respectively. For instance, the standard introduced by the EU in 2009 aimed to reduce average new vehicle CO₂ emissions per kilometer by about 15% in 2015 compared to the 2008 level. At the same time, we observe a surge in new technologies that are becoming more prominent in the automotive industry. Hybrid vehicles have become a common sight, and more radical innovations like vehicles powered by electricity are also quickly penetrating the car market. In fact, the global fleet of electric passenger vehicles has exploded from about 17,000 in 2010 to 7.2 million in 2019 (IEA, 2020).

A large literature compares taxes and standards in the car market from a variety of perspectives. In terms of welfare, fuel taxes are generally preferred to standards, as they directly affect the incentive to drive, rather than just new vehicle characteristics (e.g., Anderson and Sallee 2016). Other work has studied trading off vehicle characteristics (e.g., reducing weight or horsepower to increase fuel economy) and technology adoption as a response to regulation (e.g., Knittel 2011; Klier and Linn 2016; Kiso 2019; Reynaert 2021). In this paper we focus on the effect of different instruments to reduce emissions from car use on the direction of innovation, i.e. the invention of new technologies, using patent data.

A common view in the literature on technological change is that innovation is primarily driven by prices. Indeed, studies that focus on the invention of new technologies typically test the effects of prices on the direction of innovation. For instance, Popp (2002) shows that energy prices positively impact energy-efficient innovations and Aghion et al. (2016) show that fuel prices (which include taxes) steer technological change in a clean direction. These findings confirm the widespread belief among economists that price-based instruments are key to technological change and should be preferred to alternative policies. Our findings show that this view deserves some qualification. We find that a more precise identification of the potential impact of environmental policy instruments reveals a much more prominent

role of standards in innovation inducement, in particular in breakthrough technologies.

That standards stimulate the adoption of new automotive technologies is already demonstrated (e.g., Klier and Linn 2016). Changes in the composition of models in the car market are clearly sensitive to policies that stimulate fuel efficiency in addition to price based measures. However, only a few papers studied innovation effects of climate standards for automobiles in the market for ideas so far, but failed to find a robust impact (Crabb and Johnson 2010; Vollebergh 2010). A possible explanation is that standards were not stringent enough or even nonbinding in the considered sample periods. In other words, there was no urgency for firms to insert additional effort in the invention of breakthrough technologies because they could rely on cheaper, existing alternatives to comply. In addition, these papers measure standard stringency simply as the contemporaneous or lagged level of the standard. Measuring standard stringency in this manner might not fit the actual process of innovation.

Innovation is a dynamic process. Decisions to innovate, in particular investments in the market for ideas, are typically forward-looking and costly, which means that expectations about prices, market developments and regulation are relevant to a firm's decision making. Hence, when studying the determinants of innovation, it is important to take the role of this time dimension into account. We do this for both taxes and standards. First, we separate excise taxes from tax-exclusive fuel prices because the excise taxes usually are more predictable and less volatile than fuel prices, especially over the course of years. Moreover, taxes affect consumer markets in the first place and therefore only indirectly provide an impulse for car manufacturers decisions on innovation. By contrast, standards on cars are typically pool-based targets announced years in advance, allowing for many different response strategies by the car manufacturers (OECD, 2010). The need for breakthrough innovations only occurs if other strategies would be insufficient to comply with the (announced) regulation.

To better identify the role of standards in the innovation process we develop an intuitive measure that captures both the time dimension between the announcement of new targets and the year from which they are enforced, and its stringency by the level of the target relative to current average performance. We define stringency as the amount of CO₂ per kilometer that the average car manufacturer in a country (or the EU) needs to improve to meet the most recently announced target. We interact this with a dummy to only measure stringency

when a standard is binding, i.e. mandatory and stricter than current average performance. This measure combines the most important features of the instrument in a single dimension and is comparable across countries.

We test the impact of fuel prices and taxes as well as car standards, R&D subsidies and knowledge stocks on inventions using patents for automotive technologies as the key indicator. Our patent pool consists of triadic patent families¹ and we classify them as clean or dirty based on their IPC classification, following Aghion et al. (2016). Clean patents include those related to electric vehicles and to hydrogen fuel cells. Also patents related to hybrid vehicles are included in this category, although we do not consider those as breakthrough technologies because they still rely on the traditional technology of the internal combustion engine. Dirty patents are related to this internal combustion engine only, and can be split up in gray patents, which aim to increase fuel efficiency, and purely dirty patents. We compiled a firm-level panel for the years 2000 to 2016 and use data from 1978 to 1999 to construct pre-sample variables and weights. We use several empirical models (zero-inflated Poisson, Poisson, negative binomial) to estimate the effects of our variables of interest on clean and dirty patenting.

Our results indicate that both fuel taxes and GHG emission standards are important in clean technology inducement. The effect of taxes is not surprising and is in line with the literature (most notably, Aghion et al. 2016). However, this inducement is usually identified by tax-inclusive prices, a finding our estimations do not corroborate. The very strong effect of GHG emission standards is new.² Our results show that standards affected patenting for breakthrough technologies in particular. In other words, it is mainly the search for radically new technologies that increased due to the changes in regulation in the US and the EU around 2010. We find no evidence for the prediction that regulation shifts resources for invention away from dirty technologies, though we do find that the tax-exclusive fuel price negatively

¹A triadic patent family is a group of patents that include at least one patent at each of the three main patent authorities (those in the EU, US and Japan) and that protect a single invention.

²Aghion et al. (2016) also include a measure of standard stringency in their influential paper, namely the emissions regulation index from Dechezleprêtre et al. (2015). This measure, however, is concerned with air quality regulations on CO, NO_x and particulate matter, and not with climate regulations like standards on fuel economy and CO₂ emissions. Vollebergh (2010) uses both air quality and climate standards and elaborates on the differences between them.

affects dirty innovation. Our study focuses on the empirical effectiveness of the two policies that are in place rather than a welfare comparison between instruments.

Our work relates to two strands of literature. First, we contribute to the literature on instrument choice for environmental regulation, and especially to the debate on the regulation of GHG emissions from road transport. A common distinction in the literature is between market instruments, such as taxes, and command-and-control instruments, such as standards (see Hepburn 2006 and Goulder and Parry 2008 for literature reviews). Though economists often view standards as simple, uniform restrictions, they exist in many forms (Helfand 1991; Vollebergh and Van der Werf 2014). Indeed, the standards we study regulate car manufacturers' sales-weighted average of a particular vehicle characteristic, namely fuel economy or GHG emissions per kilometer. Whereas market instruments are generally preferred from an abatement cost perspective, this ranking is less clear when it comes to incentives for innovation (e.g., Montero 2002; Fischer et al. 2003; Requate 2005; Vollebergh 2007; Kemp and Pontoglio 2011).

The debate on instrument choice to combat climate change in the car market is centered around fuel taxes and several types of standards, though vehicle taxes, feebates and subsidies are also considered (e.g., Chen et al. 2021; Springel 2021). Anderson and Saltee (2016) use a simple model to illustrate the welfare effects of various instruments. Parry and Small (2005) discuss the optimal fuel tax and Jacobsen (2013) and Bento et al. (2020) analyze the welfare effects of the American Corporate Average Fuel Economy (CAFE) standards. Jacobsen and van Benthem (2015) highlight the importance of scrappage decisions and the effects of policies on the market for used vehicles. Holland et al. (2021) study the welfare effects of banning gasoline vehicle production in a model with decreasing costs of electric vehicles, decreasing pollution from electricity and increasing substitutability between clean and dirty vehicles. Kellogg (2018) applies Weitzman's (1974) prices versus quantities framework to optimal standard setting for automobiles when fuel prices are uncertain. Leard et al. (2017) study the interaction between standards and fuel prices. They show that the effect of fuel prices on cars' market shares was stronger before 2008 than after that year, though they find little evidence that this was due to more stringent fuel economy standards. Klier and Linn (2010, 2012, 2013, 2016) show that both standards and fuel taxes are associated with

technology adoption in the US and Europe. Our contribution to this literature is that we show the impact of these policies on inventions, in particular on patents for new future models that should guarantee compliance with the regulations such as car models based on breakthrough technologies.

Second, our work contributes to the literature on directed technical change (DTC) and the environment. The theoretical DTC literature poses that relative prices, market sizes and knowledge stocks determine the direction of innovation, and that environmental regulation can steer technical change to a cleaner path (e.g., Smulders and de Nooij 2003; Gerlagh 2008; Acemoglu et al. 2012; Acemoglu et al. 2016; Fried 2018; Hart 2019). The empirical DTC literature tests those predictions. Lanjouw and Mody (1996) are the first to study the relation between environmental policy and innovation using patent counts, which has since become the dominant way of measuring inventions. Popp (2002) is the first to estimate the effect of energy prices on clean innovation, in particular inventions that increase energy efficiency technologies. Early studies use country or industry level data, whereas later studies, such as Noailly and Smeets (2015), Aghion et al. (2016), Calel and Dechezleprêtre (2016) and Calel (2020) analyze patenting at the firm level. Most empirical studies find that energy prices or environmental policies have a significant effect on the direction of innovation, but little attention has been paid to the role of breakthrough technologies. The literature is reviewed by Popp et al. (2010) and Popp (2019).

Though a large literature studies the DTC mechanisms, work on the particular impact of regulation is limited. Some important exceptions include Johnstone et al. (2010), who analyze a variety of instruments in the energy sector, Dekker et al. (2012), who use difference-in-differences analyses at the country level to study the impact of air quality regulation on both inventions and the diffusion of inventions of emission abatement technologies, and Calel and Dechezleprêtre (2016) and Calel (2020), who use difference-in-differences analyses at the firm level to study the innovation effects of the EU Emissions Trading System. Most papers that test the effects of price instruments do not distinguish between (tax-exclusive) prices and taxes (e.g., Popp 2002; Noailly and Smeets 2015) and only a few papers have studied the innovation effects of standards, and fuel economy or GHG emission standards in particular. Vollebergh (2010) and Lee et al. (2011) find that air quality standards for automobiles have

significantly increased patenting in emission-saving technologies, while Crabb and Johnson (2010) and Vollebergh (2010) do not find the same effect for fuel economy standards. Noailly (2012) finds that building standards impact innovations for energy efficiency, whereas energy prices do not.

From this small set of papers it is clear that identification of a potential impact of regulation on innovation requires a precise link between the target of the standard, usually a specified restriction in physical space (such as miles per gallon or grams of CO₂ per kilometer driven), and inventions that contribute to the relaxation of this restriction. Moreover, it is important to identify to what extent the standard induces the need for research on breakthrough technologies as producers are usually able to exploit many substitution mechanisms to comply with new regulations (OECD, 2010). For instance, when confronted with fuel economy regulation car producers could stimulate consumers to buy vehicles with better fuel economy. In addition, research to increase fuel efficiency of existing (dirty) combustion technologies is likely to be a cheaper and a lower risk option than research on breakthrough technologies. This explains why it is important to measure the potential impact of a standard by taking its restrictiveness into account, which, in turn, depends both on the time frame announced for adaptation and its likelihood to be binding. This is precisely what our measure of the stringency of fuel economy and GHG emission standards across countries and time does, and why we find that these standards have a significant impact on clean innovation.

This paper continues as follows. Section 2 elaborates on climate regulation in the car market, discussing the innovation incentives that are provided by different policy instruments. This section also explains the importance of carefully identifying regulatory stringency and introduces our measure of standard stringency. Section 3 discusses our data. Section 4 explains our empirical strategy. Section 5 shows our results and section 6 shows our robustness checks. Section 7 concludes.

2 Regulation in the car market

2.1 Background

Road transport is heavily regulated because driving contributes to several externalities, such as local air pollution, climate change, congestion and traffic accidents (Parry et al., 2007).³ Our focus is on technologies that contribute to mitigating the climate externality of car use, so we are interested in the two types of policy instruments that address GHG emissions of car use: excise taxes on car fuels and standards that regulate either fuel economy or GHG emissions per unit of distance (Proost and Van Dender, 2011). Though taxes are generally preferred by economists from a welfare perspective (e.g., Anderson and Sallee 2016), they are often met with more resistance by the public, compared to standards. For instance, the French yellow vests protests of 2018 started as a protest against a plan to increase fuel taxes.

Most countries apply both an excise and a value added tax (VAT) to car fuels.⁴ Although an excise tax could also be levied ad valorem (e.g., Kanbur and Keen 1993), excises in transport are usually of the specific type and imposed per liter (or gallon) of fuel. This means that the excise does not vary with the tax-exclusive fuel price. A specific tax on fuel can be seen as a first best Pigouvian instrument to correct for the market failure caused by damages of climate change as a result of greenhouse gas emissions from car use.⁵ Instead, a VAT is specified as a percentage over the sum of the tax-exclusive price and the (specific) excise, and therefore does vary with the fuel price. Car fuel is typically taxed at the same VAT rate as other goods.

Fuel taxes differ considerably across countries (see also section 3). The US taxes fuels at a much lower rate than most European countries or Japan, for instance. The fuel tax rates have been relatively stable in real terms over the past two decades. They are typically either increased every few years, though some countries index the excise to the general price level. Some countries tax gasoline and diesel at the same rate, whereas others apply different rates. For instance, diesel is taxed more heavily than gasoline in the US, whereas it is the other way

³Additional reasons to tax transportation include road maintenance, oil dependency and raising revenue.

⁴An exception is the US, which does not impose VATs but instead uses sales taxes. Excise taxes are also applied, both at the federal and at the state level.

⁵CO₂ emissions are linked one-to-one to the carbon content of the fuel used for combustion in cars (Parry et al., 2007).

around in the Netherlands. There are several reasons to tax gasoline and diesel differently. One important reason is that tax competition induces governments to tax trucks, which use diesel and drive long distances, at lower rates than passenger vehicles, which are more likely to use gasoline and stay within their own tax jurisdiction. An important reason for trucks to use diesel is that a liter of diesel contains more energy than a liter of gasoline, meaning that vehicles can drive a longer distance with one liter of fuel (better fuel economy). Though diesel contains more CO₂ per liter than gasoline, its fuel economy implies that emissions per kilometer are lower than for gasoline. This is also why diesel fueled passenger cars are interesting from a climate policy perspective.

From a regulatory perspective the choice of the tax rate of the specific tax matters a lot. In particular, diesel contains more CO₂ per liter than gasoline, but its higher energy content implies that emissions per kilometer driven will be lower than for a gasoline car. So the fact that tax rate for diesel per liter is typically lower than for gasoline is in line with policies that aim to reduce CO₂ emissions per kilometer driven, whereas an equal rate per unit of carbon would not take the difference in fuel economy into account. A higher tax on diesel could be justified by concerns over air quality, for which diesel is more damaging than gasoline (Parry et al., 2007). In general, higher rates always imply higher stringency as these always increase the cost per unit and typically punish less (carbon) efficient cars relative to the more efficient ones.

The other instrument of interest is standards. Although many standards play a role in the context of the automotive industry (Vollebergh and Van der Werf, 2014), the policy standards that are typically relevant for climate change regulation of car use are those on fuel economy or on GHG emissions of cars and trucks. First, fuel economy standards, i.e. standards that restrict kilometers (miles) driven per liter (gallon) of fuel, have existed for a long time and aim to reduce the amount of fuel per kilometer driven. Higher fuel efficiency typically also reduces GHG emissions as it means less fuel use per kilometer driven. Because GHG emissions are directly related to fossil fuel use, these standards also guarantee lower emissions per kilometer driven.

Second, standards on GHG emissions, i.e. standards that specify CO₂ emissions per kilometer driven, regulate emissions more directly. Because a specific abatement technology

for GHG emissions from fossil fuel based combustion in cars is not available, however, such a standard also implicitly regulates fossil fuel driven cars and (depending on its coverage) trucks. When applied to a specific car, a fuel economy (km/liter) and a GHG (grams of CO₂/km) standard are essentially equivalent as kilometers per liter (fuel economy) and CO₂ emissions per kilometer are inversely related through a fixed amount of CO₂ per liter of fuel Anderson et al. (2011). However, the standards for automobiles are usually much more involved and differ along multiple dimensions. Standards typically set a minimum (maximum) on a car manufacturer's sales-weighted average fuel economy (GHG emissions). Hence, such a standard only has an indirect impact on fuel demand of a particular model as the overall supply of different types of models is affected. In particular, not all vehicles have to satisfy the given standard level as long as the manufacturer's weighted average does. Such a standard provides car manufacturers with the freedom to respond in different ways.

Table 1 summarizes the fuel economy and GHG emission standards for the countries that have (had) mandatory standards and that are included in our analysis.⁶ It is important to note that these standards are typically introduced as targets that car producers should meet in a specified year in the future. They can be voluntary or mandatory. Our analysis mainly focuses on the EU, the US and Japan, which are the three leaders in the automotive industry, both in terms of R&D and regulation. The US has the longest history of regulating fuel economy, introducing its first Corporate Average Fuel Economy (CAFE) standards for passenger vehicles and light trucks in the 1970s. Japan's Top Runner program, introduced in 1999, sets targets based on the current top performer in the industry. In recent years the EU has been most stringent in its GHG emission regulations. Canada, Korea and Mexico also have mandatory standards.

