ABSTRACT

Title of dissertation: BEHAVIORAL RESPONSE TO ENVIRONMENTAL TAXATION: EVIDENCE FROM THE TRANSPORTATION SECTOR

Davide Cerruti, Doctor of Philosophy, 2016

Dissertation directed by: Professor Anna Alberini
Professor Roberton Williams III
Department of Agricultural and Resource Economics

This dissertation analyzes how individuals respond to the introduction of taxation aimed to reduce vehicle pollution, greenhouse gases and traffic.

The first chapter analyzes a vehicle registration tax based on emissions of carbon dioxide (CO_2), a major greenhouse gas, adopted in the UK in 2001 and subject to major changes in the following years. I identify the impact of the policy on new vehicle registrations and carbon emissions, compared to alternative measures. Results show that consumers respond to the tax by purchasing cleaner cars, but a carbon tax generating the same revenue would further reduce carbon emissions.

The second chapter looks at a pollution charge (polluting vehicles pay to enter the city) and a congestion charge (all vehicles pay) adopted in 2008 and 2011 in Milan, Italy, and how they affected the concentration of nitrogen dioxides (NOx). I use data from pollution monitoring stations to measure the change between areas adopting the tax and other areas. Results show that in the first quarter of their

introduction, both policies decreased NOx concentration in a range of -8% and -5%, but the effect declines over time, especially in the case of the pollution charge.

The third chapter examines a trial conducted in 2005 in the Seattle, WA, area, in which vehicle trips by 276 volunteer households were recorded with a GPS device installed in their vehicles. Households received a monetary endowment which they used to pay a toll for each mile traveled: the toll varied with the time of the day, the day of the week and the type of road used. Using information on driving behavior, I show that in the first week a $0.10 toll per mile reduces the number of miles driven by around 7%, but the effect lasts only few weeks at most. The effect is mainly driven by a reduction in highway miles during trips from work to home, and it is strongly influenced by past driving behavior, income, the size of the initial endowment and the number of children in the household.

BEHAVIORAL RESPONSE TO ENVIRONMENTAL TAXATION: EVIDENCE FROM THE TRANSPORTATION SECTOR

by

Davide Cerruti

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2016

Advisory Committee:
Professor Anna Alberini, Co-Advisor
Professor Roberton Williams III, Co-Advisor
Professor Lint Barrage
Professor Qingbin Cui
Professor Sébastien Houde

Foreword

The first chapter of the following dissertation is a jointly authored work with Anna Alberini and Joshua Linn. The Dissertation Committee aknowledges that Davide Cerruti made substantial contributions to the relevant aspects of the chapter.

Acknowledgments

A doctoral dissertation is inevitably the end of a long and sometimes tortuous journey, starting even before the actual enrollement in a PhD program. This work is no exception. For this reason, I wish to thank all the people who supported me, from the very beginning to the very end.

First of all, I thank my advisors Anna Alberini and Roberton Williams III for their invaluable support, advice and conversations. Their example, intellectual curiosity and insight were fundamental in my training as academic economist and I owe them a large share of what I have learned during the time spent in the program. The remaining members of my dissertation committee, Lint Barrage, Sébastien Houde and Qingbin Cui provided extremely useful comments to the previous versions of this dissertation, for which I am very grateful.

During my time at Maryland, I had the opportunity to work with various faculty members as Research or Teaching Assistant. What I learned from those experiences has been very useful and will continue to be as I proceed in my career, and for that I thank Lint Barrage, Vivian Hoffmann, Pamela Jakiela, Howard Leathers and Lars Olson.

I also would like to thank other faculty members of the department, namely Erik Lichtenberg, Ted McConnell and Jorge Holzer, who organized and attended regularly the department environmental economics workshop, which was a terrific opportunity to present early draft of my work to a friendly and attentive audience.

In terms of the single chapters of the dissertation, I thank Anna Alberini and

Josh Linn for their contribution to the first chapter of the dissertation, which started as a coauthored project. The second chapter benifited from the information from analysts at AMMA and ARPA. Analysts at NREL and the Puget Sound Regional Council provided useful data for the writing of the third chapter.

The financial support from the department and from Resources For the Future, especially during the later stages of the program, is greatly appreciated.

I owe a great debt of gratitude to the researchers of the Electricity Policy Research Group at the University of Cambridge, and in particular to Michael Grubb, who believed in me and gave the opportunity to work as his Research Assistant from 2010 to 2011, and to David Newbery and Jevgenijs Steinbuks, who provided as well a great deal of support and advice in choosing my future career path. It is safe to say that without them my journey as PhD student and environmental economist would not have even started.

My time spent in the program gave me the opportunity to meet a terrific group of classmates and colleagues, many of which became dear friends. I want to thank in particular Mehrab Bakhtiar, Andrew Brudevold-Newman, Jaclyn Evans, Jeff Ferris, Paige Gance, Zheng Jen He, Marie Hyland, Boya Liu, Mark Miller, Ian Page, Yuandong Qi, Monica Saavoss, Elina Tselepidakis, Inge van den Bijgaart, Jikun Wang, Dan Werner, Youpei Yan, Ziyan Yang and Yichen Christy Zhou, for their company, friendship and support.

Last, but not least, I would like to deeply thank my family for their unconditional love and for their constant encouragement to follow my interests and never give up. *Grazie di tutto*!

Table of Contents

List of Tables

List of Figures

List of Abbreviations

CO_2	carbon dioxide
NO_x	nitrogen oxides
O_3	ozone
PM_{10}	particulate matter
AMMA	Agenzia Milanese Mobilità Ambiente
ARPA	Agenzia Regionale per la Protezione dell'Ambiente
ATM	Automated Teller Machine
BgBs	provinces of Bergamo and Brescia
CA	California
CAFE	Corporate Average Fuel Economy
CCTV	Closed-Circuit Television
DC	District of Columbia
DOT	Department of Transportation
GBP	British Pound
GPS	Global Positioning System
LEZ	Low Emission Zones
LPG	Liquefied Petroleum Gas
MiMb	provinces of Milan and Monza e Brianza
NHTS	National Household Travel Survey
NREL	National Renewable Energy Laboratory
NTS	UK National Travel Survey
OBU	On-Board Unit
UK	United Kingdom
US	United States
VED	Vehicle Excise Duty
VMT	Vehicle Miles Traveled
WA	Washington

Chapter 1: Charging Drivers by the Pound: The Effect of the VED System in the UK

1.1 Introduction

New vehicle fuel economy standards are becoming increasingly common around the world. Transportation accounts for about 15 percent of global greenhouse gas emissions and 25 percent of carbon dioxide (CO_2) emissions, and the standards aim to reduce fuel consumption and greenhouse gas emissions from passenger vehicles. Countries adopting standards account for about three quarters of global passenger vehicle fuel consumption, and include developed and developing countries.

Many countries are redesigning their vehicle tax systems to complement fuel economy standards. This approach is particularly common in Europe, where vehicles are subject to a sales tax at the time of the purchase and annual registration (circulation) fees. Historically a vehicles taxes depended on its weight or engine size, both of which tend to be correlated with fuel consumption rate and the CO_2 emissions rate, but starting in the early 2000s several European countries have tied taxes more directly to the CO_2 emissions rate. For example, France offers subsidies for purchasing vehicles with low CO_2 emissions rates and imposes taxes on purchasing vehicles with high emissions rates. In Germany a vehicles circulation tax increases linearly with its emissions rate. The main goal of these taxes is to provide consumers a price signal to encourage car buyers to purchase vehicles with lower CO_2 emissions rates, which implies better fuel economy.

Several recent studies have demonstrated that linking purchase and circulation taxes to CO_2 emissions rates reduces the average emissions rate of new vehicles purchased (Adamou et al., 2012, Ciccone, 2014, D'Haultfuille et al., 2014, Huse and Lucinda, 2014, Konishi and Meng, 2014, Klier and Linn, 2015, Alberini and Bareit, 2016). One key question is how annual registration fees compare with fuel taxes. Fuel taxes affect vehicle purchase and use, and penalize total CO_2 emissions (Busse et al., 2013, Grigolon et al., 2015). They are thus often thought as reasonably efficient at controlling the variable external costs of driving (emissions, congestions, road wear and tear, etc.). On the other hand, European countries already impose relatively high fuel taxes, perhaps higher than the efficient level (Parry and Small, 2005), and, given current tax levels, vehicle taxes may be preferred to fuel taxes for political reasons. It is unclear from the existing literature whether fuel taxes or vehicle taxes are more effective at reducing passenger vehicle CO_2 emissions rates.

The second question is how to structure the tax system. Some countries, such as Sweden and Germany, impose taxes that increase linearly with the vehicles CO_2 emissions rate. Other countries, such as France and the United Kingdom, impose one-time and annual taxes respectively that vary discretely with the emissions rate. An efficient CO_2-based vehicle tax would tax all vehicles in proportion to their lifetime CO_2 emissions, but this may not be administratively feasible. Instead, policy makers may aim to design the tax system to most effectively reduce CO_2 emissions rates. The literature has typically estimated the average effect of CO_2 taxes on vehicle sales, providing little information about how to achieve this objective. We analyze an annual vehicle registration tax (Vehicle Excise Duty) based on CO_2 emissions rates that the UK adopted in 2001 and changed substantially in the following years. The taxes are imposed each year the vehicle is owned. In 2009, the ranged from £0 for vehicles with emissions rates below 100 grams of CO_2/kilometer (g CO_2/km) to £405 for vehicles with emissions rates above 255 g CO_2/km (for a gasoline-powered car 255 g CO_2/km is equivalent to about 22 miles per gallon). The UK registration tax system thus penalizes heavily vehicles with very high emissions rates and provides discounts to vehicles with very low emissions rates. The tax a particular vehicle faces changed several times since the inception of the system. Those changes were more drastic in the case of high and low polluting vehicles, while the tax for vehicles with moderate emissions rates remained relatively stable. Using a highly disaggregated dataset of UK monthly new car registrations and characteristics, we estimate the effects of the taxes on new car registrations. We also predict the effect of switching from the VED to a tax proportional to carbon emission rates and to a carbon tax, which depend on both the vehicles emissions rate and the miles driven. We find that a tax that is proportional to emissions rate would actually increase carbon emissions from new vehicles by +0.12%, whereas a carbon tax would result in a reduction by 3.71%. The reminder of the paper is organized as follows. Section 1.2 provides a background of the vehicle registration tax scheme in the UK. Section 1.3 describes the vehicle registration data. Section 1.4 shows the model and the identification strategy. Section 1.5 presents the results. Section 1.6 concludes.

1.2 Background

Prior to March 1, 2001, the annual registration fee for passenger vehicles in the UKcalled the Vehicle Excise Duty (VED)depended on the size of the engine. Cars with larger engine capacity paid a higher registration fee. Starting March 1, 2001, vehicles were placed in CO_2 emissions bands, and different bands paid different VED amounts: The higher the emissions rate, the greater the VED amount.[1] Special provisions were established for smaller-engined cars that were first registered before March 1, 2001, and exceptions to this rule were granted to cars that were first registered before 1973. Initially, there were a total of four bands. Band A included all cars with emissions rates below 150 g/km, band B cars with emissions rates ranging between 151 and 165 g/km, band C vehicles in the 166-185 g/km range, and band D all others, namely 186 g/km and higher. In March 2002, band A was broken into bands AA (less than 120 g/km) and A (121-150 g/km), while all others remained unchanged. In March 2003, vehicles that emitted less than 100 g/km were placed in band AAA, those between 101 and 120 g/km in band AA, and band A continued to apply for emissions rates between 121 and 150 g/km (see table 1.1 and 1.2).

In April 2005, the bands were renamed with no changes to the emissions range for each band (see table 1.3 and 1.4). In March 2006, band F was split to form a new band F (186-225 g/km) and band G, which includes vehicles with emissions rates 226 g/km and higher. Major revisions to the system were done in May 2009,

[1] Emissions are expressed as grams of CO_2 per kilometer.

Table 1.1: VED bands classification for new cars between 2001 and 2003

<table>
<tr><th>CO_2 emissions rate</th><th>Mar-01</th><th>Mar-02</th><th>Mar-03</th></tr>
<tr><td>up to 100 g/km</td><td rowspan="3">A</td><td rowspan="2">AA</td><td>AAA</td></tr>
<tr><td>101-120</td><td>AA</td></tr>
<tr><td>121-150</td><td>A</td><td>A</td></tr>
<tr><td>151-165</td><td>B</td><td>B</td><td>B</td></tr>
<tr><td>166-185</td><td>C</td><td>C</td><td>C</td></tr>
<tr><td>186 and higher</td><td>D</td><td>D</td><td>D</td></tr>
</table>

Table 1.2: VED rates for new cars between 2001 and 2003

<table>
<tr><th>CO_2 emissions rate</th><th>Mar-01</th><th>Mar-02</th><th>Mar-03</th></tr>
<tr><td>up to 100 g/km</td><td rowspan="6">£100</td><td rowspan="3">£70</td><td>£65</td></tr>
<tr><td>101-110</td><td rowspan="2">£75</td></tr>
<tr><td>111-120</td></tr>
<tr><td>121-130</td><td rowspan="3">£100</td><td rowspan="3">£105</td></tr>
<tr><td>131-140</td></tr>
<tr><td>141-150</td></tr>
<tr><td>151-165</td><td>£120</td><td>£120</td><td>£125</td></tr>
<tr><td>166-175</td><td rowspan="2">£140</td><td rowspan="2">£140</td><td rowspan="2">£145</td></tr>
<tr><td>176-185</td></tr>
<tr><td>186-200</td><td rowspan="4">£155</td><td rowspan="4">£155</td><td rowspan="4">£160</td></tr>
<tr><td>201-225</td></tr>
<tr><td>226-255</td></tr>
<tr><td>256 and higher</td></tr>
</table>

when the existing bands were redefined (bands A-I use 10-gram intervals) and new bands were added. As shown in table 1.3, the highest band is M, with 256 g/km and higher emissions rates.[2]

Table 1.3: VED bands classification for new cars after April 2005

<table>
<tr><th>CO_2 emissions rate</th><th>Apr-05</th><th>Mar-06</th><th>May-09</th></tr>
<tr><td>up to 100 g/km</td><td>A</td><td>A</td><td>A</td></tr>
<tr><td>101-110</td><td rowspan="2">B</td><td rowspan="2">B</td><td>B</td></tr>
<tr><td>111-120</td><td>C</td></tr>
<tr><td>121-130</td><td rowspan="3">C</td><td rowspan="3">C</td><td>D</td></tr>
<tr><td>131-140</td><td>E</td></tr>
<tr><td>141-150</td><td>F</td></tr>
<tr><td>151-165</td><td>D</td><td>D</td><td>G</td></tr>
<tr><td>166-175</td><td rowspan="2">E</td><td rowspan="2">E</td><td>H</td></tr>
<tr><td>176-185</td><td>I</td></tr>
<tr><td>186-200</td><td rowspan="4">F</td><td rowspan="2">F</td><td>J</td></tr>
<tr><td>201-225</td><td>K</td></tr>
<tr><td>226-255</td><td rowspan="2">G</td><td>L</td></tr>
<tr><td>256 and higher</td><td>M</td></tr>
</table>

In sum, the number of bands was increased over the years, the thresholds set between bands changed, and the registration fee amounts were changed over the years, making the difference between low-emissions vehicle and high-emissions vehicle starker and starker. Until 2006, Diesel cars paid a slightly higher VED (between 5 and 15 pound more) for the same emissions rate.