From table 1 it is clear that regulatory standards on fuel economy and GHG emissions indeed differ in many important respects. First, the choice of the regulatory base of a standard matters a lot and contains more degrees of freedom compared to a specific tax. The traditional

⁶Australia, New Zealand and Turkey do not have mandatory targets for GHG emissions or fuel economy. Switzerland follows EU standards. The UK was still in the EU for the entire sample period. The Czech Republic, Hungary, Poland and Slovakia entered the EU in 2004, i.e. before the European standards became binding, and did not have any binding restrictions before accession. The other countries that we cover are all in the table (or have been in the EU since before 2000).

EU	
1998/1999	Voluntary agreements reached with European, Japanese and Korean carmaker associations. Target of 140 grams of CO ₂ per km (NEDC) for 2008, with intermediate targets for 2003/2004 of 165-175 grams per km. Only two firms actually achieved the 2008 target.
2009	First mandatory target set at 130 g/km (NEDC) for 2015. Phase-in from 2012.
2013	Target for 2021: 95 g/km (NEDC), based on vehicle weight.
2019	Targets set for 2025 (81 g/km NEDC) and 2030 (59 g/km NEDC), representing a 15% and 37.5% reduction from the 2021 target, respectively.
Japan	
1979	First fuel economy standards set for 1985.
1993	Standards set for 2000
1999	Top runner program introduced, targets set for 2010. New targets are based on the performance of the current top performer. Targets vary by fuel type and vehicle weight. Overall fuel economy expected to be 15.1 km/l (JC08) if all firms meet target.
2007	Targets set for 2015: overall target 16.8 km/l (JC08).
2011	Targets set for 2020: overall target 20.3 km/l (JC08).
2019	Targets set for 2030: overall target 25.4 km/l (WLTP).
US	
1978-2010	Corporate average fuel economy (CAFE) standards. Tightened until 1985, then slightly loosened and tightened before stagnating from 1990.
2006	CAFE reform adopted. Standards based on vehicle size (footprint), transition period 2008-2010 with option to comply with the old standards, new standards mandatory from 2011. Additional change: larger passenger vehicles now also included (SUVs, small passenger vans).
2010	EPA now also responsible for standards (together with NHTSA). Standards set in grams of CO ₂ per mile, still measured using CAFE test cycle. Target set at 250 g/m for 2016.
2012	Target set at 163 g/m (CAFE) for 2025.
Canada	
1977-2007	CAFC targets harmonized with US standards, but voluntary. In 1982 the MVFCSA was passed to make standards mandatory, but it was not implemented until 2007.
2007-2010	CAFC targets remained the same, but became mandatory: 8.6 liter per 100 km (CAFE).
2010	Targets set in grams of CO ₂ per km. New target for 2016: 135 g/km (CAFE).
2014	New target for 2025 (aligned with US): 98 g/km (CAFE).
Korea	
2005	First standards announced for 2006-2011. Targets set at 12.1 km/l (CAFE) for vehicles up to 1500 cc and at 9.4 for other vehicles.
2009	Standards set for 2015 (phase-in from 2012). All vehicles are required to meet the target, which is set both in terms of fuel economy and emissions: 16.7 km/l or 140 g/km (both CAFE).
2014	Targets announced for 2020: 24.1 km/l or 97 g/km (both CAFE).
Mexico	
2013	First fuel economy standards: 16.7 km/l (CAFE) in 2016. No new regulations implemented, but standards for 2016 carried over to 2017 and 2018.

Table 1: Changes in fuel economy and GHG emission standards for selected countries. CAFE, JC08, NEDC and WLTP are test cycles (more in section 3). Sources: TransportPolicy.net, ICCT and national legislature.

standards from the 1970s just regulated the sales-weighted average fuel economy of cars with a maximum weight per manufacturer. More recently, however, the table illustrates that most countries switched to attribute-based standards which also take vehicle attributes, other than fuel economy or GHG emissions, into account. For instance, the US implemented standards based on the footprint of cars (the vehicle’s wheelbase multiplied by its track width) in 2006, whereas standards in the EU and Japan are based on vehicle weight. This means that larger or heavier vehicles are allowed to have higher emissions per kilometer than smaller or lighter ones.⁷

Second, another important dimension of standard setting is (like the specific tax) its selectivity of coverage. In order to be implemented in practice, the regulator has to specify carefully which cars or trucks are or are not covered by the standard. For instance, the CAFE standards in the 1970s in the US were restricted by using a maximum car weight. Recently also so-called multipliers have been added. Multipliers allow manufacturers to weight particular cars (e.g., electric vehicles) more heavily when computing their average fuel economy or GHG emissions (Gillingham, 2021).

Third, the strictness of a particular standard matters. Strictness of a standard typically depends on the target value chosen in physical terms relative to existing performance, whether and how compliance is enforced or not, and the length of the phase-in period. Clearly, strictness firstly relates to distance to target in the regulated dimension. Furthermore, standards can simply be voluntary and do not provide any financial incentive for compliance, or they could be combined with a fine for firms that do not comply. In that case firms can choose between compliance and paying the fine. For instance, Jacobsen (2013) shows that some manufacturers chose not to comply with the American CAFE standards in the 1990s and consistently paid fines.⁸ In addition, supercredits can be used to relax stringency

⁷Japan has brackets of vehicle weight and a level of fuel economy for each bracket. The EU and the US work with a linear formula. The higher/larger a manufacturer’s average vehicle weight/footprint, the more CO₂ its (sales-weighted) average vehicle is allowed to emit per kilometer.

⁸Jacobsen (2013) distinguishes between three groups of firms: those for which the standard is nonbinding (Asian carmakers and Volkswagen), those that just comply with the standard (American carmakers), and those that consistently violate the standard and pay the fine (BMW, Mercedes, Porsche). However, fines have increased in recent years, even causing high-emission car manufacturer Fiat Chrysler Automobiles to pool its fleet with electric car producer Tesla in exchange for ”hundreds of millions of euros” to comply with EU

by providing firms with the opportunity to compensate under-performance in one year with over-performance in a different year. Finally, and very important for the stringency of a given standard, is the length of the phase-in period. Standards are typically expressed as targets in the future and announced years in advance to give firms the opportunity to adjust. Phase-in periods define the years before the target becomes enforced. For instance, in 2013 the EU announced its target of 95 grams of CO₂ per kilometer to become fully enforced in 2021, which allows firms to adjust during 7 years.⁹

To illustrate the relevance of the design features of the car standards for their stringency, figure 1 shows standards and performance for the EU, Japan and the US in the same dimension, i.e. in grams of CO₂ per kilometer measured using the NEDC test cycle (we elaborate more on test cycles in section 3). The coloured line in figure 1 is the most recently announced target. So the level of the target changes when a new target is announced, not when it becomes enforced (starting date of the announced target). The black line shows average performance, measured at the same scale as the target. To determine the stringency of a particular announced target or standard, we use two criteria: i) whether the announced target is voluntary or not; ii) if not, whether the standard is below the contemporaneous average performance level, i.e. to what extent is it really binding. The coloured line is blue and marked with diamonds if non-binding, red and marked with triangles if binding. For example, the EU had only voluntary targets until 2009, when it announced its first mandatory standards (supported by fines), which were subsequently tightened in 2013 and 2019. Japan has had mandatory standards in place for the whole period, but these were not binding for the average firm between 2013 and 2018. Standards in the US were mandatory but not binding for the average firm until 2010, when more stringent targets were announced. Figure 5 in the appendix shows a similar graph for Canada, Korea and Mexico.

standards and avoid fines (Financial Times, 2019).

⁹Another aspect of the phase-in period is that in the years between announcement and enforcement of the target, some intermediate targets become binding. For example, in 2020 the target of 95 grams per kilometer was mandatory, but only for 95% of the fleet (i.e., firms were allowed to drop the dirtiest 5% of their fleet when computing their sales-weighted average emissions).

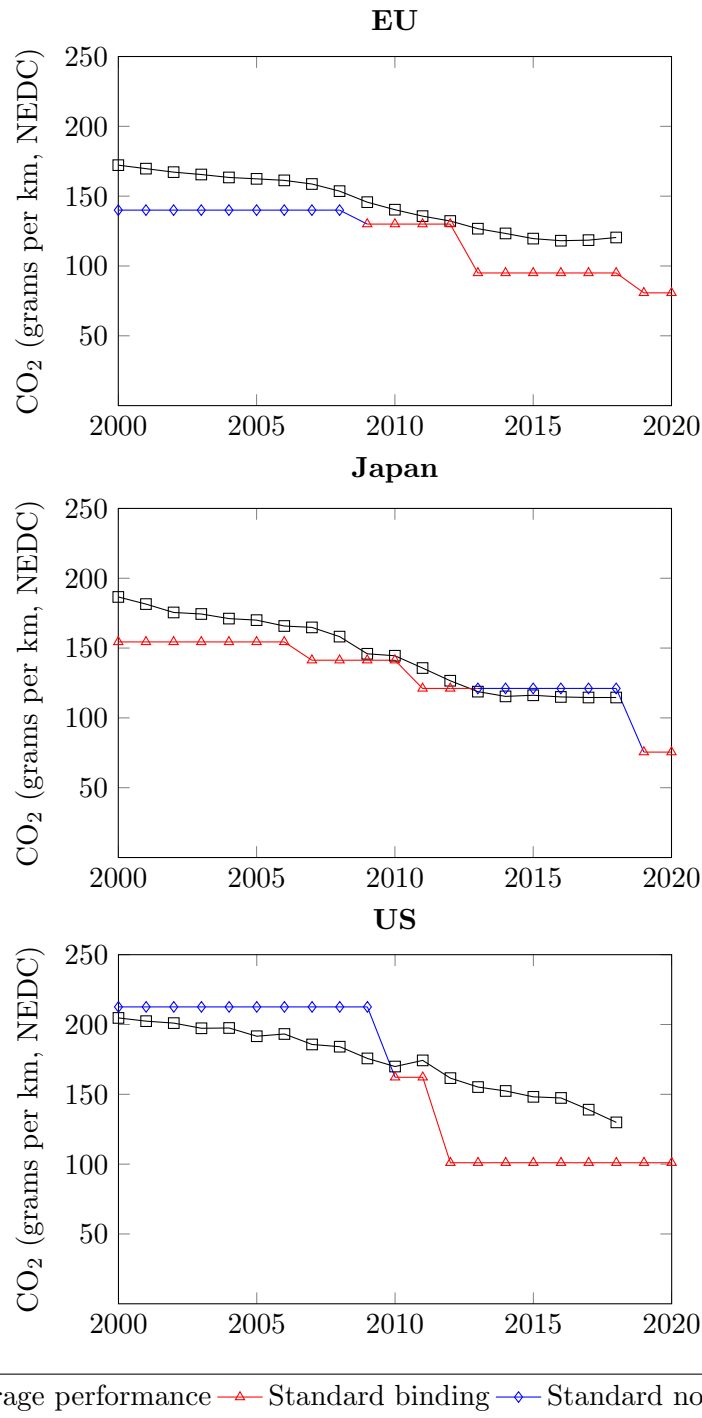


Figure 1: Average performance and standards converted to CO₂ emissions per kilometer (NEDC). Standards are defined as the most recently announced target level. They are binding if the target is mandatory and lower than average performance. Source: ICCT.

2.2 Innovation effects

In theory both taxes and standards can provide an incentive for car manufacturers to improve their vehicles' climate performance. Fuel prices, and taxes in particular, have been shown to impact consumer behavior. It is well-known that prices and taxes influence the marginal decision to use a car (e.g., Espey 1998),¹⁰ but it has been shown more recently that fuel prices and taxes may have different impacts. Davis and Kilian (2011) and Li et al. (2014) find that taxes are more effective at reducing fuel consumption than prices. Indeed, also the demand for vehicles is sensitive to fuel prices and taxes. For instance, Busse et al. (2013) and Allcott and Wozny (2014) show that fuel prices affect the demand for vehicles, depending on their fuel economy.¹¹ Possible explanations for the difference between fuel prices and taxes are that tax changes are perceived as persistent and most likely upward, whereas price changes may be temporary and less predictable. Additionally, tax changes may be more salient than tax-exclusive fuel price changes.

A shift in demand towards vehicles with better fuel economy incentivizes innovators to redirect their attention towards less dirty technologies. Indeed, Aghion et al. (2016) apply the DTC model to the automotive sector to show that researchers shift their efforts from dirty technologies (the traditional combustion engine) to gray (improved fuel efficiency) and even clean ones (hybrid, electric, hydrogen) in response to higher tax-inclusive fuel prices, as long as dirty, gray and clean vehicles are substitutes. Tax incentives in the consumer market that directly stimulate penetration of clean vehicles, e.g. by providing very low or even zero tax brakes for zero-emission cars, are likely to play a similar role here (see also D'Haultfœuille et al. 2014). If demand for clean vehicles increases, it becomes more profitable for researchers to develop clean technologies.

A similar argument can be made for standards. Car manufacturers have a variety of

¹⁰Knittel and Tanaka (2021) show that fuel prices also affect driving behavior. They estimate that 19% of the price elasticity of fuel consumption is due to improvements in fuel-conserving driving.

¹¹Allcott and Wozny (2014) find that consumers value discounted future fuel expenses only 76% as much as they do the vehicle purchase price and attribute this to bounded rationality. Busse et al. (2013) do not find evidence of such myopic behavior. Gillingham et al. (2021) use evidence from a natural experiment to show that consumers undervalue fuel economy. Leard et al. (2021) also find evidence of undervaluation of fuel economy by consumers.

margins along which they can adapt to stricter GHG emission regulations (e.g., Gillingham 2021; Reynaert 2021). How firms respond precisely depends on the design of the standard. If a fuel economy standard applies to a firm’s sales-weighted average of emissions, a natural solution for the firm would be to adjust its vehicle mix. For instance, the firm could shift from selling high-emission to more low-emission vehicles.¹² It is well established that car manufacturers have some degree of market power and thus some price-setting ability (e.g., Berry et al. 1995). In this way the standard works as an implicit tax on dirty vehicles (Bento et al., 2020). Without changing their vehicle mix, firms have also been shown to trade off vehicle characteristics (Knittel, 2011). For instance, they can reduce a particular model’s weight or horsepower to improve its fuel economy.

Alternatively, firms can adopt different car technologies to comply with more stringent regulations. Klier and Linn (2016) use vehicle-level data and estimate technology frontiers to show that observed changes in vehicle characteristics cannot be explained by trading off vehicle characteristics alone. They interpret this result as evidence that adoption of emission-saving technologies is another important response to more stringent regulation. Such technologies should be available however, and require car manufacturers to have invested in these new, cleaner technologies earlier in order to accommodate additional demand into such clean technologies.¹³

So both taxes and standards provide car manufacturers with an incentive to improve the climate performance of their fleet, which, in turn, would induce researchers to also switch away from dirty to gray and even clean innovation. However, the incentives between the instruments could be substantially different. Since the effect of a specific tax goes through the demand side of the vehicle market, the elasticity of demand for climate performance with

¹²It can do so by switching from gasoline vehicles to diesel vehicles, which have better fuel economy, though diesel vehicles face stricter air quality regulation in most countries. It could also shift from large, heavy, high-emission vehicles to smaller and lighter vehicles. The incentive to make cars smaller and lighter is reduced, however, if attribute-based standards are in place. Using bunching analysis, Ito and Sallee (2018) even show that Japanese firms increased vehicle weight in response to weight-based standards. This, of course, harms overall fuel economy.

¹³In addition to mix-shifting, changing vehicle characteristics and technology adoption, Reynaert and Sallee (2021) show that car manufacturers in the EU game policies by making their vehicles perform better in the lab where they are tested than they do on the road.

respect to fuel prices is the key determinant. Fuel economy improvements through retrofitting dirty technologies are likely to be the cheapest and most important option here.¹⁴ This solution is likely to become less important if demand side (tax) policies explicitly start to stimulate clean technologies such as electric or hydrogen cars in niche markets, for instance with reduced or even zero rates for specific technologies such as electric cars. To comply with standards, however, firms themselves face a trade-off between the potential responses outlined above. Car manufacturers are likely to weigh the costs of mix-shifting, trading off vehicle characteristics and investing in new research and its direction.

Furthermore, the innovation inducement effect is determined by the strength of the policy impulse, i.e. the stringency of the regulation. For instance, small changes in fuel tax rates or a weak or voluntary standard are likely to induce a very different response compared to higher tax rates or more strict standards on the climate performance of the overall car park of a manufacturing company. Innovation is expensive and will only take place if more research is likely to be profitable (or if it is directly subsidized, of course). The stronger the policy impulse, the more likely it is that breakthrough technologies will become part of firms' cost-efficient response and radical research in this area will pay off. For taxes this means that the more expensive it becomes to drive a dirty vehicle, the more likely consumers are to switch to a gray or even clean one. For a standard the incentive to do more research on new technologies depends on the manufacturer's needs to improve its (future) performance in the face of a stricter standard.

Finally, innovation is a dynamic process that takes time. Hence, expectations about future policies or announcements of targets for future years matter for current decisions about innovation investments. One would expect the fundamental characteristics of fuel prices as well as both instruments to play an important role here as well. One such characteristic is that fuel prices can be (very) volatile, while both taxes and standards are much more predictable because they are rarely decreased. The stronger effect of taxes relative to fuel prices due to the persistence and salience arguments mentioned above illustrate this (Davis and Kilian 2011; Li et al. 2014).¹⁵ For standards the announcement date of a new target is

¹⁴Up to some (thermodynamic) point.