[2] Starting in 2006, cars that were first registered after 2001 but before the current budget fiscal year date may be given a slightly different tax schedule than the one shown in table 1.3, which refers to new cars.

Table 1.4: VED rates for new cars after April 2005

<table>
<tr><th>CO_2 emissions rate</th><th>Apr-05</th><th>Mar-06</th><th>May-07</th><th>Mar-08</th><th>May-09</th><th>Apr-10</th></tr>
<tr><td>up to 100 g/km</td><td>£65 (£75)</td><td>£0</td><td>£0</td><td>£0</td><td>£0</td><td rowspan="4">£0</td></tr>
<tr><td>101-110</td><td rowspan="2">£75 (£85)</td><td rowspan="2">£40 (£50)</td><td rowspan="2">£35</td><td rowspan="2">£35</td><td rowspan="2">£35</td></tr>
<tr><td>111-120</td></tr>
<tr><td>121-130</td><td rowspan="3">£105 (£115)</td><td rowspan="3">£100 (£110)</td><td rowspan="3">£115</td><td rowspan="3">£120</td><td rowspan="2">£120</td></tr>
<tr><td>131-140</td><td>£110</td></tr>
<tr><td>141-150</td><td>£125</td><td>£125</td></tr>
<tr><td>151-165</td><td>£125 (£135)</td><td>£125 (£135)</td><td>£140</td><td>£145</td><td>£150</td><td>£155</td></tr>
<tr><td>166-175</td><td rowspan="2">£150 (£160)</td><td rowspan="2">£150 (£160)</td><td rowspan="2">£165</td><td rowspan="2">£170</td><td rowspan="2">£175</td><td>£250</td></tr>
<tr><td>176-185</td><td>£300</td></tr>
<tr><td>186-200</td><td rowspan="4">£165 (£170)</td><td rowspan="2">£190 (£195)</td><td rowspan="2">£205</td><td rowspan="2">£210</td><td rowspan="2">£215</td><td>£425</td></tr>
<tr><td>201-225</td><td>£550</td></tr>
<tr><td>226-255</td><td rowspan="2">£210 (£215)</td><td rowspan="2">£300</td><td rowspan="2">£400</td><td rowspan="2">£405</td><td>£750</td></tr>
<tr><td>256 and higher</td><td>£950</td></tr>
</table>

VED rates for gasoline vehicles. If rates are different for diesel vehicles, they are reported in parenthesis.

Changes in April 2010 apply only for the first year.

1.3 Data

Our original data source is a large dataset compiled by R.L. Polk & Co where the unit of observation is a make-model-trim-variant.[3] For each such unit, we have the number of new registrations in each month from January 2005 to October 2010.

In this paper attention is restricted to gasoline or diesel passenger vehicles, and so we exclude vans and a number of other vehicles (bus, city van, combi van, panel van with double cabin, recreational van, and rigid van) from the source dataset. The excluded vehicles account for only 1.30% of the original sample. We are also forced to drop from the analysis variants with no price information (0.27% of the sample).

Although the original data are at the monthly level, we choose to use the policy period as the time interval for our analysis. The policy period is essentially the fiscal year, and comes with a set of VED bands and rates. The policy periods (a total of 6) are reported in table 1.5. Among other things, table 5 shows that, in forming the policy periods, we deleted a month if the VED bands and rates are effective from a date in the middle of that month (rather than the first of the month).

Because we tally the number of sales for each make-model-trim-variant over each policy period, we are forming a panel dataset where the cross-sectional unit is

[3] A unique variant is an observation with a given make-model-trim, number of doors, market segment, body type, whether two- or four-wheel drive, transmission type and number of gears, fuel type, engine size, weight, length, height, number of cylinders, horsepower, fuel consumption rate and price.

Table 1.5: Policy periods

policy period	beginning	end	duration in months
1	May-05	Feb-06	9.2
2	Apr-06	Feb-07	10.2
3	Apr-07	Feb-08	10.2
4	Apr-08	Apr-09	12.2
5	May-09	Mar-10	10.1
6	Apr-10	Oct-10	6.1

Policy periods used in the analysis. The duration is expressed in months (30 days)

the make-model-trim-variant. This is an unbalanced panel because some of these cross-sectional units enter or exit the market during our study period. In total, we have 55 makes, 507 different make-models, 3130 make-model-trims, 36110 different make-model-trim-variants and the maximum length of the panel is 6. Each model belongs to a unique market segment (Mini, Small, Lower-Medium, Medium, Upper-Medium, Large).

Towards the end of our study period, new vehicle registrations declined sharply, especially from 2008 (figure 1.1). This is likely due to the beginning of the recession which had a negative effect on automobile sales across Europe. In our sample, a make-model has an average of about 21,245 new registrations, while a single variant has an average of about 299. About 12.52% of the variants had no new registrations at all within a given policy period.

In terms of variation in the VED rates within make-model and policy period, on average each make-model has 28.23 different variants, and 6.45 variants with

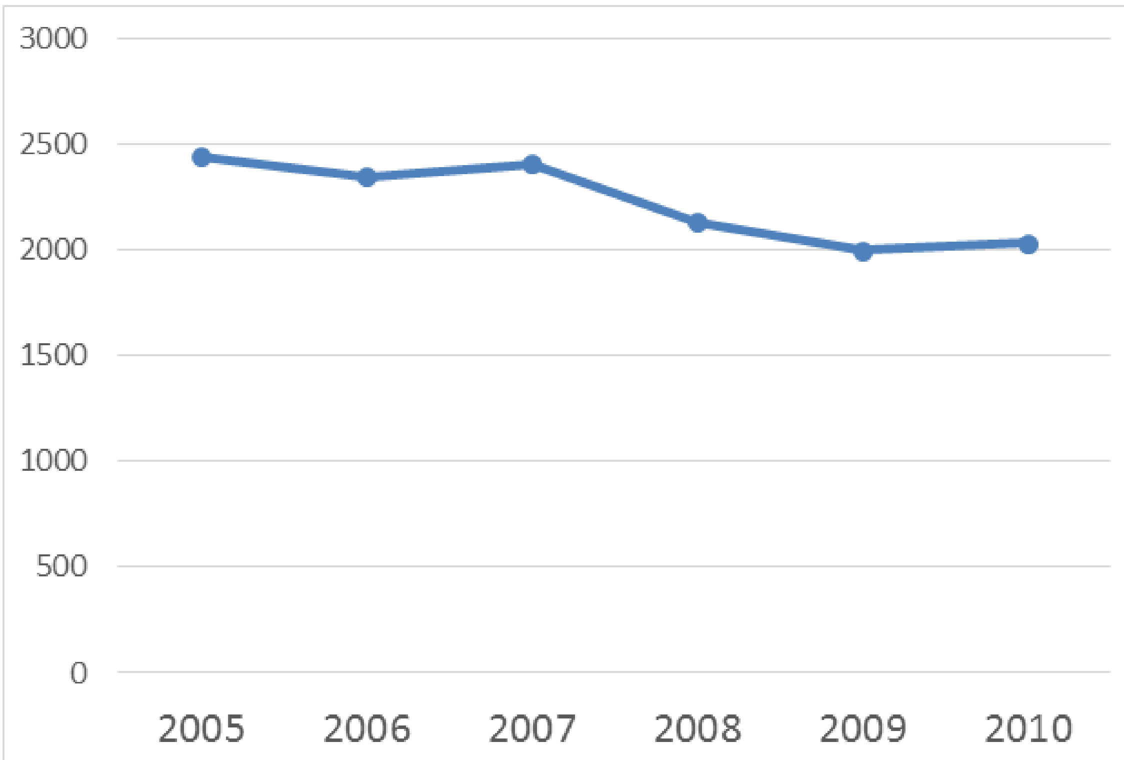

Figure 1.1: New passenger vehicle registrations in the UK between 2005 and 2010 (thousand vehicles).

Source: European Automobile Manufaturers Association (ACEA) different VED rates. Only 6.77% of the make-models have only one variant. Having enough variation in VED rates within make-model and policy period is essential to identify the effect of the VED on vehicle registrations.

On average, vehicles became cleaner over time: figure 1.2 shows the trends in terms of average CO_2 emissions rates from the variants in our dataset. Figure 1.2 is based on data not weighted by sales, and shows that in less than six years, average emissions rates declined by about 30 g/km, from about 190 g/km to a little over 150 g/km.[4]

In order to estimate the relationship between vehicle characteristics, driving costs and miles driven, we use data from the UK National Travel Survey (NTS). The UK NTS is conducted each year, and collects information from households about individual trips taken during a specified period (travel diaries component), car ownership and characteristics (including miles driven each year and odometer reading), and household sociodemographics. We have 6 waves of the UK NTS, from 2005 to 2010 so as to match the vehicle sales dataset, covering 81,855 households.[5] We consider only gasoline and diesel cars owned by households with information on carbon emissions (48.57% of all vehicles and 30.56% of all households covered by the survey).[6]

[4]The reduction in average emission rates comes from improvements in vehicle technology. We find unlikely that such improvement is caused by manufacturers direct response to changes in the VED, because UK is just a part of the EU car marker and other European countries have different tax schemes.

[5]We form a multi-year cross-sections type of dataset, as the households interviewed as part of the UK NTS are different every year.

[6]Before the introduction of the VED in 2001 information on carbon emission rates was not

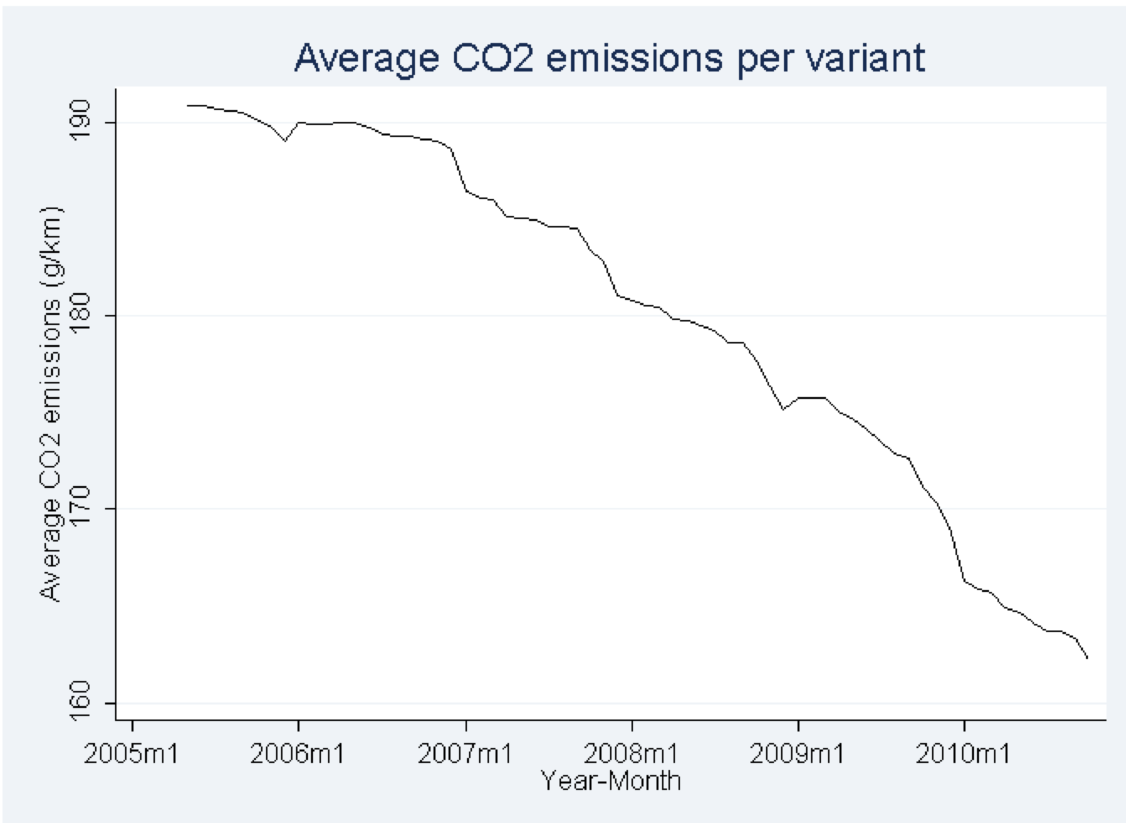

Figure 1.2: Average CO_2 emissions per variant. Values not sales-weighted.

The NTS dataset reports the exact CO_2 emissions rate of each vehicle, as specified by the automaker, when available, but does not contain information on fuel economy. We construct the vehicles fuel economy using data from the European Environmental Agency on passenger cars emissions and fuel economy, taking advantage of the fact that the emissons are perfectly proportional to the vehicles fuel consumption rate (in liters per 100 km) and that the proportional factor is different for diesel and gasoline (see Appendix A.1).

mandatory for new cars. For this reason, about 62.01% of the passenger cars in the NTS do not have information on carbon emissions rates and were dropped from the final sample.

1.4 Empirical strategy

We tally the number of sales for each policy period and fit the regression equation:

$$\ln \text{Sales}_{imt} = \alpha + \beta \text{VEDTAX}_{imt} + \gamma P_{imt} + \delta \text{FUELCOST}_{imt} + \theta x_{im} + \epsilon_{imt} \quad (1.1)$$

where i denotes the trim-variant, m the make-model, t the policy period, and Sales the number of units sold normalized by the number of months in the policy period.[7] In the right-hand side of equation 1.1, VEDTAX is the annual registration tax (in hundred 2005 GBP), P is the manufacturer-suggested retail price of the vehicle, and FUELCOST is the fuel cost per 100 km, which depends on the price of gasoline or diesel fuel and the fuel economy of the car.[8] Car prices and fuel costs are expressed in 2005 constant GBP. Vector x contains vehicle attributes thought to influence sales, namely engine size, horsepower, weight, length and height, body

[7] If a make-model-trim-variant is introduced after the beginning of the policy period, or is withdrawn from the market during that policy period, we normalize the sales by the number of months that the make-model-trim-variant is in actual existence within that policy period. In some cases the VED changes occurred in the middle of the month. We remove from the original dataset the months in which this occurs, as we cannot assign the registration to a specific tax rate, and adjust the normalization of the sales accordingly. The months removed are March 2006, March 2007, and March 2008.

[8] In the last year of our study period, the structure of the VED was changed so that for certain cars the amount due the first year after registration was different from that due in subsequent years. If that is the case, we simply use the amount due in the first year.

type dummies, number of doors, type of transmission and number of gears, gasoline or diesel fuel, cylinders, and whether two- or four-wheel drive.

We interpret equation 1.1 as a reduced-form model. The inclusion of make-model-by-time fixed effects (the αs) means that we rely on variation in the VED across models within a make and over time to identify the effect of the VED itself. This is akin to assuming that tradeoffs between the VED and other attributes within a make-model determine the final car purchase after the consumer has first selected the make and model.

We first estimate equation 1.1 by ordinary least squares, assuming that all car attributes (including price) and the VED band and tax amounts are exogenous. Earlier analyses with similar models allow for the price to be endogenous, and use instrumental variable techniques to obtain consistent estimates of the coefficient on price (Berry, 1994, Berry et al., 1995, Vance and Mehlin, 2009, Adamou et al., 2012, 2014, Huse and Lucinda, 2014, Konishi and Meng, 2014, Alberini and Bareit, 2016). Since we are more interested in the coefficient on the VED registration tax than in that on price, we simply re-estimate equation 1.1 after omitting price from set of regressors.