¹⁵These papers investigate the differential effect of taxes and prices on driving. The same difference may be present for consumers' choice to buy a vehicle, though we are not aware of any studies that investigate this

relevant because this determines a particular date in the future for the firms to comply. As mentioned above, a firm’s response to this announcement might indeed involve investment in more research on alternative technologies.¹⁶ The likelihood that firms will invest more into research is likely to be higher if compliance will be difficult without selling a larger share of models using breakthrough technologies in the future.

So both fuel taxes and climate standards are expected to positively affect clean innovation and dirty innovation negatively, while the predicted effect on gray innovation is ambiguous. Several factors, such as stringency and timing, determine the strength of these effects and can lead to differences between the effects of the two policies. Determining which of the policy instruments has a stronger impact would require a detailed modelling of the firm’s decision making, taking into account the many margins along which firms can respond to both policies. This is beyond the scope of our paper.

2.3 Stringency of fuel taxes and standards

The subsection above discusses the importance of properly measuring regulatory stringency, both for taxes and for standards, when estimating the effect of policies on innovation. Instruments often vary along multiple dimensions, which are difficult to compare. Indeed, Brunel and Levinson (2013, 2016) identify this multidimensionality of instruments as one of the primary obstacles to measuring stringency.¹⁷ The challenge is to capture those features that really determine regulatory stringency in a one-dimensional measure or to control for different dimensions separately.¹⁸

link.

¹⁶Dekker et al. (2012) show the relevance of such expectations for innovators in the context of regulation by air quality standards on coal fired power plants.

¹⁷Brunel and Levinson (2013, 2016) also discuss simultaneity, industrial composition and capital vintage as obstacles. Simultaneity refers to the issue that the consequences of stringent regulation may also be its determinants (highly polluted cities may be more likely to introduce stringent pollution regulation). Industrial composition is a particular example of simultaneity. The capital vintage argument poses that many environmental regulations feature different rules for older sources of pollution than for new ones.

¹⁸Brunel and Levinson (2013, 2016) discuss several potential solutions, namely composite indices (e.g., Johnstone et al. 2012), pollution abatement expenditures (e.g., Brunnermeier and Cohen 2003), and shadow prices (e.g., Van Soest et al. 2006). All these approaches have disadvantages or are infeasible for us. We use the approach they label as direct assessment of regulation. Importantly, the broader the scope of a measure,

We carefully measure the stringency of both instruments to test the prediction that the instruments induce clean innovation. Stringency measures should be comparable across countries and capture the instruments' key determinants that matter for the innovation incentive. Our measures indeed focus on the three key design elements that matter for innovation, as discussed in the previous subsections: the instrument's tax or regulatory base, its level of stringency in the regulated dimension, and its timing. Fuel taxes can be relatively easily compared in these three dimensions. First, fuel taxes have the same tax base over time, i.e. fuel use of either gasoline or diesel which is directly linked to the regulated dimension (either fuel efficiency or GHG emissions). Second, its stringency is defined by the level of the tax rate which can easily be converted to one currency and adjusted for differences in purchasing power between countries and over time. Third, tax changes happen over a relatively short time frame (typically about one year between announcement and implementation). They depend on the political process and are usually persistent and bounded from below, as discussed before.

For climate standards, however, this comparison is much more involved. We develop a simple and intuitive measure of the stringency of fuel economy and GHG emission standards. In order to do so we first convert all standards to the same units and test cycle, namely grams of CO₂ per kilometer measured using the New European Driving Cycle (NEDC). We elaborate more on this process in the next section. We then decompose the target into its main features: its level relative to current performance, whether or not it is binding, and its time horizon (see also figure 1). Combining these features we compute stringency for each country and year as the yearly reduction in CO₂ per kilometer that the average car manufacturer in a country needs to achieve to comply with the most recently announced target. It is computed as follows for country c in year t .

$$Stringency_{ct} = Binding_{ct} \frac{Actual_{ct} - Target_{ct}}{TargetYear_{ct} - t}, \quad (1)$$

where *Actual* is the current actual performance of the average car manufacturer, *Target* is the level of the most recently announced target, *TargetYear* is the year in which firms need to comply with the target, and *Binding* is an indicator that is one if the target is mandatory.

the more difficult it is to keep comparability. That is, comparing countries' overall environmental stringency is harder than comparing the stringency of particular regulations.




Figure 2: Stringency for the EU, Japan and the US defined as the average required annual reduction in CO₂ emissions per kilometer (NEDC) to meet the most recently announced target, computed following equation (1). Source: ICCT.

and more stringent than current performance and zero otherwise. Let us illustrate with an example. In 2009 the EU announced its first mandatory target: 130 grams per kilometer in 2015. Average emissions in 2009 were 145.7 grams per kilometer, so the standard was binding. The required annual decrease is thus 2.62 grams per year (15.7 grams in 6 years).

Figure 2 shows our stringency indicator for the EU, Japan and the US, which are the dominant jurisdictions both in terms of environmental policy and in terms of innovation. Figure 6 in the appendix shows a similar graph for Canada, Korea and Mexico. The difference between the EU and the US on the one hand and Japan on the other hand is striking. Stringency in Japan was relatively stable and positive until 2008, while standards in the EU and the US were not yet binding. Japan’s stringency then declined due to improved performance and a lack of ambitious targets, whereas the EU and US started introducing binding targets and gradually increasing their stringency. Meanwhile Canada largely followed the US, Mexico had low stringency even in the brief period in which its regulation was binding, and Korea’s stringency drastically increased in recent years as the 2020 target was approaching but performance lagged behind.

To identify the impact on a dynamic variable like innovation requires accounting for anticipation effects as discussed in the previous subsection. Defining stringency as the required

annual improvement does exactly this. We start counting from the year in which a policy is announced, which is when car manufacturers receive the impulse to start innovating. If performance does not improve, the need for a breakthrough innovation becomes stronger as the target year approaches.

3 Data

3.1 Patents

This paper uses patent counts as a measure of innovation. Patent counts is the most widely used measure of innovation in the empirical economic literature despite its drawbacks, which include the fact that not all innovations are patented and that the values of patents are highly heterogeneous. Furthermore, patents measure the outcome of the innovation process, rather than the inputs, such as R&D investment and number of researchers. A patent represents an invention and does not indicate much about technology diffusion and adoption. Despite these drawbacks, patents are often used, as there are no better alternatives. R&D investment, for instance, is the closest substitute but it is often not available at the firm level and cannot easily be classified by technology. Concerns about firms' propensity to patent and patent values will be addressed by including firm fixed effects and using only a select group of patent families.

Patent data has been collected from the October 2020 version of the Patstat database, which is maintained by the European Patent Office (EPO).¹⁹ Patstat includes over 100 million patent records from 90 patent issuing authorities around the world. The oldest patents go back to the 19th century. This database provides a wealth of information about each invention's technologies, applicants and inventors. The patent data is used to establish a panel of counts by firm and year for each category (dirty, gray, clean). The sample includes 3646 distinct patent holders and the years 2000 until 2016. We chose 2000 as the starting year to have enough pre-sample data available to create weights and to control for fixed effects and because fuel economy performance data, which we need to compute our standard stringency measure, is only available from that year onward. We selected 2016 as the end year because

¹⁹For more information about the Patstat database see <https://www.epo.org/searching-for-patents/business/patstat.html>.

this is the last year for which the October 2020 version of Patstat is complete. This is due to the delay between patent applications, grant decisions and updating of the database (Aghion et al., 2016).

Selected patents are classified as dirty, gray or clean based on their International Patent Classification (IPC) code, which categorizes patents by the type of technology they protect. In selecting patents we used the same criteria as Aghion et al. (2016), who base their selection on work by the OECD (e.g., Vollebergh 2010) and conversations with patent experts. Dirty inventions relate to the internal combustion engine. A subgroup of these dirty patents are gray, meaning that they aim to improve the fuel efficiency of the internal combustion engine in particular (making dirty less dirty). The other subgroup is classified as purely dirty. Clean inventions are those related to hybrid, electric and hydrogen vehicles and fuel cells. We classify hybrid vehicle technologies as clean to keep consistency with the classification used by Aghion et al. (2016), though hybrid vehicles use fossil fuels and could therefore also be classified as gray. Electric and hydrogen powered vehicles are purely clean, as long as renewable energy is used to generate their power source. We also refer to these as breakthrough technologies.²⁰ Table 2 shows the IPC codes belonging to each category and table 3 shows some examples of the technologies.

Patents are counted at the family level to prevent double counting of inventions. A patent family includes all patent applications that cover the same invention. A family may include multiple applications at the same patent office and it may contain applications for the same invention in multiple countries. Each family has at least one priority patent, which is the first application of a certain invention. There are multiple ways of defining a patent family (Martinez, 2010). We use DOCDB families, which are provided by Patstat. These families are constructed by patent examiners and may have multiple priorities. We count each family in the year of the earliest application within the family, as is standard in the literature. According to Griliches (1990), firms file a patent application early in the invention process, which means that the priority date is close to the date of invention.

A well-known issue with patent data is that the value distribution of patents is skewed

²⁰We use the term breakthrough because these technologies allow vehicles to drive without emitting any GHG emissions. We do not classify individual patents or patent families as breakthrough or radical as is done by, for instance, Acemoglu et al. (2021).

Technology	IPC codes
Dirty patents	
Internal combustion engine	F02B, F02D, F02F, F02M, F02N, F02P
Gray patents	
Fuel efficiency of internal combustion engines	F02M 39-71, F02M 3/02-05, F02M 23, F02M 25, F02D 41, F02B 47/06
Clean patents	
Hybrid vehicles	B60K 6, B60W 20, B60L 7/1, B60L 7/20
Electric vehicles	B60L 11, B60L 3, B60L 15, B60K 1, B60W 10/08, B60W 10/24, B60W 10/26
Hydrogen vehicles	B60W 10/28, B60L 11/18
Fuel cells	H01M 8

Table 2: Technologies and IPC codes of the patents used in this paper. In selecting the IPC codes and classifying them as clean, dirty and gray we follow Aghion et al. (2016). Note that gray patents are a subgroup of dirty patents. To find the exact definition of each code, see <https://www.wipo.int/classifications/ipc/ipcpub/>.

IPC code	Definition
Dirty	
F02B 1/08	Engines characterised by fuel-air mixture compression with separate admission of air and fuel into cylinder
Gray	
F02M 23/06	Apparatus for adding secondary air to fuel-air mixture dependent on engine speed
Clean	
B60L 1/04	Supplying electric power to auxiliary equipment of electrically-propelled vehicles fed by the power supply line
H01M 8/18	Fuel cells; manufacture thereof: regenerative fuel cells, e.g. redox flow batteries or secondary fuel cells

Table 3: Some examples of the technologies belonging to the different classes of patents used in this paper.

(e.g., Scherer 1965; Schankerman and Pakes 1986). There are many patents that have little economic value and the distribution has a long tail with some inventions that are highly valuable. In order to exclude inventions with little value we follow Aghion et al. (2016) and select only triadic patent families. These families include at least one application at the EPO, one at the Japan Patent Office (JPO) and one at the American USPTO. The idea here is that applying for a patent is costly, which means it is only worth the cost if the invention is likely to be profitable. Furthermore, this approach limits the sample to inventions with a potential for international application, because protecting an invention abroad is only worth the cost if there is a possibility to use or sell it internationally.

Figure 3 shows the total number of dirty and clean patents per year for the period 1978-2016. Dirty patents peak in 2010 and clean and total patents in 2011, which is broadly consistent with the findings of Probst et al. (2021) and with aggregate country level data from the OECD.²¹ The figure clearly shows that the number of patent applications falls after its peak for both clean and dirty patents. The graph also shows that yearly clean patents have overtaken dirty ones in 2009 and stayed higher since. The total number of patent families selected for the sample period 2000-2016 is 34,622, of which 16,627 are classified as clean and 19,392 are dirty.²² Almost half (9,223) of the dirty patents are aimed at increasing fuel efficiency and are thus classified as gray, and the rest (10,169) are purely dirty. About half of the clean families relate to hydrogen vehicles or fuel cells (8,497). The other half mostly exists of patents for electric vehicles (7,816). A total of 3,670 patents relate to hybrid technologies, most of which (2,923) are also classified as electric.

Most patent applications mention multiple applicants. We assigned a patent family to the organization or individual that is mentioned as applicant on the highest number of applications within that family. In case of a tie we assigned an equal fraction of the family to all those applicants that were tied. We then manually checked the list of applicants and matched those that were mentioned multiple times with slightly different names (e.g. Toyota Motor

²¹The aggregate data is available at <https://stats.oecd.org/> under “Science, technology and patents”. Total triadic patent families in the category “Selected environment-related technologies” peaked in 2011. The same is true for patent by Japanese and German applicants (those by American applicants peaked in 2012).

²²Some patent families contain IPC codes that fall into both categories and are thus counted both as clean and as dirty. This is the case for 1,397 families (about 4% of all families), most of which (1,098) are for hybrid car technologies.

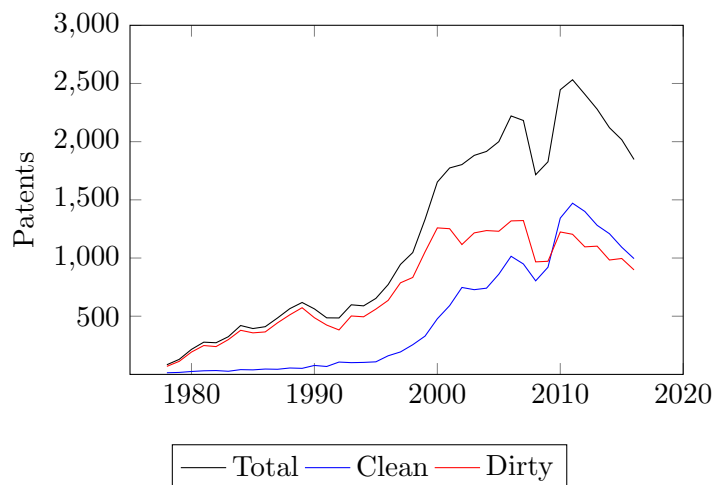


Figure 3: Number of triadic patent families per year. Total consists of all families used for this paper, but is not the sum of clean and dirty as some patent families fall into both categories. The sample period used is 2000-2016. Source: Patstat.

Corporation and Toyota Motor Europe). This reduced the number of different applicants from 4,279 to 3,646. The majority of applicants are companies (72%), and the remaining applicants are either individuals (20%), universities (4%), non-profit organizations (2%) or unknown (1%). All of the largest patent holders are companies. Interestingly, it is not just the car manufacturers themselves who innovate. In fact, though most car manufacturers do appear in our patent data, the majority of firms that apply for car-related patents are suppliers of components or research firms. For instance, two of the firms that hold most patents in our data set are Toyota, a car manufacturer that does much of its R&D in-house, and Bosch, a large supplier of vehicle components. As car manufacturers look for ways to improve fuel economy, they may decide to innovate or to buy new technologies from their suppliers.

3.2 Fuel taxes and prices

Data on fuel taxes and prices is obtained from the International Energy Agency's (IEA) Energy Prices and Taxes database.²³ This database contains annual energy prices for many countries and several energy sources, such as diesel, gasoline, LPG and electricity.

²³For more information about the IEA database see <https://www.iea.org/subscribe-to-data-services/prices-and-taxes>.

When selecting the database that is in each country's local currency it also distinguishes between excise taxes and value added taxes (VATs).²⁴ We take tax-exclusive and tax-inclusive fuel prices and the excise tax and convert them to 2015 US dollars using the OECD's purchasing power parities (PPP) conversion rates. For the purpose of this paper we are interested in car fuels, so for each country we take the (unweighted) average of the diesel price and the gasoline price.²⁵ We include fuel prices and taxes for 30 countries.²⁶

Figure 4 shows the evolution of excise taxes and tax-exclusive fuel prices for selected countries as used in our estimations. Excise taxes vary considerably across countries. The US has a low excise tax and European countries generally have high taxes. Japan is in between. Excises are relatively stable in the time dimension. This is not surprising, as many countries index their excise to the general price level and our measurement of the tax corrects for differences in purchasing power (both across countries and over time). Tax-exclusive fuel prices are much more volatile and follow a highly similar trend in all countries. A large part of the time series variation is due to (tax-exclusive) oil prices, which means that most countries experience the same shocks. Level differences between countries are due to variation in transportation costs and purchasing power.

Fuel taxes and prices are measured at the national level, whereas patents are measured at the firm level. We therefore follow Noailly and Smeets (2015) and Aghion et al. (2016)

²⁴The database provides separate excises and VATs for all countries in our data set except the US. We thus supplement this data set with data from the American Highway Statistics, published by the Federal Highway Administration (see <https://www.fhwa.dot.gov/policyinformation/statistics/2019/>). The US has federal excise taxes of 18.4 and 24.4 cents per gallon on gasoline and diesel, respectively. These have been at the same level since 1993. States also collect excise taxes, so we add the average state-level excise (weighted by volume taxed) to the federal one.

²⁵There are several types of gasoline (leaded, unleaded, regular, premium). We take premium unleaded 95 if it is available for a country in all years (2000-2016), and otherwise regular unleaded. Prices and taxes for at least one of those two fuels are available for all countries. Diesel prices are available for all countries in our data set.

²⁶The countries are Australia, Austria, Belgium, Canada, Chile, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Japan, Korea, Luxembourg, Mexico, the Netherlands, New Zealand, Norway, Poland, Portugal, the Slovak Republic, Spain, Sweden, Switzerland, Turkey, the United Kingdom and the United States.