We perform a number of robustness checks. For example, since our identification strategy relies on the variation in emissions rates and hence the VED tax within a make, we re-run equation 1.1 after (i) including an interaction between the VED amount and the number of trim-variants within a make-model, and (ii) removing the models the top 1% distribution of all sales.

One concern whenever one exploits variation in taxes and/or the specifics of a government program is whether these are anticipated by the public and car purchases in the period preceding a new tax or policy regime reflect peoples attempts to take advantage of loopholes or simply avoid higher tax rates. Based on our reading of Her Majestys Treasury documentation of each years budget and on examining news coverage about the budget and VED debate prior to the final budget approval, we believe that anticipation effects are unlikely to be important in this setting (see Appendix A.2). For good measure, however, we re-run equation 1.1 with quarterly data and after dropping the month preceding and the month following the time when the new VED becomes effective.

1.5 Results

Regression results are presented in table 1.6. Table 1.7 shows the elasticity derived from the coefficient of table 1.6. All standard errors are clustered at the market segment level. (For good measure, we also computed standard errors clustered at the make level, and found that they are very similar.)

The first column of table 1.6 shows the key coefficients from our main specification using equation (1). We run various robustness checks on our main results, which are found in columns (2)-(5) of table 1.6. Specifically, column (2) adds make-time fixed effects to equation (1). To check if the vehicle price endogeneity has an effect on fuel cost and VED semi-elasticities, column (3) omits the price from the equation. Column (4) drops the top 1% of vehicle models in terms of sales to make sure our results are not driven exclusively by few, very popular models. In column (5) we drop the month before and after any change in VED changes to check that our results are not driven by anticipation effects. Finally, in column (6) we add an interaction term between the VED amount and the number of variants within a model.

In general, we found that the coefficients on our key regressors were similar across columns (2)-(5). Our discussion of the results below is thus based on column (1). The results from the specification in column (6) suggest that consumers respond more to the VED if a model has a large number of trim-variants.

Next, we compare the effect on carbon emissions and registrations of the VED with two alternative policies: a tax proportional to the carbon emission rate of the

Table 1.6: Main results, semi-elasticities

	(1)	(2)	(3)	(4)	(5)	(6)
PRICE	−5.95e−05***	7.45e−06		−5.94e−05***	−5.02e−05***	−6.07e−05***
	(9.34e−06)	(4.74e−06)		(1.02e−05)	(1.07e−05)	(9.33e−06)
FUELCOST	−0.196***	−0.212***	−0.182***	−0.172***	−0.267***	−0.202***
	(0.0561)	(0.0456)	(0.0601)	(0.0532)	(0.0633)	(0.0542)
VEDTAX	−0.261***	−0.244***	−0.259***	−0.269***	−0.238***	−0.093*
	(0.0575)	(0.0527)	(0.0577)	(0.0512)	(0.0572)	(0.0550)
VEDTAX						−0.0032***
∗ VARIANTS						(0.0009)
Observations	55,811	55,811	55,811	50,705	53,448	55,811
R−squared	0.391	0.273	0.389	0.399	0.399	0.392

Results using model 1.1. The dependent variable is the log of the normalized number of units sold. Only coefficients for sale price, fuel cost and VED are reported. Column 1: main specification. Column 2: using make-time fixed effects. Column 3: without sales price. Column 4: dropping top 1% models sold. Column 5: dropping months before and after changes in VED. Robust standard errors in parentheses, clustered by vehicle segment. Column 6: Interaction between VED rates and number of variants within make-model. P-values: *** p¡0.01, ** p¡0.05, * p¡0.1

Table 1.7: Main results, elasticities

	(1)	(2)	(3)	(4)	(5)	(6)
Price elasticity	-0.885^{***}	0.111		-0.912^{***}	-0.743^{***}	
	(0.1390)	(0.0705)		(0.1560)	(0.1580)	
Fuel cost elasticity	-1.208^{***}	-1.303^{***}	-1.121^{***}	-1.077^{***}	-1.643^{***}	
	(0.3450)	(0.2800)	(0.3700)	(0.3330)	(0.3900)	
VED elasticity	-0.422^{***}	-0.394^{***}	-0.418^{***} -0.447^{***}	-0.383^{***}		
	(0.0930)	(0.0852)	(0.0933)	(0.0850)	(0.0918)	
Observations	55,811	55,811	55,811	50,705	53,448	

Elasticities calculates using the coefficients of table 1.6 and the sample average. Column 1: main specification. Column 2: using make-time fixed effects. Column 3: without sales price. Column 4: dropping top 1% models sold. Column 5: dropping months before and after changes in VED. Robust standard errors in parentheses, clustered by vehicle segment. P-values: *** p¡0.01, ** p¡0.05, * p¡0.1

car and a carbon tax to be paid as a lump sum at the end of each year. We wish to emphasize the difference between the two: With the former, the same amount is paid regardless of how many miles a car is driven, while a driver would be able the reduce his liability under the latter scheme by purchasing a car that emits less and/or driving fewer miles.

We use a proportional tax rate of 0.83 GBP per g CO_2/km and a carbon tax of 72 GBP per ton CO_2. These rates are chosen such that the total revenue collected from these taxes will be equal to the revenue under the VED (see Appendix A.3).

Comparing the total real tax amount paid during the vehicle lifetime shows the difference between the three policies. Figure 1.3 compares VED and the proportional tax: a proportional tax would increase the tax payment for clean cars, which generally enjoy especially low rates under the current VED scheme. High-polluting vehicles would pay more or less under a proportional tax depending on the date of the first registration: new vehicle registered in 2005 would pay more under the proportional tax, but vehicles registered in 2010 would pay less. That is because the VED rates for high polluting vehicles increased dramatically between 2005 and 2010.

Figure 1.4 shows similar results with the carbon tax, with the difference that the amount to pay under a carbon tax is also based on mileage. We use the mileage equation illustrated in Appendix A.1 to take into account the change in mileage due to the introduction of the carbon tax.

Again, under the carbon tax clean vehicles end up paying more than under the current VED scheme, and most of the very polluting vehicles pays much more, in