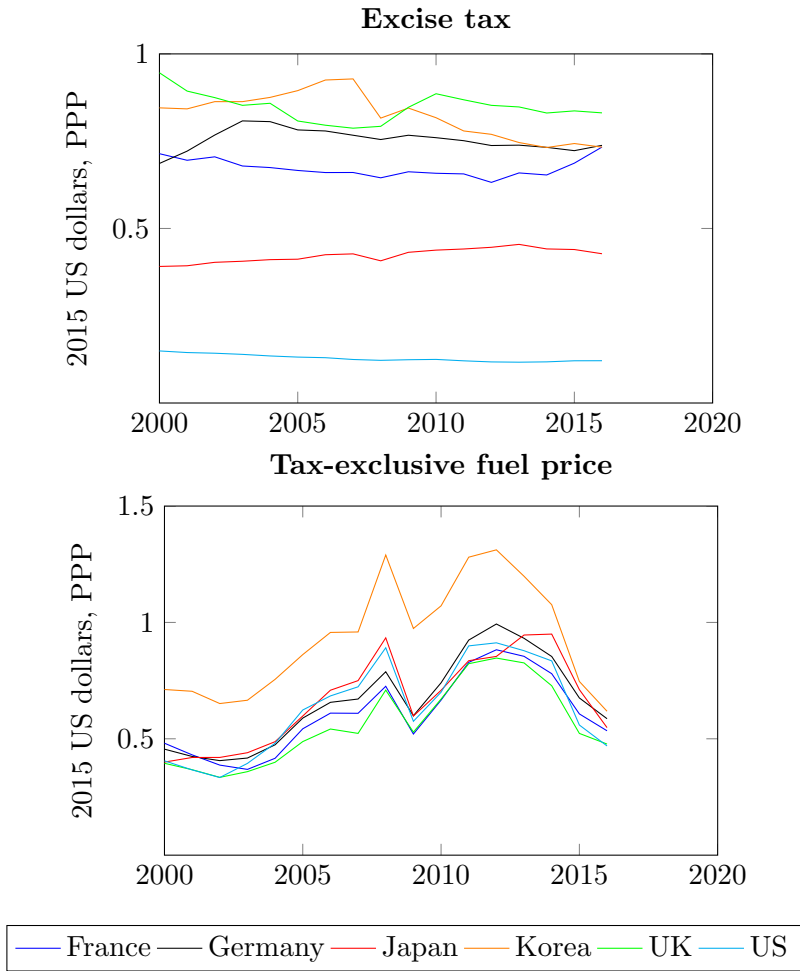


Figure 4: Excise tax and tax-exclusive fuel price in 2015 US dollars (PPP) for selected countries. Both are computed as the average of diesel and gasoline. Source: IEA.

and create firm level excise taxes as follows (prices are computed in the same manner).

$$FuelTax_{it} = \sum_c w_{ci}^F FuelTax_{ct}, \quad (2)$$

where c indicates a country and w_{ic}^F is a weight that captures firm i 's exposure to country c 's market. These weights are computed as follows.

$$w_{ci}^F = \frac{w_{ci}^P GDP_c^{0.35}}{\sum_{c'} w_{c'i}^P GDP_{c'}^{0.35}}, \quad (3)$$

where w_{ci}^P is the share of firm i 's pre-sample patent applications that were filed in country c .²⁷ The pre-sample period runs from 1978 until 1999. We use all pre-sample patent applications by the firms in our data set to construct these weights, so not only those in triadic families concerning car technologies. This is to increase the number of firms for which pre-sample data is available. The idea behind these weights is that a firm will protect its inventions in those places where it believes it is going to sell its product. Hence, a firm's patent portfolio should be a reasonable proxy for its sales distribution by country and thus its exposure to these countries' fuel prices. This is a common method of measuring firms' exposure to country-level variables (e.g., Noailly and Smeets 2015; Aghion et al. 2016). We follow Dechezleprêtre et al. (2021) and further weight by pre-sample GDP to reflect market size.²⁸ Aghion et al. (2016) show that geographical sales weights and patent portfolio weights are highly correlated for some large automotive companies. This approach incorporates firms' home bias: the fact that their home market is often disproportionately important (Dernis et al., 2001). We use pre-sample portfolios to establish weights to assure their weak exogeneity. Almost half (1722) of the 3,646 firms in the sample do not have any pre-sample patents. These firms are given the average weights of the other firms.

3.3 Standards

Data on standards and their announcement dates is gathered from TransportPolicy.net, which is a website created by the International Council on Clean Transportation (ICCT) and Diesel-

²⁷Note that a patent application is filed at a national patent agency. A patent family may include patents that are filed at different agencies. These weights thus add up to 1.

²⁸The exponent 0.35 is added because patent weights already partially reflect market size (see Dechezleprêtre et al. 2021). Setting the exponent to 0 or to 1 does not change our results in a meaningful way. Pre-sample GDP is computed as the average over the years 1995-1999.

Net.²⁹ It contains an overview of regulations on transportation for many countries and includes the standards that we study in this paper. We restrict our attention to standards for passenger vehicles (PVs), rather than those for light commercial vehicles (LCVs) or heavy duty vehicles, because PV standards are more common and typically more stringent than LCV standards. In addition, new standards for LCVs are often announced at the same moment as PV standards and are therefore likely to be closely correlated to our proposed measure of stringency. PVs are also a larger source of CO₂ emissions.³⁰

Standards are measured in different units across countries (see also Table 1). The EU, for instance, sets standards on GHG emissions, measured in grams of CO₂ per kilometer, whereas the US initially used standards measured in miles per gallon of fuel, and now has standards on emissions in grams of CO₂ per mile. Japan has standards expressed in kilometers per liter. In addition, countries use different test cycles to determine a car’s fuel economy or emissions. The four most used ones are the American CAFE test cycle, the New European Drive Cycle (NEDC), the Japanese JC08 cycle and the Worldwide harmonized Light vehicles Test Cycles (WLTC). Cycles differ, among other things, in the amount of city driving relative to highway driving and the average speed during the test (An and Sauer, 2004). This means that even if the American standard (measured using the CAFE cycle) is converted from grams per mile to grams per kilometer, it is still not equivalent to the European standard (measured using NEDC). To circumvent this issue we use a conversion tool for standards, developed by the ICCT.³¹ This tool allows one to convert standards across test cycles and units. We use this

²⁹Some additional data was needed for Japan and was acquired from the website of the Energy Conservation Center Japan (eccj.or.jp). Japan’s targets vary by weight group, which means that they are not straightforward to aggregate. One would need to weight by the (expected) distribution of vehicle weights. However, the regulatory documents of the ECCJ mention the expected overall fuel economy if the standards are met. We circumvent the aggregation problem by using this overall expectation as the target.

³⁰For instance, 12% of total EU emissions comes from PVs versus 2.5% from LCVs. See <https://ec.europa.eu/clima/policies/transport/vehicles/regulation>.

³¹This tool is available on the ICCT website. See <https://theicct.org/chart-library-passenger-vehicle-fuel-economy> under “Sources and tools”. The conversion tool requires the fleet average diesel penetration (average of the two countries) as an input (Kühlwein et al., 2014). The share of diesels in new car registrations has been negligible in the US and Japan, whereas in the EU it has risen from around 30% to around 50% in the sample period (Díaz et al., 2017). Since we convert American and Japanese standards to the European NEDC, we use a diesel penetration of 20%. None of the results hinge on this parameter choice.

tool to convert all standards to grams per kilometer (NEDC). Canada, Korea and Mexico also have mandatory standards. For several reasons we have not included these countries in all specifications.³²

Converting targets into the same units and test cycle does not give us a reliable measure of the stringency of the standards, however. As described in section 2, we use data on average performance in terms of CO₂ emissions per kilometer to compute the required reduction to meet the most recently announced target. The average performance data is gathered from the ICCT, which collects the data from government agencies.³³ This data concerns the sales weighted average of newly sold cars and is converted by the ICCT to the NEDC test cycle. It is thus comparable with the targets discussed above.

Based on this data we compute our country specific stringency measures for the different standards at the country level. To translate these stringency measures into firm specific levels we take the same approach we use for fuel prices, i.e. using weights that reflect market exposure. However, the weights are somewhat different to those for taxes, as some countries are left out for standard stringency. Portfolio shares are computed over all pre-sample patents filed in EU countries, the US and Japan.³⁴ The standard stringency measure for each firm i is thus computed as follows.

$$Stringency_{it} = \sum_{r=1}^3 w_{ir}^S Stringency_{rt}, \quad (4)$$

where $r = 1, 2, 3$ represent the EU, Japan and the US and

$$w_{ir}^S = \frac{w_{ir}^P GDP_r^{0.35}}{\sum_{r'} w_{ir'}^P GDP_{r'}^{0.35}}. \quad (5)$$

³²None of these three countries has had mandatory standards for the entire sample period (2000-2016), meaning we cannot include the level of their target. Canada and Mexico have some missing data points in their performance data (2011-2015 missing for Canada and 2015 for Mexico). We use linear interpolation to fill in these values and compute stringency. Korea has had a target that was mandatory for all vehicles rather than for a sales-weighted average, making it difficult to compare across countries. Since we do not face any of these issues for the EU, Japan and the US, we mainly focus on those countries.

³³This data is available at <https://theicct.org/chart-library-passenger-vehicle-fuel-economy>.

³⁴In the specifications with Canada, Mexico and South Korea these countries are also used to create weights. The weight assigned to the EU is based on all pre-sample patent applications in countries in our data set that use the EU standard.

Since stringency is 0 when a standard is non-binding, and we take natural logarithms in our regressions, we use $\log(1 + \textit{Stringency})$.

3.4 Knowledge stocks, spillovers, R&D subsidies, GDP

The same patent data discussed above is used to establish separate stocks of clean and dirty knowledge, which proxy for a firm’s R&D experience and productivity. These stocks are updated each year using the perpetual inventory method:

$$K_{Cit} = (1 - \delta)K_{Cit-1} + P_{Cit} \quad \text{and} \quad K_{Dit} = (1 - \delta)K_{Dit-1} + P_{Dit}, \quad (6)$$

where δ is the depreciation rate of knowledge, which accounts for the fact that some knowledge becomes obsolete over time. We set the knowledge depreciation rate at 20%, which is a value often assumed in the literature (e.g., Aghion et al. 2016).

All explanatory variables are included as their natural logarithm. Since knowledge stocks are zero until the first patent application, we follow the standard approach in the literature and add an arbitrary constant (equal to one) to all knowledge stocks to avoid taking the logarithm of zero. In addition, we add three dummies to the regression: one to indicate that $K_{Cit} = 0$, one for $K_{Dit} = 0$ and one in case both are equal to zero (Blundell et al., 1999).

In addition to firm level knowledge stocks, we account for knowledge accumulation at the country level. Jaffe et al. (1993), for instance, show the importance of location in knowledge spillovers. Since patents are recorded at the firm level and many firms operate from several countries, we need a method to create stocks at the country level. Since we are interested in geographical spillovers, it makes sense to consider the location where R&D takes place, rather than the location where firms sell their product.

We follow Aghion et al. (2016), who create weights that link firms to countries based on the location of firms’ inventors. Each patent file mentions its applicants and inventors, as well as their country of residence. We analyze patents at the applicant level, but use inventor location to create weights. For each firm in our data set we pool all the inventors that are mentioned on the patents we selected. We then define the (time-invariant) weight w_{ci}^K of firm i for country c as the proportion of firm i ’s inventors that are located in country c , according to the patent applications. Note that these weights are different from those used for prices, taxes and standards, which are based on patent applications (rather than inventor location).

The spillover stock of firm i for technology s is then defined as follows.

$$Spillover_{ist} = \sum_c w_{ci}^K Spillover_{cist} \quad \text{where} \quad Spillover_{cist} = \sum_{j \neq i} w_{cj}^K K_{jst}, \quad (7)$$

where K_{jst} is the firm-level knowledge stock. In words, the spillover stock for firm i in country c exists of the cumulative knowledge stocks of all other firms that have inventors located in country c , weighted by the proportion of each firm's inventors that are located in country c .

Another potential determinant of innovation that we control for are R&D expenditures by national governments. We take this data from the IEA's Energy Technology RD&D Statistics, which specifies government R&D budgets for a variety of technologies.³⁵ We use government support for energy efficiency in the category transport (category 13). These subsidies are mainly aimed at gray technologies, i.e. improving the efficiency of the internal combustion engine (Aghion et al., 2016). We do not have data on which firms received subsidies, so we use weights again to measure the exposure to government support. Since governments subsidize firms that do research in their country, we use the inventor location weights w^K , rather than the market exposure weights that we use for taxes and standards.

Our final control variable is GDP per capita, which is meant to capture overall economic conditions at the country level. We collect this data from the OECD database. We measure GDP in 2015 dollars (PPP). We use the same market exposure weights that we use for taxes and prices to transform GDP per capita to a firm-level variable.

4 Empirical strategy

4.1 Choice of estimator

Our dependent variable is a count of patents, so we use count models to estimate our regressions. Our regression equation is the following.

$$\begin{aligned} P_{ist} = \exp & \left(\beta_{s1} \log(S_{it-1}) + \beta_{s2} \log(FT_{it-1}) + \beta_{s3} \log(FP_{it-1}) + \beta_{s4} \log(RD_{it-1}) \right. \\ & + \beta_{s5} \log(KC_{it-1}) + \beta_{s6} \log(KD_{it-1}) + \beta_{s7} \log(SC_{it-1}) + \beta_{s8} \log(SD_{it-1}) \\ & \left. + \beta_{s9} \log(PP_i) + \beta_{s10} \log(PPD_i) + w_{it} \gamma_s \right) + u_{ist}, \end{aligned} \quad (8)$$

³⁵For more information about the IEA database see <https://www.iea.org/data-and-statistics/data-product/energy-technology-rd-and-d-budget-database-2>.

where i denotes a firm, s denotes a technology (either clean or dirty in most regressions), and t is the year. S is standard stringency, FT is the fuel excise, FP is the tax-exclusive fuel price, and RD are R&D expenditures. KC and KD are the firm’s own clean and dirty knowledge stocks, respectively. SC and SD measure exposure to stocks of knowledge at the country level to reflect spillovers of clean and dirty knowledge, respectively. PP and PPD are time-invariant variables based on pre-sample patenting to control for firm fixed effects. Finally, w includes control variable GDP per capita, dummies for when firm-level knowledge stocks are zero (three dummies: one for no clean knowledge, one for no dirty knowledge, one for when both stocks are zero), and a complete set of year dummies, u_{ist} is the idiosyncratic error term.

We use three models to estimate the effects of our variables of interest on the direction of innovation. The standard models for count data are Poisson and negative binomial.³⁶ The Poisson distribution has a single parameter λ and its mean and variance are equal.

$$\mathbb{E}(P) = \text{Var}(P) = \lambda, \tag{9}$$

where we parameterize λ for individual i , technology s and year t as $\lambda_{ist} = \exp(X_{it}\beta_s + \eta_i + \nu_t)$.

The negative binomial distribution adds overdispersion to the Poisson distribution. That is, its variance can be larger than its mean. We specify its mean and variance as follows.

$$\mathbb{E}(P) = \lambda \tag{10}$$

$$\text{Var}(P) = \lambda(1 + \delta), \tag{11}$$

where δ is the overdispersion parameter. If it is zero the two distributions are equal. We use the negative binomial model as our baseline because our patent data shows overdispersion. The negative binomial distribution is thus a better fit than Poisson. An advantage of the negative binomial model is that its coefficients are easy to interpret as elasticities, like for the Poisson model. We specify overdispersion to be constant across observations.³⁷

The third model we use is zero-inflated Poisson (ZIP), which is appropriate in the presence of excess zeros in the dependent variable. It simultaneously estimates the extensive margin

³⁶A linear model could predict a negative number of patents for certain inputs, whereas Poisson and negative binomial models never predict fewer than zero patents.

³⁷That is, we use the NB1 model from Cameron and Trivedi (2013), chapter 3.

(whether or not to innovate) and the intensive margin of innovation (how much to innovate). We define π as the probability that a firm does not engage in research. The idea is that there are two potential explanations for observing zero patents: the firm did not do any research or its research was not successful.³⁸ The ZIP distribution's mean and variance are as follows.

$$\mathbb{E}(P) = (1 - \pi)\lambda \tag{12}$$

$$\text{Var}(P) = (1 - \pi)\lambda + \pi(1 - \pi)\lambda^2. \tag{13}$$

It thus accounts for a specific type of overdispersion, namely the presence of excess zeros. We estimate the binary decision on whether or not to do research with a Logit model using the same explanatory variables that we use for the intensive margin (which is Poisson).³⁹ This means that the number of parameters to estimate is doubled, compared to the Poisson model. A disadvantage of the ZIP model is that the estimated coefficients are not as easy to interpret as those of the Poisson and negative binomial models. That is because each variable affects both the intensive and the extensive margin. To see the overall effect one needs to analyze the marginal effects.

4.2 Identification

The main challenge in identifying the effect of regulation on innovation is to specify regulatory stringency for both taxes and standards in an appropriate manner, as discussed in section 2 and 3. In addition, we also identify two econometric challenges. First is the potential presence of unobserved firm characteristics that lead to differences in the propensity to patent. We cannot use standard fixed effects models as not all our variables satisfy the strict exogeneity assumption. We address this by using pre-sample data to proxy the firm

³⁸Hence, we can specify the probability of observing a particular number of patents P_{ist} as follows:

$$\Pr[P = P_{ist}|X_{it}] = \begin{cases} \pi + (1 - \pi)Poi(0, \lambda) & \text{if } P_{ist} = 0 \\ (1 - \pi)Poi(P_{ist}, \lambda) & \text{if } P_{ist} > 0. \end{cases}$$

³⁹Hence, the probability that firm i does not engage in research into technology s in year t is modeled as

$$\pi_{ist} = \frac{\exp(X_{it}\beta_s^E + \eta_i + \nu_t)}{1 + \exp(X_{it}\beta_s^E + \eta_i + \nu_t)},$$

where β^E are the parameters for the extensive margin. The intensive margin parameters are denoted by β^I .

fixed effects. Second is the potential endogeneity of our regressors, as innovations may impact future standards, taxes and fuel prices. We address this by using lagged variables where necessary.

A well known problem in patent regressions with fixed effects is that knowledge stocks, which are an important determinant of new innovations, are not strictly exogenous (Blundell et al., 1995). This violates the assumptions of the standard fixed effects count data model (Hausman et al., 1984) and using this model would lead to biased results. In dealing with fixed effects we follow the approach of Blundell et al. (1995, 1999) and Noailly and Smeets (2015), who use firms' pre-sample patenting behavior to proxy for the firm fixed effect. In addition, we include a dummy variable that is one for firms that do not have any pre-sample patents to account for the possibility that these firms are structurally different (in equation (8) PP counts pre-sample patents and PPD is the dummy). Blundell et al. (1995) defend this approach by arguing that under the assumption that a firm's innovative search process is stationary and follows an AR1 process, average yearly search activity is proportional (up to a constant) to the firm fixed effect η_i . If average pre-sample patenting is a good proxy for search activity, then it can be used to proxy for the fixed effect. They then show that the inclusion of these pre-sample variables strongly reduces serial correlation, which suggests that the fixed effect has been eliminated.

For pre-sample patenting we use two specifications. The first is to use average yearly patenting in any category (not only clean and dirty car technologies) from the first year in which a firm files a patent until 1999, the last year of the pre-sample period. We include all technologies to increase the number of firms that have pre-sample data available and because overall patenting is arguably a good proxy for the propensity to patent, which is what the fixed effect captures. This is also the approach taken by Noailly and Smeets (2015). The second specification is to include both the clean and dirty knowledge stock in year 0, i.e. the last year of the pre-sample period (1999 in our case). This approach is taken by Aghion et al. (2016) in the regressions they label as BGVR. Dummies for no pre-sample data are included in both specifications.

Our second concern is the potential endogeneity of our explanatory variables, especially standard stringency and fuel prices and taxes. New innovations may influence future regu-

lation and vehicle demand, which potentially affects fuel prices. We thus use the first lag of our explanatory variables in our main specifications. This solves the issue of endogeneity and also makes sense in the context of patenting. A change in regulation may lead firms to respond by investing in R&D, but it takes time for this investment to result in patents. Hence, a lag of one year is reasonable. We show specifications with different lag structures in our robustness analysis.

A final decision considers which years to include in our analysis for each firm. If a firm applies for a patent in a particular year, we know it is active. However, when a firm does not apply for a patent, it could be that the firm ceased to exist (or did not yet exist) or that it was active but just did not do any research (or its research did not result in any patents). In our main specification, we include only the years in which we know a firm existed. That is, for each firm we find the first and last year in which it applied for a patent (any category), and we use those as start and end point of the active period. We show the alternative of using a balanced panel, which includes many more zeros, in our robustness checks. In each regression we only include the firms that have at least one patent in the relevant category. So firms with only dirty patents are not included in clean patent regressions.

5 Results

5.1 Main results

Table 4 shows our main results for counts of clean and dirty patents (columns 1-3 and 4-6, respectively). All of our estimations include fixed effects as described above, a full set of year dummies, GDP per capita, and dummies that indicate a knowledge stock of zero. We use the negative binomial model for our baseline results. Column 1 shows that both our indicator of standard stringency and the excise tax have a positive and significant effect on clean patenting. The coefficient of 0.19 for stringency can be interpreted as an elasticity: a 10% increase in $1 + \textit{Stringency}$ increases clean patenting by almost 2%. Similarly, a 10% increase of the fuel excise is associated with an almost 3% increase in clean patenting. The positive and significant impact of our two main variables of interest on clean patenting confirms the prediction that both taxes and standards induce clean innovation. The point estimates do

	Clean			Dirty		
	(1)	(2)	(3)	(4)	(5)	(6)
Standard stringency	0.191*** (0.067)	0.190*** (0.063)		0.096 (0.066)	0.115* (0.067)	
Fuel excise tax	0.284*** (0.107)		0.245** (0.112)	-0.137 (0.151)		-0.172 (0.157)
Tax-exclusive fuel price	-0.316 (0.415)		-0.554 (0.399)	-0.582* (0.305)		-0.751** (0.326)
R&D subsidy	-0.143 (0.168)	-0.298* (0.173)	0.007 (0.164)	0.123 (0.183)	0.152 (0.192)	0.191 (0.194)
Clean knowledge stock	1.034*** (0.024)	1.029*** (0.025)	1.022*** (0.026)	-0.077** (0.035)	-0.072* (0.038)	-0.080** (0.035)
Dirty knowledge stock	-0.002 (0.016)	0.001 (0.017)	0.002 (0.017)	1.090*** (0.044)	1.084*** (0.049)	1.089*** (0.045)
Clean spillover	0.193* (0.099)	0.0873 (0.089)	0.143 (0.098)	0.030 (0.091)	0.088 (0.104)	0.003 (0.094)
Dirty spillover	-0.189** (0.090)	-0.105 (0.086)	-0.154* (0.090)	-0.022 (0.076)	-0.063 (0.082)	0.002 (0.081)
Pre-sample average	0.082*** (0.013)	0.079*** (0.013)	0.086*** (0.014)	0.044** (0.022)	0.047** (0.022)	0.047** (0.021)
Pre-sample zero	0.485*** (0.053)	0.467*** (0.053)	0.511*** (0.057)	0.292*** (0.059)	0.294*** (0.058)	0.306*** (0.059)
$\ln(\delta)$	-0.023 (0.105)	-0.018 (0.105)	-0.012 (0.106)	0.075 (0.138)	0.085 (0.142)	0.080 (0.140)
Observations	25280	25280	25280	23439	23439	23439
Log likelihood	-16842.3	-16853.2	-16872.7	-15265.4	-15271.3	-15273.0

Table 4: Baseline regression results for clean and dirty patents. Estimation is done using the negative binomial model (with overdispersion parameter δ). Robust standard errors are clustered at the firm level and reported in parentheses. All regressions include knowledge stocks, spillover stocks, GDP per capita, dummies for knowledge stocks that are equal to zero (to account for firm fixed effects), pre-sample patents and a dummy for no pre-sample patents, and a full set of time dummies. All (time-variant) regressors are included as their first lag. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

not change much when we leave out either the fuel tax and price (column 2) or stringency (column 3).

The tax-exclusive fuel price and R&D subsidies have no significant impact. This is not entirely surprising, as the tax-exclusive price mainly captures level differences in prices, whilst price shocks, which are highly similar across countries, are absorbed by the time fixed effects. R&D subsidies are measured rather imprecisely, as we do not observe which firms are subsidized, and this might explain the absence of a significant effect. The marginally significant and negative impact of the R&D subsidy in column 2 is apparently related to the absence of the fuel price and might be also affected by the fact that these subsidies are mainly aimed at gray technologies (Aghion et al., 2016), which shifts resources away from clean initiatives.

Firms' own clean knowledge stocks have a strong and positive effect on clean patenting, whereas own dirty knowledge is insignificant. This confirms the DTC prediction that innovation is path dependent. A firm that has gained expertise in a particular area is likely to continue on the same path. The clean spillover has a positive effect, albeit only significant at 10%, while the dirty spillover has a negative and stronger significant effect. This indicates that local spillover effects are indeed a determinant of innovation. Pre-sample patenting is highly significant in all specifications, which suggests it captures the propensity to patent. The coefficient for the average number of yearly patents is positive, meaning that firms that applied for more patents in the pre-sample period, were also more likely to patent during the sample period. The dummy that indicates firms with no pre-sample patents is positive, which means that these firms are more likely to patent than other firms during the sample period. To a large extent, this captures patenting by firms that did not exist during the pre-sample period.

Surprisingly, we find no evidence that standards and excise taxes affect dirty patenting (columns 4-6). DTC theory predicts a negative effect, as firms would shift resources from dirty to clean innovation when regulation makes clean driving more attractive. However, we only find a negative impact of the tax-exclusive fuel price and only with significance at the 10% level (and at 5% when stringency is left out). Although this is the expected sign, it is surprising that the exclusive price has this effect whilst the excise tax does not. The impact of own knowledge stocks is as expected with a strong positive effect of dirty knowledge

and a smaller, negative effect of clean knowledge. This is another confirmation of the path dependence prediction. We also do not find a significant effect of the spillover stocks on dirty patenting. These results on own clean and dirty knowledge stocks are consistent with the literature, which has shown that past innovations are an important predictor of future patenting (Blundell et al., 1995).

Our results on standard stringency contribute a new insight to the literature. The most closely related work to ours is Aghion et al. (2016), who study the innovation effects of tax-inclusive fuel prices. They find a strong, positive effect of fuel prices on clean patents (coefficient around 1.0) and no impact from standards. Our coefficient for the fuel excise is smaller, but also significantly positive, while the coefficient for the tax-exclusive price is not distinguishable from zero, however. For dirty patents Aghion et al. (2016) find a significant coefficient of around -0.5 . The coefficient we find is of similar size but for the tax-exclusive price and not for the excise. We show the results of regressions that include the tax-inclusive price in our robustness analysis.

Our findings provide evidence that not only fuel taxes but also GHG emission standards strongly induce clean innovation for automobile technologies and the search for breakthrough technologies. Apparently our stringency indicator captures an important aspect of firms' innovation decisions. To interpret the magnitude of the effect it is best to consider the average value of $\log(1 + \textit{Stringency})$ in our clean sample (used in columns 1-3), which is 0.68, and its standard deviation, which is 0.59. A one standard deviation increase in $\log(1 + \textit{Stringency})$ is associated with an 11% increase in patenting for clean technologies. Our finding that regulation did stimulate clean patenting, but did not direct much research away from the dirty options is also interesting. Apparently regulation was not strong enough to reveal such an asymmetric impact and car manufacturers still had confidence in the future of the dirty market during our sample period.

5.2 Disaggregated technologies

In this subsection we take a closer look at the technologies that make up our clean and dirty categories. As noted before we label clean technologies, in particular electric and hydrogen vehicles and fuel cells, as breakthrough technologies. Electric and hydrogen vehicles are

	Electric	Hydrogen/fuel cell	Hybrid	Gray	Purely dirty
	(1)	(2)	(3)	(4)	(5)
Standard stringency	0.297*** (0.105)	0.223*** (0.080)	0.130 (0.120)	0.115* (0.065)	0.0482 (0.068)
Fuel excise tax	0.202 (0.194)	0.150 (0.154)	-0.379* (0.226)	-0.005 (0.166)	-0.246 (0.202)
Tax-exclusive fuel price	-0.758 (0.778)	0.185 (0.466)	-2.214*** (0.545)	0.417 (0.477)	-1.024*** (0.357)
R&D subsidy	-0.058 (0.294)	-0.193 (0.222)	0.150 (0.380)	-0.039 (0.333)	0.135 (0.221)
Clean knowledge stock	0.751*** (0.039)	1.096*** (0.040)	0.594*** (0.046)	0.019 (0.055)	-0.057 (0.061)
Dirty knowledge stock	0.241*** (0.040)	-0.215*** (0.038)	0.472*** (0.040)	1.037*** (0.035)	0.966*** (0.083)
Clean spillover	0.317** (0.133)	0.056 (0.116)	0.082 (0.180)	0.257** (0.120)	-0.129 (0.095)
Dirty spillover	-0.291** (0.131)	-0.036 (0.096)	-0.092 (0.155)	-0.283*** (0.010)	0.150* (0.090)
Pre-sample average	0.090*** (0.026)	0.028 (0.019)	-0.065* (0.037)	0.022 (0.030)	0.042 (0.028)
Pre-sample zero	0.522*** (0.102)	0.436*** (0.068)	0.355*** (0.138)	0.630*** (0.105)	0.191** (0.077)
$\ln(\delta)$	0.421*** (0.139)	0.034 (0.133)	0.179 (0.186)	0.144 (0.162)	-0.164 (0.181)
Observations	11257	16848	5512	10342	18986
Log likelihood	-7081.9	-11316.3	-3475.9	-6779.6	-11576.9

Table 5: Regression results for disaggregated technology classes. Estimation is done using the negative binomial model (with overdispersion parameter δ). Robust standard errors are clustered at the firm level and reported in parentheses. All regressions include knowledge stocks, spillover stocks, GDP per capita, dummies for knowledge stocks that are equal to zero, pre-sample patents and a dummy for no pre-sample patents (to account for firm fixed effects), and a full set of time dummies. All (time-variant) regressors are included as their first lag. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

completely clean, provided that the electricity and hydrogen that powers them is generated using renewable energy. Although hybrid vehicles still use gasoline, they also use technology strongly related to electric vehicles and therefore we also categorize them as clean. All dirty patents relate to the internal combustion engine but we consider the patents that aim to improve fuel efficiency as gray and the others as purely dirty. One could argue that hybrid technologies belong in the gray category, but we follow Aghion et al. (2016) and classify them as clean. As described in section 3 some patent families belong to multiple categories.

Table 5 shows the results of our estimations for the disaggregated technologies. We include the same explanatory variables as in our baseline, so we do not compute separate knowledge stocks for each technology class.⁴⁰ Columns 1 through 3 show that the positive effect of stringency that we find for clean patenting is mostly driven by technologies for electric and hydrogen powered vehicles. Interestingly, the effect of the fuel excise is not significant at the 5% level for any of the three clean categories, whereas we found a significantly positive effect for the aggregate measure of clean innovation in table 4. This difference could be due to the decreased sample size in our disaggregated estimations.

Our results for hybrid technologies are quite different from those for the other clean categories (see column 3). Although the standards also have a positive impact here, it is much smaller and insignificant. Moreover, both the fuel excise and the tax-exclusive fuel price have a negative impact. In particular, the impact of the tax-exclusive fuel price on patenting is strong and highly significant. Apparently a rise in this price affects research in this category negatively as is more or less similar to the research impacts on the purely dirty technologies (see Column 5). The only difference between these two categories is the highly significant positive coefficient for clean knowledge stocks for the hybrid category. This justifies our choice to still categorize this as a clean, although not as a breakthrough technology.

Column 4 shows the results for gray technologies, i.e. technologies that aim to make the internal combustion engine more fuel efficient. Our stringency indicator also has a positive impact here, though only significant at the 10% level. This might be explained by the fact that an increase in fuel efficiency of the (very) large share of current cars sold on the market contributes more strongly to compliance than the other technologies that still had to find their

⁴⁰The differences in the number of observations come from the fact that we include only firms with at least one patent in the relevant category.

way to consumers. Furthermore, we find no evidence of an effect for excise taxes, exclusive prices and R&D subsidies.

Our results clearly provide evidence for a strong impact of regulation on breakthrough technologies, while neither fuel excises nor tax-exclusive fuel prices seem to play a role. Though our estimation of the pooled clean technologies shows a positive and significant impact of the fuel excise, no such evidence could be found for the clean technologies separately. Interestingly, research related to fuel efficiency, our gray category, resembles to some extent the pattern of research in in electric and hybrid cars. At the same time research on hybrid and purely dirty technologies resembles the (negative) predictions on research effort by the DTC literature, in particular through the tax-exclusive fossil fuel price channel. In our robustness analysis we further investigate to what extent our results are sensitive to the specification of policy instruments, especially in the time dimension.

6 Robustness analysis

6.1 Alternative estimators

This subsection shows our main regressions estimated using the Poisson and ZIP estimators. The overdispersion parameter δ is substantially greater than 0 in our baseline regressions, meaning that the negative binomial model fits our data better than the Poisson model (which is a negative binomial model with $\delta = 0$). Table 10 shows our results for the baseline estimations with clean and dirty patents, estimated using the Poisson and ZIP models. Looking at the Poisson results in columns 1 and 2, the main results still hold. Stringency and the excise tax have a significant impact on clean patenting, but not on dirty patenting.