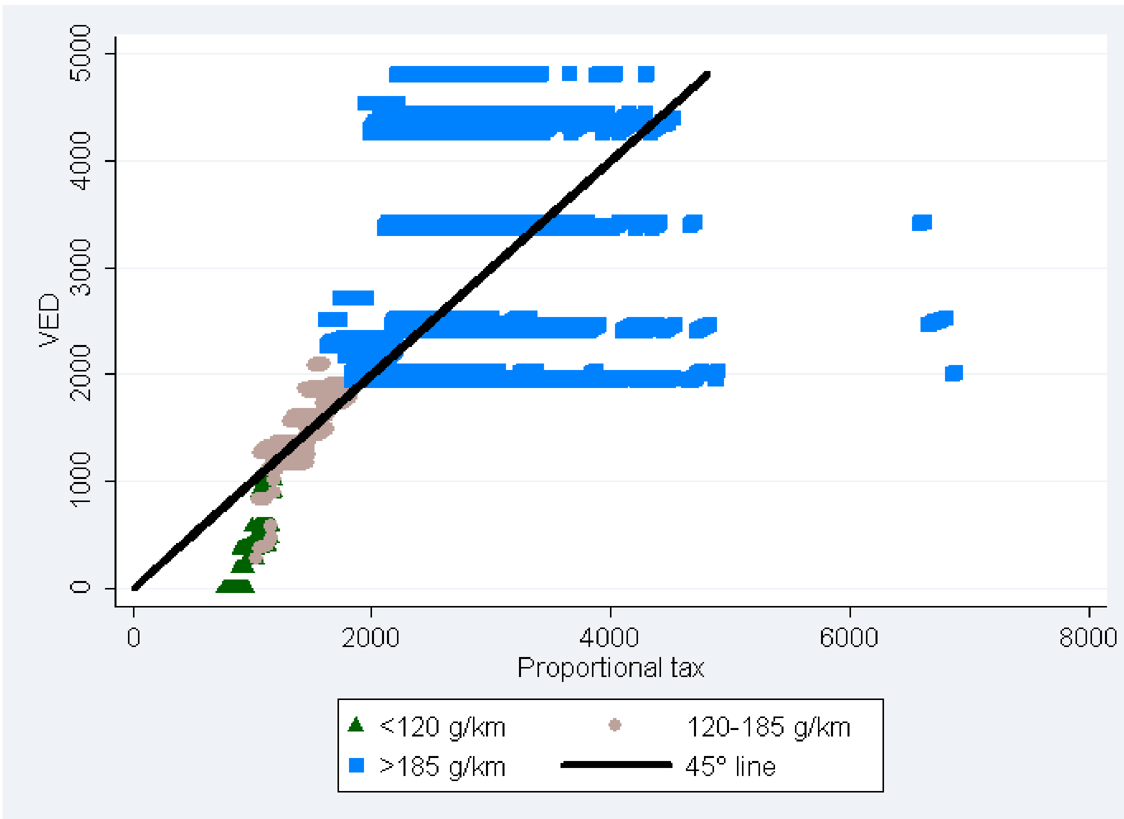

Figure 1.3: Real total tax payment under VED and proportional tax, 2005-2010, expressed in GBP. Each point represents a make-model-trim-variant at a given policy period. Both tax schemes generate the same aggregate real revenue.

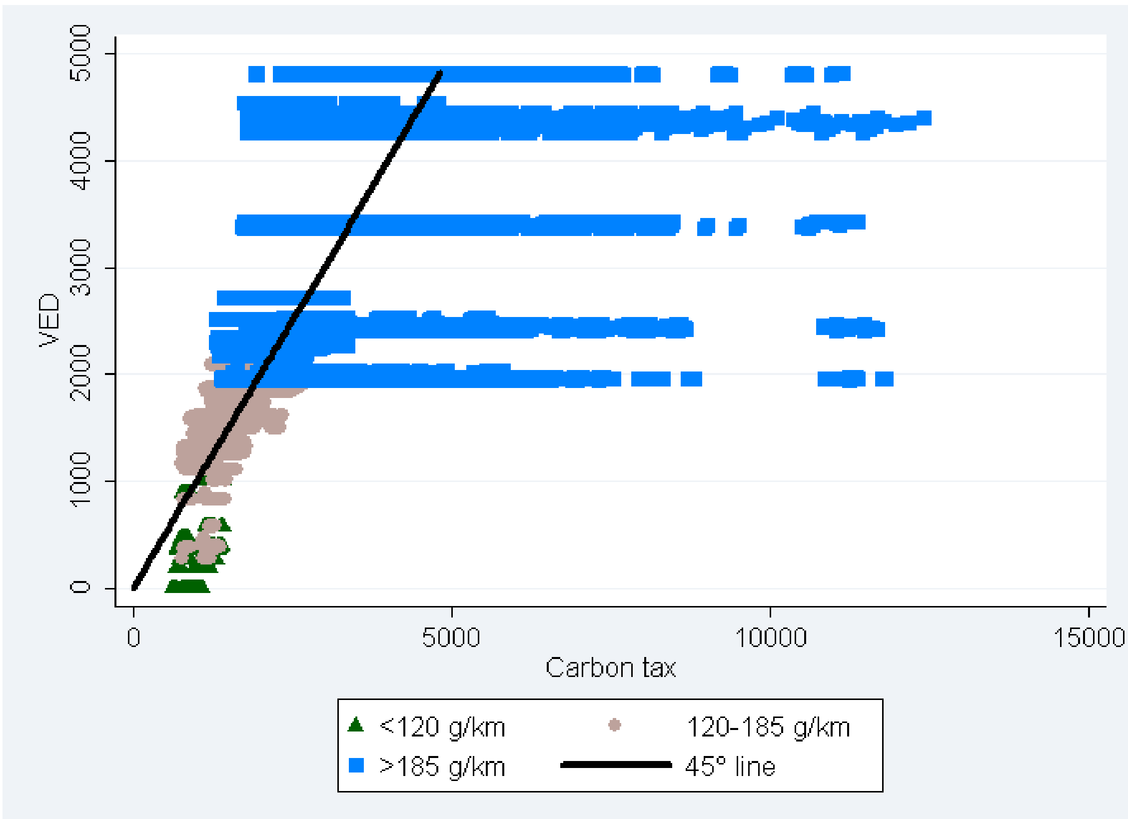

Figure 1.4: Real total tax payment under VED and carbon tax, 2005-2010, expressed in GBP. Each point represents a make-model-trim-variant at a given policy period. Both tax schemes generate the same aggregate real revenue.

certain cases more than twice, than under the VED. Figure 1.5 explains why this is the case: vehicles with a high emission rate tend to have a higher predicted mileage, mainly because they have larger engines and our mileage equation lets miles driven depend on the age of the vehicle, fuel type and engine size. An important caveat of this approach is that because we do not have data on drivers in the Polk dataset, we are assuming that mileage is completely dependent on vehicle characteristics.

To understand the effect of the proportional tax and the carbon tax on registrations and carbon emissions, we use the results of model 1.1 to predict registrations and market shares of each make-model-trim-variant in each period. Because our model implicitly assumes that changes in registrations due to the tax policy occur only within a model in a given period, we normalize predicted registrations at the make-model-period level.

We then estimate the total carbon emissions from new registrations during the lifetime of the vehicle. These emissions are a function of predicted registrations, emissions rate per km and miles driven. To predict the total mileage driven by each variant, we use the data from the UK NTS to estimate the relationship between mileage, vehicle characteristics and fuel cost, as illustrated in Appendix A.1. Our framework implies that mileage and registrations are estimated separately, and because we do not have information on the buyers, predicted miles depend only on vehicle characteristics and fuel costs.

Results on registration shares are presented in table 1.8. We calculated shares by VED band to underline the changes from changing policies. In general changes in shares across different policies are very small, but both in the case of the proportional