Columns 3 and 4 show our results for the zero-inflated Poisson model. We simultaneously estimate the extensive margin, i.e. the decision whether or not to innovate, and the intensive margin, which shows, given that innovation takes place, the number of patents that are filed. The extensive margin is estimated using a logit model and predicts the probability that a firm does not engage in research. Hence, a positive coefficient for the extensive margin implies a positive effect on the probability not to patent (i.e., a negative effect on the probability to patent). For the intensive margin a positive coefficient implies that, conditional on doing

	Clean	Dirty	Clean		Dirty	
	Poisson		Zero-inflated Poisson			
	(1)	(2)	Intensive	Extensive	Intensive	Extensive
	(1)	(2)	(3)		(4)	
Standard stringency	0.150*** (0.072)	0.084 (0.070)	0.158* (0.081)	0.041 (0.127)	0.085 (0.072)	0.186 (0.123)
Fuel excise tax	0.291** (0.132)	-0.244 (0.204)	0.013 (0.138)	-0.534** (0.235)	-0.308** (0.144)	-0.899*** (0.246)
Tax-exclusive fuel price	-0.500 (0.497)	-0.493 (0.391)	-0.739 (0.497)	-1.124 (0.837)	-0.134 (0.382)	-0.136 (0.791)
R&D subsidy	-0.050 (0.236)	0.160 (0.218)	0.098 (0.289)	0.632 (0.431)	-0.055 (0.234)	-0.009 (0.425)
Clean knowledge stock	1.048*** (0.025)	-0.091*** (0.032)	0.853*** (0.023)	-1.075*** (0.062)	-0.008 (0.029)	-0.037 (0.074)
Dirty knowledge stock	-0.015 (0.017)	1.115*** (0.045)	0.036** (0.016)	-0.143*** (0.055)	0.931*** (0.028)	-1.099*** (0.066)
Clean spillover	0.166 (0.119)	0.013 (0.101)	0.140 (0.123)	-0.184 (0.192)	0.091 (0.088)	0.119 (0.154)
Dirty spillover	-0.183* (0.106)	-0.013 (0.095)	-0.162 (0.111)	0.091 (0.185)	-0.086 (0.075)	-0.105 (0.136)
Observations	25280	23439	25280		23439	
Log likelihood	-18467.5	-16672.7	-17188.4		-15301.0	

Table 6: Regression results for alternative estimators. Estimation is done using the Poisson (columns 1 and 2) and ZIP (columns 3 and 4) models. Robust standard errors are clustered at the firm level and reported in parentheses. All regressions include knowledge stocks, spillover stocks, GDP per capita, dummies for knowledge stocks that are equal to zero, pre-sample patents and a dummy for no pre-sample patents (to account for firm fixed effects), and a full set of time dummies. All (time-variant) regressors are included as their first lag. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

research, the expected number of patents increases with this variable. For stringency we observe only a marginally significant effect along the intensive margin (p value of 0.051). The coefficient does have the expected sign. There is no effect along the extensive margin. In other words, more stringent standards do not lead to more or fewer firms actively doing research, but the firms that already do research do so more intensively.

Interestingly, the excise tax does not have a significant effect along the intensive margin, but does have a positive effect on the probability to patent (a negative effect on the probability not to patent). This suggests that whereas standards work through the intensity of research at already active firms, excise taxes work through the extensive margin by attracting firms that were not actively engaged in inventions before. Standards do not affect dirty patenting. Excise taxes have a negative effect on dirty patenting along the intensive margin and a positive effect on the probability to patent along the extensive margin. These effects are opposite to each other.⁴¹

Tables 7 and 8 show the results of our regressions with disaggregated technologies using the Poisson and ZIP model, respectively. Consistent with our baseline results, we find that the breakthrough technologies for hydrogen fuel cells and especially electric vehicles are responsible for the positive effect of standards on clean innovation. In the ZIP regressions we see that this effect comes mainly from the intensive margin, although there is some positive effect on entry of new firms as well. Fuel taxes do not have a significant effect on electric and hydrogen vehicle technologies when the technology classes are separated. They do have a negative effect on hybrid technologies along the intensive margin in the ZIP estimation. As we found before, standards do not seem to impact dirty patenting. The excise tax does not affect gray patenting in either model, but has a negative effect on purely dirty patenting along the intensive margin and a positive effect on the probability to patent along the extensive margin in the ZIP estimation. Table 7 shows that these opposing effects lead to a negative but insignificant overall effect in the Poisson model.

The regressions with different estimators show largely similar results to the ones we find in our baseline. Both standards and excise taxes seem to play an important role, especially for clean innovation, although the role of the excise taxes is less convincing at the disaggregated

⁴¹When computing the marginal effect it turns out that the intensive margin dominates, i.e., the marginal effect is negative (-0.13 for the average firm and -0.19 when we take the average marginal effect).

	Electric	Hydrogen/fuel cell	Hybrid	Gray	Purely dirty
	(1)	(2)	(3)	(4)	(5)
Standard stringency	0.304** (0.121)	0.156* (0.089)	0.106 (0.133)	0.092 (0.076)	0.030 (0.075)
Fuel excise tax	0.322 (0.269)	0.092 (0.205)	-0.480* (0.275)	-0.137 (0.207)	-0.386 (0.279)
Tax-exclusive fuel price	-0.261 (1.066)	-0.233 (0.554)	-3.016*** (0.652)	0.727 (0.533)	-1.203** (0.469)
R&D subsidy	0.285 (0.520)	-0.195 (0.287)	0.126 (0.513)	-0.013 (0.413)	0.192 (0.243)
Clean knowledge stock	0.817*** (0.044)	1.141*** (0.043)	0.590*** (0.049)	-0.017 (0.050)	-0.097 (0.059)
Dirty knowledge stock	0.193*** (0.038)	-0.229*** (0.037)	0.491*** (0.044)	1.084*** (0.039)	1.025*** (0.079)
Clean spillover	0.238 (0.178)	0.046 (0.148)	-0.011 (0.232)	0.273** (0.130)	-0.181 (0.115)
Dirty spillover	-0.218 (0.187)	-0.036 (0.117)	-0.033 (0.209)	-0.318*** (0.113)	0.213* (0.122)
Observations	11257	16848	5512	10342	18986
Log likelihood	-8191.0	-12430.3	-3882.2	-7471.1	-12374.4

Table 7: Regression results for disaggregated technologies. Estimation is done using the Poisson model. Robust standard errors are clustered at the firm level and reported in parentheses. All regressions include knowledge stocks, spillover stocks, GDP per capita, dummies for knowledge stocks that are equal to zero, pre-sample patents and a dummy for no pre-sample patents (to account for firm fixed effects), and a full set of time dummies. All (time-variant) regressors are included as their first lag. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Electric	Hydrogen/fuel cell	Hybrid	Gray	Purely Dirty
	(1)	(2)	(3)	(4)	(5)
Intensive margin					
Standard stringency	0.341*** (0.110)	0.098 (0.100)	0.215 (0.150)	0.062 (0.090)	0.080 (0.081)
Fuel excise tax	-0.104 (0.257)	-0.051 (0.211)	-0.830*** (0.289)	-0.309 (0.231)	-0.498** (0.254)
Tax-exclusive fuel price	-0.876 (0.908)	-1.016* (0.544)	-3.568*** (0.839)	0.977 (0.613)	-0.879* (0.481)
R&D subsidy	0.366 (0.488)	-0.016 (0.364)	0.001 (0.579)	0.037 (0.433)	-0.089 (0.264)
Clean knowledge stock	0.675*** (0.043)	0.870*** (0.050)	0.468*** (0.070)	-0.010 (0.054)	-0.010 (0.054)
Dirty knowledge stock	0.174*** (0.036)	-0.148*** (0.036)	0.431*** (0.056)	0.991*** (0.043)	0.809*** (0.069)
Extensive margin					
Standard stringency	0.251* (0.151)	-0.194 (0.135)	0.380 (0.245)	0.008 (0.226)	0.195 (0.132)
Fuel excise tax	-0.637* (0.362)	-0.387 (0.322)	-1.029 (0.674)	-0.968 (0.590)	-0.857*** (0.342)
Tax-exclusive fuel price	-0.946 (1.224)	-1.815* (1.026)	-2.518 (2.149)	-0.217 (1.627)	-0.225 (0.927)
R&D subsidy	0.718 (0.589)	0.551 (0.569)	0.136 (0.961)	0.652 (0.824)	-0.240 (0.486)
Clean knowledge stock	-0.480*** (0.084)	-0.834*** (0.085)	-0.455*** (0.123)	-0.128 (0.139)	-0.064 (0.098)
Dirty knowledge stock	-0.343*** (0.067)	0.217*** (0.068)	-0.303*** (0.094)	-0.512*** (0.097)	-0.889*** (0.089)
Observations	11257	16848	5512	10342	18986
Log likelihood	-7414.0	-11574.7	-3624.0	-7065.1	-11457.8

Table 8: Regression results for disaggregated technologies. Estimation is done using the ZIP model. Robust standard errors are clustered at the firm level and reported in parentheses. All regressions include knowledge stocks, spillover stocks, GDP per capita, dummies for knowledge stocks that are equal to zero, pre-sample patents and a dummy for no pre-sample patents (to account for firm fixed effects), and a full set of time dummies. All (time-variant) regressors are included as their first lag. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

level. This strengthens our confidence in our preferred estimator, the negative binomial model.

6.2 Standard stringency

This subsection explores alternative specifications of standard stringency to assess whether they match our baseline result. Tables 9 and 10 show the results for clean and dirty patenting, respectively. Our baseline regressions include the first lag of our stringency measure, which is defined as the required yearly reduction in CO₂ emissions per kilometer (NEDC) in order to meet the most recently announced target. We use regulations from five countries and the EU for our baseline. We start by analyzing two different lag structures. We then restrict our attention to the three main regulators when it comes to standards. Finally, we employ three alternative measures of stringency.

In column 1 of both tables 9 and 10 we use the second lag of our indicator, rather than the first lag. This change has little impact on the results. For clean patents we still find a significant estimate of around 0.2. For dirty patents we still find a coefficient of around 0.1 that is not significant at the 5% level. The same holds when we use contemporaneous stringency measure instead of the first lag. This finding is not entirely surprising, as innovation takes time, and one innovation may take more time than another one. The time between the policy impulse and the moment of invention can range from less than one year to multiple years. Our preferred specification includes a lag to mitigate concerns about endogeneity. This robustness check suggests that that is not an issue, perhaps because we analyze anticipation effects, which already account for the period between announcement and implementation of a target.

In column 3 we include our stringency indicator computed using only the EU, Japan and the US, i.e. leaving out Canada, Korea and Mexico. We include this specification because the data on either standards or performance for the latter three countries is imperfect (see section 3). This adaption does not alter our results in any meaningful way, which is not surprising because the countries that are left out generally have small weights compared to the EU, Japan and the US. Hence, for most firms in our data set the stringency indicator hardly changes when Canada, Korea and Mexico are left out.

	Clean					
	(1)	(2)	(3)	(4)	(5)	(6)
Standard stringency ($t - 2$)	0.206*** (0.075)					
Standard stringency (t)		0.194*** (0.071)				
Standard stringency (EU, JP, US)			0.184*** (0.069)			
Performance (EU, JP, US)				3.603*** (1.184)		
Target (EU, JP, US)					-1.046** (0.408)	
Distance to target						0.088*** (0.030)
Fuel excise tax	0.283** (0.110)	0.279*** (0.107)	0.287*** (0.106)	0.571*** (0.160)	0.092 (0.122)	0.300*** (0.107)
Tax-exclusive fuel price	-0.304 (0.423)	-0.389 (0.406)	-0.324 (0.411)	-1.040*** (0.340)	-0.166 (0.438)	-0.232 (0.426)
R&D subsidy	-0.123 (0.168)	-0.154 (0.169)	-0.139 (0.166)	-0.159 (0.163)	-0.160 (0.178)	-0.140 (0.167)
$\ln(\delta)$	-0.020 (0.106)	-0.029 (0.105)	-0.022 (0.105)	-0.032 (0.104)	-0.013 (0.107)	-0.021 (0.106)
Observations	23911	25280	25280	25280	25280	25280
Log likelihood	-16024.5	-16838.1	-16846.0	-16843.7	-16856.6	-16844.1

Table 9: Regression results for alternative stringency specifications and clean patents. Estimation is done using the negative binomial model (with overdispersion parameter δ). Robust standard errors are clustered at the firm level and reported in parentheses. All regressions include knowledge stocks, spillover stocks, GDP per capita, dummies for knowledge stocks that are equal to zero, pre-sample patents and a dummy for no pre-sample patents (to account for firm fixed effects), and a full set of time dummies. All (time-variant) regressors are included as their first lag, unless specified otherwise. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Dirty					
	(1)	(2)	(3)	(4)	(5)	(6)
Standard stringency ($t - 2$)	0.144*					
	(0.074)					
Standard stringency (t)		0.080				
		(0.068)				
Standard stringency (EU, JP, US)			0.090			
			(0.067)			
Performance (EU, JP, US)				-1.207		
				(1.411)		
Target (EU, JP, US)					-0.845**	
					(0.373)	
Distance to target						0.045
						(0.029)
Fuel excise tax	-0.143	-0.148	-0.138	-0.285	-0.295*	-0.130
	(0.172)	(0.151)	(0.151)	(0.237)	(0.178)	(0.151)
Tax-exclusive fuel price	-0.651*	-0.623**	-0.598*	-0.691*	-0.457	-0.540*
	(0.371)	(0.303)	(0.305)	(0.355)	(0.332)	(0.310)
R&D subsidy	0.133	0.129	0.127	0.268	0.055	0.123
	(0.181)	(0.182)	(0.183)	(0.192)	(0.189)	(0.183)
$\ln(\delta)$	0.082	0.075	0.075	0.079	0.081	0.075
	(0.137)	(0.138)	(0.138)	(0.140)	(0.140)	(0.139)
Observations	22084	23439	23439	23439	23439	23439
Log likelihood	-14412.2	-15267.2	-15266.4	-15270.5	-15263.4	-15265.4

Table 10: Regression results for alternative stringency specifications and dirty patents. Estimation is done using the negative binomial model (with overdispersion parameter δ). Robust standard errors are clustered at the firm level and reported in parentheses. All regressions include knowledge stocks, spillover stocks, GDP per capita, dummies for knowledge stocks that are equal to zero, pre-sample patents and a dummy for no pre-sample patents (to account for firm fixed effects), and a full set of time dummies. All (time-variant) regressors are included as their first lag, unless specified otherwise. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In columns 4 through 6 we include three different measures of standard stringency. Column 4 shows the results when we only use the first lag of countries' average performance to measure stringency (reflecting the same data that we use to compute our stringency measure). Under the assumption that firms do just enough to meet the target and do not anticipate potential future impacts of stricter standards, they have no reason to outperform the standard and average performance would be a reasonable proxy for stringency. This assumption is questionable, however, as consumers value fuel economy and firms may have an incentive to improve it even in the absence of regulation. In fact, performance can even be seen as an inverse proxy of stringency, as firms that perform badly have the largest incentive to innovate when a binding standard is introduced because of the (discounted) costs of future non-compliance and the (high) probability of a time lag in innovation. We only include the EU, Japan and the US in this specification, as the other countries do not have data available for all years.⁴² Note that a higher value for performance (measured as CO₂ emissions per kilometer) indicates higher emissions for a country's car fleet, making performance worse. Indeed, in our regression for clean patents we see that performance has a positive coefficient, which means that worse performance (higher CO₂ per kilometer) is associated with more clean patenting. In other words, the firms that are furthest away from the target will innovate more. Hence, average performance is certainly a relevant attribute to measure standard stringency, but it should not be used as a stringency measure independently. As a stringency measure, performance makes the strong assumption that firms exactly satisfy the target each year and it does not account for anticipation effects. Our stringency measure uses the distance between performance and the target and does account for the time dimension. It should therefore be preferred. For dirty patents we no longer find a significant effect.

Column 5 shows the results when we measure stringency as the most recently announced target at a given date. Again, as the target is expressed in CO₂ emissions per kilometer, a low value for the target corresponds to a more stringent standard. In this specification we include only the EU, Japan and the US, as the other countries did not have standards for all years. It is not obvious what value to use for target when there is no target, especially when lower values imply higher stringency. This stringency measure does not account for

⁴²This incompleteness is not a problem for our baseline specification as no standard was in place for the years for which data is missing. Hence, stringency was 0 in those years due to the binding indicator.

the time between announcement and enforcement, the difference with current performance, and whether a standard is voluntary or mandatory. Despite these limitations, column 5 in table 9 shows that the effect on clean patenting is consistent with our baseline result. The negative estimate implies that a more stringent target (lower value for target) increases clean patenting. Since we already showed that performance also affects innovation, we prefer our stringency measure, which accounts both for the target and for performance. Quite surprisingly, we also find a negative coefficient and thus a positive effect of the target on dirty patenting in column 5 of table 10.

In column 6 we leave out the time dimension and measure stringency as the distance between current performance and the most recently announced target, if the standard is binding for the average firm (so equation (1) but setting the denominator to 1). Contrary to our preferred indicator, stringency in this specification only increases when a new target is announced.⁴³ This measure does not take into account, however, that a requirement on firms to cut emissions by a certain amount in 3 years is more stringent than a requirement to cut the same amount in 4 years. Column 6 (both tables) shows that this alternative indicator gives similar results to our baseline for both clean and dirty patents. The effect on clean patents is positive and significant although there is no significant effect on dirty technologies.

The results in this subsection strengthen our trust in the standard stringency measure we have developed. It is an intuitive measure that is highly robust to different specifications in terms of lags and the countries we include. Its separate components also show the expected signs in most regressions and are highly significant, but do not provide information on all the relevant attributes of the instrument, i.e. stringency compared to current performance, timing and whether or not the target is binding, at the same time.

Though we account for several important features of the standards, we are unable to include all of them. For instance, even though the *Binding* indicator distinguishes between voluntary and mandatory standards, we do not include a measure of the (expected) enforcement of mandatory policies. If firms believe a standard will not be enforced (or could be gamed upon by optimizing the test cycle) they have less incentive to innovate. This is clearly relevant, but including actual enforcement is infeasible at this stage. A closely related di-

⁴³It also increases when performance gets worse over time, but this is rarely the case (see figures 1 and 5).

mension is the level of fines. Some firms may be better off paying fines each year than they are if they invest in innovation. Including fines could be an interesting future extension to our stringency measure. Other dimensions that we do not take into account are phase-in periods, supercredits, compensation schemes and exceptions. Furthermore, if a standard is not binding for the average firm, it may still be binding for firms that perform worse than the average. We do not account for this type of heterogeneity. Due to the limitations outlined above we do not use our stringency indicator to predict the exact number of patents that will result from a particular new standard. Rather, we use our measure to show that these standards have a clear impact on the direction of innovation. For this purpose we believe that the strengths outweigh the limitations and that our measure is insightful as such.

6.3 Fuel taxes and prices

In this subsection we include the tax-inclusive fuel price, rather than the excise tax and the tax-exclusive price separately. As noted before the previous literature did establish its DTC confirmation based on an estimation that included the tax-inclusive price and a weak measure for standard stringency (see section 2).⁴⁴ Table 11 shows the results.

In column 1 we see that if we include the tax-inclusive fuel price in our specification we find a positive effect on clean patenting like Aghion et al. (2016) at first sight, although our standard stringency measure is again very robust. Interestingly, if we compare a specification of our model without our standard stringency measure and R&D subsidies (column 2 for clean and column 5 for dirty) to be as close as possible to table 5 in Aghion et al. (2016), we find a consistent outcome, but with a much smaller (for clean) and insignificant coefficient.⁴⁵ This suggests that using the tax-inclusive fuel price, rather than the excise and tax-exclusive price separately, reduces the explanatory power of our model.

There are several potential explanations for these different findings. First, the effect of

⁴⁴We compare our findings to Aghion et al. (2016) who use the same patent selection criteria, but apply their analysis to a sample period from 1986 until 2005. They also employ a different estimator in their main regressions. Their estimator uses a control function to deal with unobserved differences in firm characteristics. However, they also show estimates for the estimator introduced by Blundell et al. (1999) that is similar to the one we use (they refer to this is the BGVR estimator).

⁴⁵Aghion et al. (2016) also have a measure of (air quality) standard stringency and R&D subsidies, but these are not included in their estimations shown in their table 5.

	Clean			Dirty		
	(1)	(2)	(3)	(4)	(5)	(6)
Tax-inclusive fuel price	0.375*	0.249	0.295**	-0.361	-0.482	-0.060
	(0.228)	(0.238)	(0.133)	(0.284)	(0.304)	(0.169)
Standard stringency	0.201***			0.095		
	(0.064)			(0.064)		
R&D subsidy	-0.256			0.128		
	(0.169)			(0.189)		
GDP per capita	0.441	-0.115		-1.042**	-1.148**	
	(0.496)	(0.420)		(0.521)	(0.535)	
Pre-sample average	0.082***	0.087***	0.088***	0.044**	0.046**	0.062***
	(0.013)	(0.014)	(0.013)	(0.022)	(0.021)	(0.019)
Pre-sample zero	0.471***	0.499***	0.501***	0.295***	0.305***	0.303***
	(0.053)	(0.058)	(0.057)	(0.057)	(0.058)	(0.060)
$\ln(\delta)$	-0.018	-0.006	-0.005	0.072	0.079	0.112
	(0.106)	(0.106)	(0.108)	(0.136)	(0.139)	(0.148)
Observations	25280	25280	25280	23439	23439	23439
Log likelihood	-16848.2	-16882.4	-16882.6	-15265.5	-15275.0	-15294.5

Table 11: Regression results for tax-inclusive fuel prices. Estimation is done using the negative binomial model (with overdispersion parameter δ). Robust standard errors are clustered at the firm level and reported in parentheses. All regressions include knowledge stocks, spillover stocks, GDP per capita (unless indicated otherwise), dummies for knowledge stocks that are equal to zero, pre-sample patents and a dummy for no pre-sample patents (to account for firm fixed effects), and a full set of time dummies. All (time-variant) regressors are included as their first lag. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the fuel price may have become weaker over the years since we cover a different sample period (2000-2016 compared to 1986-2005). Second, the effect may still be there but we lack the statistical power to show it in our regressions. Surprisingly, when we leave out our control variable GDP per capita, the coefficient becomes significant at the 5% level (column 3). This is surprising, as fuel prices and GDP per capita are not correlated and GDP per capita was not significant in column 2. Leaving out GDP per capita in our baseline regressions does not lead to any meaningful changes.

Columns 4 through 6 show our results for dirty patenting. Again, we find coefficient estimates that are consistent with Aghion et al. (2016) in terms of their sign, namely negative, but not in terms of significance. They find a significantly negative relation between fuel prices and dirty patenting, whereas our coefficients are not significantly different from zero. Leaving out GDP per capita does not change this outcome for dirty patents like it does for clean ones. These findings confirm (again) the very robust role of our measure of standard stringency, but also the importance of separating fuel prices and fuel excises in order to properly identify the role of policy-driven impulses for innovation.

6.4 Firm fixed effects

This subsection assesses the extent to which our results might depend on our method of dealing with firm fixed effects. As described in section 4, we use two variables to proxy for unobserved differences in firm characteristics, namely the average number of yearly patents in any category during the pre-sample period 1978-1999, and a dummy that is 1 if a firm has no pre-sample patents. With this approach we follow Noailly and Smeets (2015), who argue that average yearly patents are a reasonable proxy for the propensity to patent. In this subsection we explore a different strategy, namely the one used by Aghion et al. (2016) in their BGVR specifications. They include each firm's clean and dirty knowledge stocks in the last year of the pre-sample period as a proxy for the fixed effect, as well as dummies for zero pre-sample knowledge stocks.

Table 12 shows the results. In columns 1 and 4 we include both yearly patents and pre-sample knowledge stocks, so we combine our approach with the one from Aghion et al. (2016). In columns 2 and 5 we include only the pre-sample knowledge stocks. In columns

	Clean			Dirty		
	(1)	(2)	(3)	(4)	(5)	(6)
Standard stringency	0.190*** (0.070)	0.216*** (0.072)	0.219*** (0.070)	0.109* (0.063)	0.126** (0.061)	0.134** (0.061)
Fuel excise tax	0.210** (0.100)	0.116 (0.084)		0.097 (0.099)	0.049 (0.084)	
Tax-exclusive fuel price	-0.181 (0.453)	-0.170 (0.425)		-0.101 (0.337)	-0.188 (0.327)	
Tax-inclusive fuel price			0.123 (0.176)			0.119 (0.157)
R&D subsidy	-0.183 (0.174)	-0.300* (0.154)	-0.350** (0.153)	0.077 (0.149)	-0.022 (0.147)	-0.055 (0.143)
Pre-sample KC	-0.098*** (0.032)	-0.106*** (0.031)	-0.108*** (0.031)	0.121* (0.064)	0.125** (0.058)	0.125** (0.058)
Pre-sample KD	-0.083** (0.038)	-0.051 (0.040)	-0.054 (0.040)	-0.348*** (0.032)	-0.321*** (0.031)	-0.324*** (0.031)
No KC_0 dummy	-0.292*** (0.077)	-0.292*** (0.070)	-0.298*** (0.073)	0.057 (0.081)	-0.001 (0.074)	-0.002 (0.073)
No KD_0 dummy	0.067 (0.089)	-0.021 (0.080)	-0.019 (0.081)	-0.780*** (0.069)	-0.660*** (0.060)	-0.665*** (0.060)
Pre-sample average	0.0819*** (0.015)			0.057*** (0.019)		
Pre-sample zero	0.511*** (0.051)			0.543*** (0.052)		
$\ln(\delta)$	-0.046 (0.109)	-0.019 (0.110)	-0.018 (0.110)	-0.143 (0.129)	-0.133 (0.129)	-0.132 (0.129)
Observations	25280	25280	25280	23439	23439	23439
Log likelihood	-16800.2	-16889.1	-16890.7	-14932.9	-15012.1	-15012.3

Table 12: Regression results for alternative fixed-effects specifications. Estimation is done using the negative binomial model (with overdispersion parameter δ). Robust standard errors are clustered at the firm level and reported in parentheses. All regressions include knowledge stocks, spillover stocks, GDP per capita, dummies for knowledge stocks that are equal to zero, pre-sample patents and a dummy for no pre-sample patents (unless indicated otherwise), and a full set of time dummies. All (time-variant) regressors are included as their first lag. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3 and 6 we include only the pre-sample knowledge stocks and we include tax-inclusive fuel prices rather than the price and tax separately (so both the fixed effects specification and the fuel price are as in the BGVR regressions in Aghion et al. 2016). The results for standard stringency do not change for clean patenting. We still find a coefficient of around 0.2 that is highly significant. For dirty patents stringency turns significant at the 5% level when we use only the alternative fixed effects method. The size of the coefficient does not change much compared to our baseline results. The effect of the excise tax on clean patenting does depend on our fixed effects approach. When we use only the pre-sample knowledge stocks, i.e. column 2, the excise is no longer significant. The coefficient is still positive, and slightly smaller than before. As in table 4, the effect of the tax-inclusive fuel price is positive but insignificant. When we use both methods (column 1) the coefficient remains significant. Changes to our fixed effects specification have little effect on our finding for fuel taxes for dirty patents.

The results of this robustness check lead us to conclude that stringency is robust to the specification of the fixed effect, but that the excise tax is not. It is important to note, however, that the fixed effect specification mainly affects the precision of our estimate, rather than its sign or its size.

6.5 Firm and year inclusion

Our next robustness check concerns our decision to include only firms that are active innovators in particular technologies and only the years for which we are sure that firms are active. In our baseline regressions we only include firms that have patented at least once in the relevant category during our sample period. That is, in the regression for clean patents we only include firms with at least one clean patent. We also restrict the years that we include for each firm to the period between its first patent application and its last (for any technology). We do this in order to exclude years in which a firm did not yet exist or had stopped to exist. In this subsection we include all firms and all years in our regressions.

Table 13 shows the results. In columns 1 and 4 we include all firms and all years. In other words, we use a balanced sample of 3646 firms and 16 years. In columns 2 and 5 we include all firms, but only in their active years. In columns 3 and 6 we include only active firms, but all years. None of our main results depend on our decisions to include only active firms and

	Clean			Dirty		
	(1)	(2)	(3)	(4)	(5)	(6)
Standard stringency	0.194*** (0.064)	0.191*** (0.065)	0.198*** (0.066)	0.083 (0.064)	0.089 (0.064)	0.091 (0.066)
Fuel excise tax	0.387*** (0.101)	0.358*** (0.104)	0.305*** (0.104)	-0.101 (0.161)	-0.129 (0.158)	-0.115 (0.153)
Tax-exclusive fuel price	-0.365 (0.408)	-0.446 (0.417)	-0.280 (0.405)	-0.381 (0.318)	-0.568* (0.315)	-0.404 (0.308)
R&D subsidy	-0.154 (0.167)	-0.151 (0.169)	-0.149 (0.167)	0.160 (0.190)	0.166 (0.188)	0.130 (0.184)
$\ln(\delta)$	-0.029 (0.107)	-0.036 (0.106)	-0.014 (0.105)	0.066 (0.139)	0.062 (0.140)	0.080 (0.137)
Firms	All	All	Active	All	All	Active
Years	All	Active	All	All	Active	All
Observations	58336	40473	33952	58336	40473	33712
Log likelihood	-19339.1	-18058.3	-18000.8	-17997.9	-16669.8	-16539.4

Table 13: Regression results for alternative firm and year inclusion criteria. Estimation is done using the negative binomial model (with overdispersion parameter δ). Robust standard errors are clustered at the firm level and reported in parentheses. All regressions include knowledge stocks, spillover stocks, GDP per capita, dummies for knowledge stocks that are equal to zero, pre-sample patents and a dummy for no pre-sample patents (to account for firm fixed effects), and a full set of time dummies. All (time-variant) regressors are included as their first lag. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Clean			Dirty		
	(1)	(2)	(3)	(4)	(5)	(6)
Standard stringency	-0.083 (0.128)	0.287*** (0.071)	0.135** (0.060)	0.453*** (0.128)	0.049 (0.083)	0.067 (0.058)
Fuel excise tax	0.070 (0.168)	0.257** (0.129)	0.253** (0.102)	0.167 (0.139)	-0.185 (0.168)	-0.084 (0.132)
Tax-exclusive fuel price	-0.644 (0.506)	-0.285 (0.509)	-0.203 (0.406)	-0.221 (0.702)	-0.502 (0.501)	-0.433 (0.293)
R&D subsidy	-0.209 (0.298)	-0.224 (0.199)	-0.217 (0.167)	0.679* (0.382)	0.091 (0.227)	0.055 (0.186)
$\ln(\delta)$	-0.255 (0.159)	0.089 (0.107)	0.227** (0.112)	-0.058 (0.186)	0.108 (0.144)	0.393*** (0.133)
Years	2000-2008	2009-2016	2000-2016	2000-2008	2009-2016	2000-2016
Restricted firms	No	No	Yes	No	No	Yes
Observations	12411	12869	17870	12018	11421	16597
Log likelihood	-7673.0	-8935.1	-12508.2	-7825.6	-7257.0	-11581.9

Table 14: Regression results for restricted samples. Estimation is done using the negative binomial model (with overdispersion parameter δ). Robust standard errors are clustered at the firm level and reported in parentheses. All regressions include knowledge stocks, spillover stocks, GDP per capita, dummies for knowledge stocks that are equal to zero, pre-sample patents and a dummy for no pre-sample patents (to account for firm fixed effects), and a full set of time dummies. All (time-variant) regressors are included as their first lag. Stars indicate significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

active years, as the estimates for standard stringency and the excise tax are very similar to those found in our baseline, both in terms of size and significance.

6.6 Restricted years and firms with weights

In our final robustness check we restrict the sample in two ways. First we restrict the sample period to see whether the Japanese standards or the European and American standards were most innovation-inducing. We then restrict the firms we include to only those that have pre-sample patent portfolios to see whether this influences our results.

As figure 2 in section 2 shows, standards in the EU and the US only became binding for the average firm in 2009 and 2010, respectively. Stringency then rapidly increased in both jurisdictions. Japanese standards, on the other hand, were binding until 2012, and stringency was rather stable before declining from 2009. We estimate our main regressions separately for the years in which only Japanese standards were binding (2000-2008) and for the years in which European and American standards became binding (2009-2016). Table 14 shows the results. Column 1 shows that stringency did not significantly impact clean patenting until 2008. Interestingly, we find that the excise tax did not have a significant impact either. Column 2 shows that the period from 2009 until 2016 is driving our main results. These are the years in which the EU and the US started tightening regulation. Quite surprisingly, we find in column 4 that the Japanese standards had a highly significant and positive effect on dirty patenting, until 2008. In the subsequent period (column 5) this effect disappears.

Not all firms in our sample filed a patent before the year 2000. These firms were either founded in later years or were not yet innovative in the pre-sample period. These firms were assigned the average market exposure weights to compute stringency, taxes and prices. This group of firms is not random. Young firms are in this group by construction and may be structurally different innovators compared to old firms. This concern is especially relevant for clean innovators, as these are newer technologies. We therefore run our main regressions with just the firms that have pre-sample patents (and thus have their own weights). Column 3 shows that our main results for stringency and excise taxes do not change much for clean patenting. The coefficient for stringency is slightly smaller, but still significant at the 5% level, while the coefficient for the tax is approximately equal to our baseline. Column 6 shows that our results for dirty patenting are also the same as our baseline. Neither standards nor taxes have a significant impact. These results strengthen our confidence in the chosen methodology.

7 Conclusion

We find strong evidence that fuel economy and GHG emission standards induce clean innovation. Our indicator of stringency, which takes into account the distance between current performance and the target, as well as timing and whether a standard is binding, has a significantly positive effect on clean innovation in our baseline regressions. This effect is driven

by inventions in electric and hydrogen vehicle technologies. It is robust to a variety of specifications. We also find that excise taxes on fuel positively affect clean innovation, which is in line with the existing literature. Our results for excise taxes are not as robust as those for standards. We do not find evidence of a negative impact of either policy on dirty patenting, as is predicted by the DTC literature. We do find that tax-exclusive fuel prices negatively affect dirty patenting. We find no evidence that R&D subsidies impact innovation. Existing stocks of knowledge do contribute to further innovation.

These results show that not only prices and price instruments, but also non-market or command-and-control instruments are an important driver of innovation. Our results have implications for policy discussions. Since both standards and taxes contribute to clean innovation, we cannot conclude that one policy should be preferred over the other from an innovation perspective. It is important, however, to take the innovation-inducing effects of both instruments into account when deciding on which policy to implement. We contribute to the literature on directed technical change by showing the importance of standards for innovation, and to the literature on regulation of emissions from road transport by showing the effects of the two most important policies on clean and dirty innovation.

Furthermore, we show that it is important to measure regulatory stringency in a manner that is appropriate for innovation decisions. Our indicator of standard stringency accounts for anticipation effects, which are important for dynamic processes like innovation. Interesting extensions of our work would be to include fines in a similar stringency measure to test their effect, and to measure stringency at the manufacturer level (an issue with this is that innovation also takes place at other firms, like component suppliers and research firms). We leave this for future work. It will also be interesting to test the effects of the stringent standards that have been introduced since 2016 and that we could not assess because of data limitations for patents.

References

Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012). The environment and directed technical change. *American Economic Review*, 102(1):131–66.