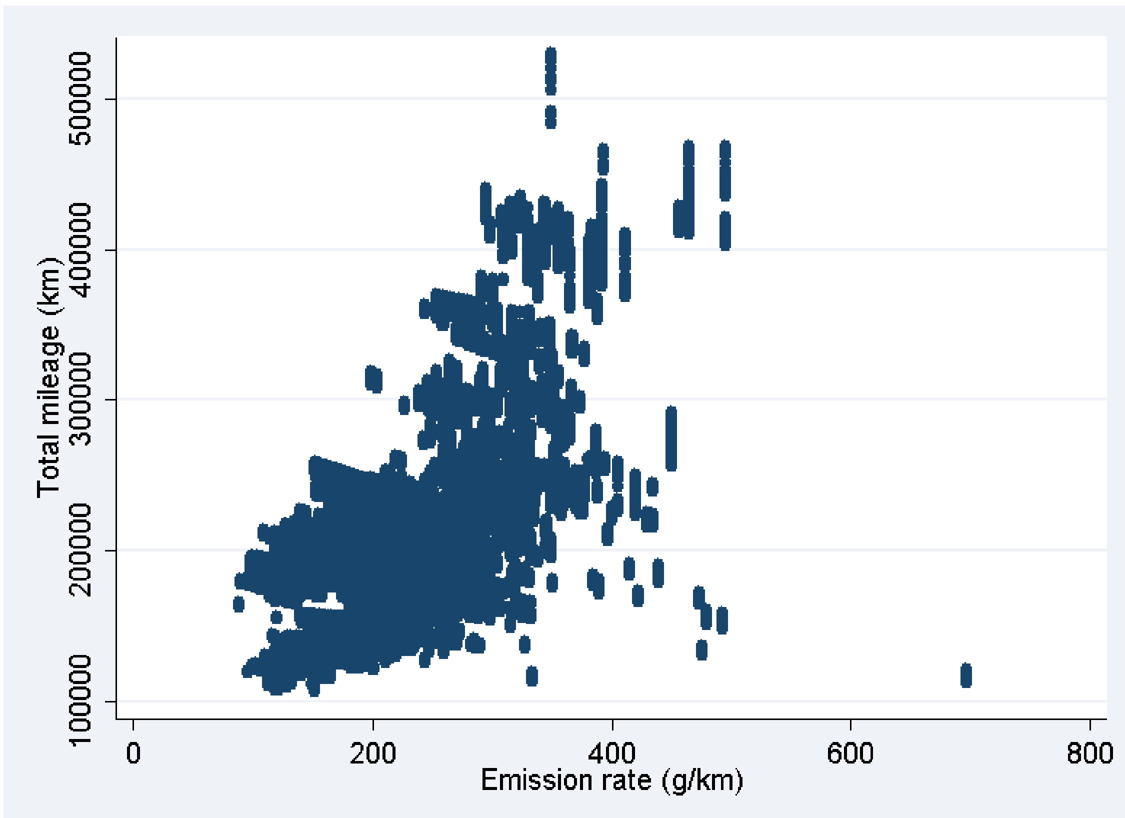

Figure 1.5: Predicted km driven during vehicle lifetime and carbon emission rate. Each point represents a make-model-trim-variant at a given policy period.

and the carbon tax, the share of clean vehicles decreased, as under the VED those variants paid little or no tax while still generating emissions. On the other hand, the shares of vehicles around the median variant increased. For what concerns high polluting vehicles, with the proportional tax we observe an increase of their market share, while under the carbon tax we see a decline. The explanation is that the VED disproportionately penalizes very highly polluting vehicles, which does not occur under a proportional scheme. Moreover, the carbon tax takes into account the number of kilometers driven. Still, the differences in market shares across different policies are quite small.

Tables 1.9 and 1.10 shows the predicted total emissions and the average emissions per vehicle across different policies. In order to isolate the effect of the registrations from the effect of mileage, we considered the effect of the carbon tax in two cases: first, when the carbon tax affects both registrations and mileage, and when it affects only registrations (which means that we assume that the miles driven are inelastic with respect to fuel prices and the carbon tax).[9]

Our results show that switching to a proportional tax would increase total emissions from new vehicles by 0.12%, while a carbon tax would decrease them by 3.71%. The effect of the carbon tax through changes in registration only is much more modest: a decline in emission of 0.36%. This is because switching from the current VED scheme to a proportional tax or a carbon tax would encourage

[9] Note that even when considering only changes in registrations there is still a change in the distribution of mileage, as different model-trim-variants have different predicted miles. What doesnt change is the mileage conditional to the variant.

Table 1.8: Predictions on new vehicle registrations

VED bands	CO_2 range (g/km)	Shares (VED)	Shares (Proportional tax)	Shares (Carbon tax)
A	88-100	0.56%	0.51%	0.49%
B	102-100	3.40%	3.29%	3.25%
C	111-120	11.41%	10.43%	10.04%
D	121-130	8.68%	8.80%	8.57%
E	131-140	12.35%	12.48%	12.44%
F	141-150	12.79%	12.74%	12.79%
G	151-165	18.15%	18.24%	18.35%
H	166-175	7.06%	7.29%	7.45%
I	176-185	6.73%	6.79%	6.94%
J	186-200	7.61%	8.03%	8.17%
K	201-225	5.79%	5.77%	5.89%
L	226-255	2.77%	2.91%	2.94%
M	256-495	2.70%	2.71%	2.68%

Predictions of the shares of new registration for each VED band, across different policies.

Table 1.9: Predictions on CO_2 emissions

VED bands	CO_2 range (g/km)	Emissions (VED)	Emissions (Proportional tax)	Emissions (Carbon tax)	Emissions (Carbon tax, same mileage)
A	88-100	930,314	840,036	784,474	803,242
B	102-110	5,141,714	4,941,456	4,755,357	4,853,414
C	111-120	22,526,911	20,539,151	19,154,287	19,724,716
D	121-130	17,764,433	17,997,275	16,929,736	17,455,475
E	131-140	25,513,831	25,684,938	24,689,392	25,395,112
F	141-150	27,962,553	27,759,839	26,822,373	27,605,694
G	151-165	44,257,369	44,283,674	42,835,905	44,228,167
H	166-175	18,042,403	18,571,813	18,277,828	18,847,969
I	176-185	18,731,762	18,836,383	18,479,627	19,114,185
J	186-200	22,977,066	24,207,045	23,613,501	24,496,113
K	201-225	19,913,020	19,790,669	19,311,664	20,090,056
L	226-255	11,894,702	12,451,300	11,956,699	12,545,763
M	256-495	16,667,384	16,711,493	15,359,296	16,244,185
Total		**252,323,463**	**252,615,072**	**242,970,139**	**251,404,090**
Change			**+0.12%**	**-3.71%**	**-0.36%**

Predictions of carbon emissions from new vehicles for each VED band, across different policies.

Table 1.10: Predictions on CO_2 emissions per vehicle

VED bands	CO_2 range (g/km)	Emissions per vehicle (VED)	Emissions per vehicle (Proportional tax)	Emissions per vehicle (Carbon tax)	Emissions per vehicle (Carbon tax, same mileage)
A	88-100	17.49	17.39	16.95	17.36
B	102-110	15.87	15.74	15.34	15.66
C	111-120	20.71	20.66	20.03	20.62
D	121-130	21.48	21.45	20.73	21.37
E	131-140	21.68	21.59	20.83	21.43
F	141-150	22.95	22.86	22.01	22.65
G	151-165	25.58	25.48	24.49	25.29
H	166-175	26.82	26.74	25.73	26.54
I	176-185	29.20	29.12	27.95	28.91
J	186-200	31.70	31.64	30.32	31.45
K	201-225	36.08	35.97	34.41	35.80
L	226-255	45.00	44.93	42.67	44.77
M	256-495	64.87	64.64	60.12	63.58
Overall		**26.48**	**26.51**	**25.50**	**26.38**

Predictions of carbon emissions per vehicle for each VED band, across different policies.

the purchase of cleaner vehicles within each VED band, as shown by the lower emissions per vehicle for each band. The VED, however, provides a strong incentive to switch to cleaner bands: as a result, there is a decline in the share of the cleanest vehicles under the proportional tax and the carbon tax, and in the case of the proportional tax there is also an increase in market shares of very highly polluting vehicles because they would be penalized less heavily than they are under the VED scheme. In the case of the proportional tax, the two effects go in opposite directions and they virtually offset each other. Their magnitudes are however modest. The carbon tax results suggest that a meaningful decrease in emissions can be obtained only through a change in miles driven rather than through changing the composition of new sales. An important caveat to this conclusion is that by construction our model assumes that changes in market shares occur only within each make-model, while the shares across models remains unchanged. As next step, we will try using a more flexible specification of model 1.1 to check the robustness of these findings.

1.6 Conclusions

In the last years, several European countries introduced various forms of vehicle registration and sales taxes linked to carbon emissions, with the goal of reducing carbon emissions. Earlier studies have shown that these policies generally are effective in pushing consumers towards cleaner vehicles. We have explored how how these policies compare with fuel taxes, which are the other main fiscal policy affecting emissions from vehicles.