- Acemoglu, D., Akcigit, U., and Celik, M. A. (2021). Radical and incremental innovation: The roles of firms, managers and innovators. *American Economic Journal: Macroeconomics*.
- Acemoglu, D., Akcigit, U., Hanley, D., and Kerr, W. (2016). Transition to clean technology. *Journal of Political Economy*, 124(1):52–104.
- Aghion, P., Dechezleprêtre, A., Hemous, D., Martin, R., and Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1):1–51.
- Allcott, H. and Wozny, N. (2014). Gasoline prices, fuel economy, and the energy paradox. *The Review of Economics and Statistics*, 96(5):779–795.
- Anderson, S. T., Parry, I. W. H., Sallee, J. M., and Fischer, C. (2011). Automobile Fuel Economy Standards: Impacts, Efficiency, and Alternatives. *Review of Environmental Economics and Policy*, 5(1):89–108.
- Anderson, S. T. and Sallee, J. M. (2016). Designing policies to make cars greener. *Annual Review of Resource Economics*, 8(1):157–180.
- Bento, A. M., Jacobsen, M. R., Knittel, C. R., and van Benthem, A. A. (2020). Estimating the costs and benefits of fuel-economy standards. *Environmental and Energy Policy and the Economy*, 1:129–157.
- Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.
- Blundell, R., Griffith, R., and Reenen, J. V. (1995). Dynamic Count Data Models of Technological Innovation. *The Economic Journal*, 105(429):333–344.
- 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.
- Brunel, C. and Levinson, A. (2013). Measuring environmental regulatory stringency. *OECD Trade and Environment Working Papers*, 2013/05.

- Brunel, C. and Levinson, A. (2016). Measuring the stringency of environmental regulations. *Review of Environmental Economics and Policy*, 10(1):47–67.
- Brunnermeier, S. B. and Cohen, M. A. (2003). Determinants of environmental innovation in US manufacturing industries. *Journal of environmental economics and management*, 45(2):278–293.
- Busse, M. R., Knittel, C. R., and Zettelmeyer, F. (2013). Are consumers myopic? Evidence from new and used car purchases. *American Economic Review*, 103(1):220–56.
- Calel, R. (2020). Adopt or innovate: Understanding technological responses to cap-and-trade. *American Economic Journal: Economic Policy*, 12(3):170–201.
- Calel, R. and Dechezleprêtre, A. (2016). Environmental policy and directed technological change: evidence from the European carbon market. *Review of economics and statistics*, 98(1):173–191.
- Cameron, A. C. and Trivedi, P. K. (2013). *Regression analysis of count data*, volume 53. Cambridge university press.
- Chen, C.-W., Hu, W.-M., and Knittel, C. R. (2021). Subsidizing fuel-efficient cars: Evidence from china’s automobile industry. *American Economic Journal: Economic Policy*.
- Crabb, J. M. and Johnson, D. K. (2010). Fueling innovation: the impact of oil prices and cafe standards on energy-efficient automotive technology. *The Energy Journal*, 31(1).
- Davis, L. W. and Kilian, L. (2011). Estimating the effect of a gasoline tax on carbon emissions. *Journal of Applied Econometrics*, 26(7):1187–1214.
- Dechezleprêtre, A., Hémous, D., Olsen, M., and Zanella, C. (2021). Induced automation: evidence from firm-level patent data. *University of Zurich, Department of Economics, Working Paper*, (384).
- Dechezleprêtre, A., Neumayer, E., and Perkins, R. (2015). Environmental regulation and the cross-border diffusion of new technology: Evidence from automobile patents. *Research Policy*, 44(1):244 – 257.

- Dekker, T., Vollebergh, H. R., de Vries, F. P., and Withagen, C. A. (2012). Inciting protocols. *Journal of Environmental Economics and Management*, 64(1):45 – 67.
- Dernis, H., Guellec, D., and van Pottelsberghe, B. (2001). Using patent counts for cross-country comparison of technology output. *STI Review*, 27.
- D’Haultfœuille, X., Givord, P., and Boutin, X. (2014). The Environmental Effect of Green Taxation: The Case of the French Bonus/Malus. *The Economic Journal*, 124(578):F444–F480.
- Díaz, S., Miller, J., Mock, P., Minjares, R., Anenberg, S., and Meszler, D. (2017). Shifting gears: The effects of a future decline in diesel market share on tailpipe CO₂ and NO_x emissions in Europe. ICCT White Paper.
- Espey, M. (1998). Gasoline demand revisited: an international meta-analysis of elasticities. *Energy Economics*, 20(3):273–295.
- Financial Times (2019). Fiat Chrysler pools fleet with Tesla to avoid EU emissions fines. <https://www.ft.com/content/7a3c8d9a-57bb-11e9-a3db-1fe89bedc16e>.
- Fischer, C., Parry, I. W., and Pizer, W. A. (2003). Instrument choice for environmental protection when technological innovation is endogenous. *Journal of Environmental Economics and Management*, 45(3):523 – 545.
- Fried, S. (2018). Climate policy and innovation: A quantitative macroeconomic analysis. *American Economic Journal: Macroeconomics*, 10(1):90–118.
- 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.
- Gillingham, K. (2021). Designing fuel-economy standards in light of electric vehicles. Working Paper 29067, National Bureau of Economic Research.
- Gillingham, K. T., Houde, S., and van Benthem, A. A. (2021). Consumer myopia in vehicle purchases: Evidence from a natural experiment. *American Economic Journal: Economic Policy*, 13(3):207–38.

- Goulder, L. H. and Parry, I. W. H. (2008). Instrument choice in environmental policy. *Review of Environmental Economics and Policy*, 2(2):152–174.
- Griliches, Z. (1990). Patent statistics as economic indicators: A survey. *Journal of Economic Literature*, 28(4):1661–1707.
- Hart, R. (2019). To everything there is a season: Carbon pricing, research subsidies, and the transition to fossil-free energy. *Journal of the Association of Environmental and Resource Economists*, 6(2):349–389.
- 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.
- Helfand, G. E. (1991). Standards versus standards: The effects of different pollution restrictions. *The American Economic Review*, 81(3):622–634.
- Hepburn, C. (2006). Regulation by Prices, Quantities, or Both: A Review of Instrument Choice. *Oxford Review of Economic Policy*, 22(2):226–247.
- Holland, S. P., Mansur, E. T., and Yates, A. J. (2021). The electric vehicle transition and the economics of banning gasoline vehicles. *American Economic Journal: Economic Policy*, 13(3):316–44.
- IEA (2020). Global EV outlook 2020. IEA, Paris. <https://www.iea.org/reports/global-ev-outlook-2020>.
- Ito, K. and Sallee, J. M. (2018). The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel Economy Standards. *The Review of Economics and Statistics*, 100(2):319–336.
- Jacobsen, M. R. (2013). Evaluating us fuel economy standards in a model with producer and household heterogeneity. *American Economic Journal: Economic Policy*, 5(2):148–87.
- Jacobsen, M. R. and van Benthem, A. A. (2015). Vehicle scrappage and gasoline policy. *American Economic Review*, 105(3):1312–38.

- 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):577–598.
- Johnstone, N., Haščič, I., and Popp, D. (2010). Renewable energy policies and technological innovation: evidence based on patent counts. *Environmental and resource economics*, 45(1):133–155.
- Johnstone, N., Haščič, I., Poirier, J., Hemar, M., and Michel, C. (2012). Environmental policy stringency and technological innovation: evidence from survey data and patent counts. *Applied Economics*, 44(17):2157–2170.
- Kanbur, R. and Keen, M. (1993). Jeux sans frontières: Tax competition and tax coordination when countries differ in size. *The American Economic Review*, 83(4):877–892.
- Kellogg, R. (2018). Gasoline price uncertainty and the design of fuel economy standards. *Journal of Public Economics*, 160:14–32.
- Kemp, R. and Pontoglio, S. (2011). The innovation effects of environmental policy instruments — a typical case of the blind men and the elephant? *Ecological Economics*, 72:28 – 36.
- Kiso, T. (2019). Environmental policy and induced technological change: evidence from automobile fuel economy regulations. *Environmental and Resource Economics*, 74(2):785–810.
- Klier, T. and Linn, J. (2010). The price of gasoline and new vehicle fuel economy: Evidence from monthly sales data. *American Economic Journal: Economic Policy*, 2(3):134–53.
- Klier, T. and Linn, J. (2012). New-vehicle characteristics and the cost of the corporate average fuel economy standard. *The RAND Journal of Economics*, 43(1):186–213.
- Klier, T. and Linn, J. (2013). Fuel prices and new vehicle fuel economy—comparing the united states and western europe. *Journal of Environmental Economics and Management*, 66(2):280–300.

- Klier, T. and Linn, J. (2016). The effect of vehicle fuel economy standards on technology adoption. *Journal of Public Economics*, 133:41–63.
- 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. and Tanaka, S. (2021). Fuel economy and the price of gasoline: Evidence from fueling-level micro data. *Journal of Public Economics*, 202:104496.
- Kühlwein, J., German, J., and Bandivadekar, A. (2014). Development of test cycle conversion factors among worldwide light-duty vehicle CO₂ emission standards. ICCT White Paper.
- Lanjouw, J. O. and Mody, A. (1996). Innovation and the international diffusion of environmentally responsive technology. *Research Policy*, 25(4):549–571.
- Leard, B., Linn, J., and McConnell, V. (2017). Fuel prices, new vehicle fuel economy, and implications for attribute-based standards. *Journal of the Association of Environmental and Resource Economists*, 4(3):659–700.
- Leard, B., Linn, J., and Zhou, Y. C. (2021). How Much Do Consumers Value Fuel Economy and Performance? Evidence from Technology Adoption. *The Review of Economics and Statistics*, pages 1–45.
- Lee, J., Veloso, F. M., and Hounshell, D. A. (2011). Linking induced technological change, and environmental regulation: Evidence from patenting in the U.S. auto industry. *Research Policy*, 40(9):1240 – 1252.
- Li, S., Linn, J., and Muehlegger, E. (2014). Gasoline taxes and consumer behavior. *American Economic Journal: Economic Policy*, 6(4):302–42.
- Martinez, C. (2010). Insight into different types of patent families. *OECD Science, Technology and Industry Working Papers*, No. 2010/02.
- Montero, J. P. (2002). Permits, standards, and technology innovation. *Journal of Environmental Economics and Management*, 44(1):23 – 44.

- Noailly, J. (2012). Improving the energy efficiency of buildings: The impact of environmental policy on technological innovation. *Energy Economics*, 34(3):795 – 806.
- Noailly, J. and Smeets, R. (2015). Directing technical change from fossil-fuel to renewable energy innovation: An application using firm-level patent data. *Journal of Environmental Economics and Management*, 72:15–37.
- OECD (2010). *Taxation, Innovation and the Environment*. OECD Publishing, Paris.
- Parry, I. W. and Small, K. A. (2005). Does Britain or the United States have the right gasoline tax? *American Economic Review*, 95(4):1276–1289.
- Parry, I. W. H., Walls, M., and Harrington, W. (2007). Automobile externalities and policies. *Journal of Economic Literature*, 45(2):373–399.
- Popp, D. (2002). Induced innovation and energy prices. *American Economic Review*, 92(1):160–180.
- Popp, D. (2019). Environmental policy and innovation: A decade of research. Working Paper 25631, National Bureau of Economic Research.
- Popp, D., Newell, R. G., and Jaffe, A. B. (2010). Energy, the environment, and technological change. In *Handbook of the Economics of Innovation*, volume 2, pages 873–937. Elsevier.
- Probst, B., Touboul, S., Glachant, M., and Dechezleprêtre, A. (2021). Global trends in the innovation and diffusion of climate change mitigation technologies.
- Proost, S. and Van Dender, K. (2011). What Long-Term Road Transport Future? Trends and Policy Options. *Review of Environmental Economics and Policy*, 5(1):44–65.
- Requate, T. (2005). Dynamic incentives by environmental policy instruments—a survey. *Ecological Economics*, 54(2):175 – 195.
- Reynaert, M. (2021). Abatement Strategies and the Cost of Environmental Regulation: Emission Standards on the European Car Market. *The Review of Economic Studies*, 88(1):454–488.

- Reynaert, M. and Sallee, J. M. (2021). Who benefits when firms game corrective policies? *American Economic Journal: Economic Policy*, 13(1):372–412.
- Schankerman, M. and Pakes, A. (1986). Estimates of the value of patent rights in European countries during the post-1950 period. *The Economic Journal*, 96(384):1052–1076.
- Scherer, F. M. (1965). Firm size, market structure, opportunity, and the output of patented inventions. *American Economic Review*, 55(5):1097–1125.
- Smulders, S. and de Nooij, M. (2003). The impact of energy conservation on technology and economic growth. *Resource and Energy Economics*, 25(1):59 – 79.
- Springel, K. (2021). Network externality and subsidy structure in two-sided markets: Evidence from electric vehicle incentives. *American Economic Journal: Economic Policy*.
- Van Soest, D. P., List, J. A., and Jeppesen, T. (2006). Shadow prices, environmental stringency, and international competitiveness. *European Economic Review*, 50(5):1151–1167.
- Vollebergh, H. R. (2007). Differential impact of environmental policy instruments on technological change: a review of the empirical literature. *Available at SSRN 991612*.
- Vollebergh, H. R. (2010). Fuel taxes, motor vehicle emission standards and patents related to the fuel-efficiency and emissions of motor vehicles. COM/ENV/EPOC/CTPA/CFA(2008)32/FINAL. Environmental Directorate, OECD, Paris.
- Vollebergh, H. R. J. and Van der Werf, E. (2014). The Role of Standards in Eco-innovation: Lessons for Policymakers. *Review of Environmental Economics and Policy*, 8(2):230–248.
- Weitzman, M. L. (1974). Prices vs. quantities. *The Review of Economic Studies*, 41(4):477–491.

Appendix

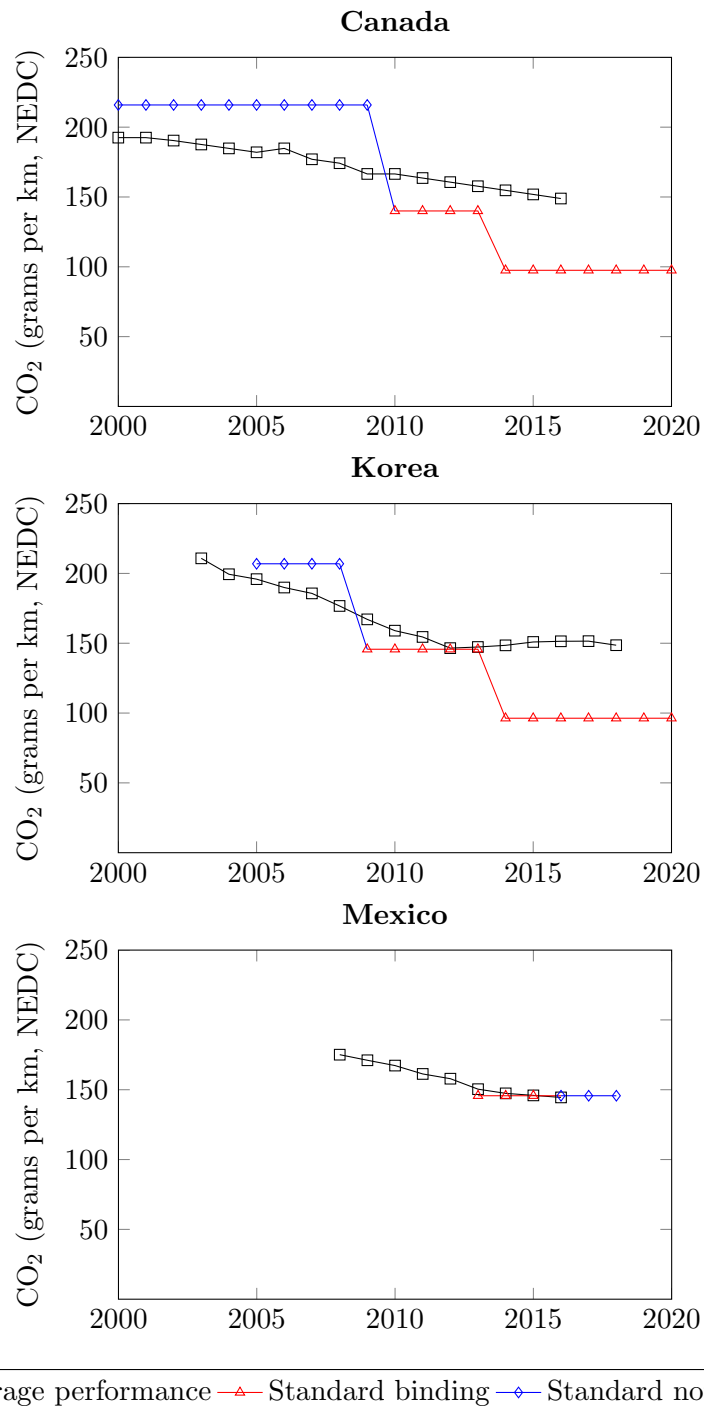


Figure 5: Average performance and standards converted to CO₂ emissions per kilometer (NEDC). Standards are defined as the most recently announced target level. They are binding if the target is mandatory and lower than average performance. Source: ICCT.




Figure 6: Stringency for Canada, Korea and Mexico defined as the average required annual reduction in CO₂ emissions per kilometer (NEDC) to meet the most recently announced target, computed following equation (1).