Our results suggests that the elasticity of new vehicle registrations to the VED is about -0.422. This coefficient is statistically significant and robust to a series of robustness checks, and its magnitude is close to the registration tax elasticities measured in other European countries by the previous literature. We use these estimates to predict the effect of switching from the VED to a proprtional tax based on emissions rates and to a carbon tax. While the tax rates of the two alternative policies would guarantee the same revenue, switching to a proportional tax increases total emissions by 0.12%, while swtiching to a carbon tax would reduce them by 3.71%, mainly due to the decline in mileages.

Our study has a number of limitations. We do not have registration data and information on emissions before 2001 when the VED was introduced so our identification relies on the variations in the registration tax across time and VED bands. Our predictions are about the effect on vehicles shares and emissions after an increase in the VED rates, but we are not able to make in-sample predictions on the effect of the introduction of VED. Another caveat is that we did not consider explicitly the hypothesis of a supply response to the VED manufacturers changing vehicle characteristics to fit them into a particular VED band. Our assumption is that manufacturers operate in the European market as a whole and do not change vehicles characteristics for relatively small policy changes in one country. Finally, our model assumes that changes in the shares of registrations due to fiscal policies occur only within a given make-model. A restrictive model allows to ignore various confounding factors, but it might underestimate the effect of vehicle taxes on registrations and emissions, and it might be the cause of discrepancies between actual and predicted registrations.

Chapter 2: Short and long run effects of vehicle charges: evidence from Milan, Italy

2.1 Introduction

In recent decades, many cities throughout the world introduced restrictions on vehicle circulation to address increasing traffic levels and stricter environmental regulations. Such measures often encounter fierce opposition due to the inconvenience they cause to drivers, and thus they need to prove themselves effective to justify their existence.

While the economic literature has extensively discussed the implications of vehicle restrictions and road pricing (Walters, 1961, Vickrey, 1963, Keeler and Small, 1977, Newbery, 1988, Goddard, 1997, Newbery, 2005, Schmutzler, 2011), it is only recently that economists have started to analyze empirically the impact of existing policies (Eskeland and Feyziolu, 1997, Santos et al., 2000, Santos, 2004, Davis, 2008, Carrillo et al., 2013, Gallego et al., 2013, Lin et al., 2014, Wolff, 2014). Most of the current literature analyzes driving bans, showing that in certain cases such policies have unintended consequences that can exacerbate environmental and congestion issues instead of mitigating them. Some studies compare the short and the long-run effects of driving bans, finding that those policies tend to be less effective in the long run.

In recent years, several cities around the world adopted or considered vehicle charge schemes, usually congestion charges, to contain traffic and vehicle emissions in the city center.[1] A natural question is whether these policies can contribute to pollution reduction and whether the adverse effects documented with driving bans are applicable to them.

Specifically, this paper analyzes the effect on nitrogen oxides (NOx) concentrations of a vehicle pollution charge (*Ecopass*) and a vehicle congestion charge (*Area C*), adopted in the center of Milan, Italy, respectively between 2008 and 2011 and from 2012 until now. The difference between the two policies is that *Ecopass* charges a different fee to drivers based on the pollution class of the vehicle, while under *Area C* some vehicles are banned and the others pay an identical fee regardless their pollution class.

While other studies in the literature look at whether *Ecopass* or *Area C* reduced traffic or pollution concentration (Rotaris et al., 2010, Danielis et al., 2011, Percoco, 2013, 2014, Gibson and Carnovale, 2015), to the best of my knowledge this is the first study explicitly comparing the two policies and focusing on short-run, long-run, and seasonal effects on pollution.

Comparing treated and non-treated pollution monitoring stations in a differ-

[1] The most notable examples are London, Milan, Oslo, Singapore, Stockholm. Cities that recently proposed a road pricing scheme are Beijing, Guangzhou, Helsinki, San Francisco and São Paulo.

ence in difference framework, I estimate a change in NOx concentration of -8.6% under *Ecopass* and between -5.1% and -8.0% under *Area C* in the first quarter of their introduction. For both policies, the effect decreases significantly after the first quarter, decreasing further across time during *Ecopass* while remaining more stable under *Area C*. I attribute such differences mainly to the increasing share of exempt new vehicles under *Ecopass*, while under *Area C* the share of exempt vehicles does not increase with time. The results are robust to a series of robustness checks, including an event study approach that exploits an exogenous suspension of *Area C*.

The paper is organized as follows: section 2.2 describes the characteristics of the driving restriction policies adopted in Milan and in other cities; section 2.3 describes the data and shows the pollution patterns inside and outside the city center; section 2.4 describes the main empirical specification used in the paper; section 2.5 shows the empirical results and the robustness checks; section 2.6 suggests an explanation of the results using data on vehicle stock and entrances; section 2.7 is the conclusion.

2.2 Background

In the last two decades, several cities around the world have introduced policies controlling vehicle circulation. In most cases their explicit goal was reducing either congestion, or pollution, or both.

Driving bans were the first policies to be introduced. Generally they are a partial vehicle ban based on the license plate number, and they have been implemented either permanently (Mexico City, Bogotá, Quito, São Paulo), during periods of pollution peaks (several European cities) or during other special events (Beijing during the Olympic games).

Congestion charges, which require all vehicles to pay a fee to enter the city, are the other most common policy. Their increasing popularity is due to two factors: first, they are a source of revenue for the city and second, they are able to price both pollution and congestion externalities at the same time, being potentially a better option in terms of welfare gains than policies not charging drivers, like high occupancy vehicle lanes (Bento et al., 2014). They have been implemented or are on study in several cities, among them Beijing, Guangzhou, Honk Kong, London, Milan, Oslo, São Paulo, San Francisco, Singapore, and Stockholm.

Finally, low emission zones (LEZ) forbid circulation only to high polluting vehicles. They are widely adopted in Europe, especially in Germany, as a consequence of stricter air quality regulations approved by the European Union.

Due to the wide range of possible driving restriction policies and their differences in implementation, getting a general sense on their effectiveness is difficult.

Some policies have positive effects on air quality, while others are poorly designed and produce unintended consequences. Several factors, often related to the local context, contribute to the success or the failure of a specific measure. It is then hard to find a solution fitting any situation, and even harder to predict *ex ante* the outcome of a specific policy.

In his study of the LEZ in Germany, Wolff (2014) suggests that LEZ reduced pollution concentration, even in nearby areas not directly regulated. Santos et al. (2000) and Santos (2004) simulate the impact of a congestion charge on traffic and pollution in eight British cities, finding a positive reduction of both with a stronger impact on traffic. Lin et al. (2014) compare the impact of driving bans in Bogotá, Beijing, and Tianjin, presenting a formal model of intertemporal substitution of driving. Gallego et al. (2013) examine the duration of pollution reduction, showing that the public transportation reform in Santiago de Chile and the partial vehicle ban in Mexico City caused a reduction in pollution in the short run. The Mexico City case is examined also by Eskeland and Feyziolu (1997) and Davis (2008), finding that the vehicle ban increased both overall gasoline consumption and car stock without any significant benefit to air quality. Finally, Carrillo et al. (2013) study the impact of a driving ban in Quito, finding a decrease in carbon monoxide concentration.

The main lessons from the past literature is that reduction in pollution should not be taken for granted and the magnitude of any actual decrease in pollution might change across time.

The city of interest of this study is Milan, Italy. Milan is the second largest city in Italy and capital of the Lombardy region, one of the richest and most productive

Italian regions. The high population density, the high rate of economic activity and the geographical and atmospheric conditions make air pollution in Milan a particularly pressing problem. In 2007 and in the previous years Milan exceeded the European pollution concentration limits for nitrogen oxides (NOx), ozone (O_3) and particulate matter (PM_{10}) (ARPA, 2007).

Concerned by the constant high level of pollution concentration in Milan, in January 2008 the city administration introduced a new entrance policy for the city center. The measure, called *Ecopass*, was a pollution charge. It affected only vehicles classified as highly polluting and it did not forbid entrance to the city center but rather imposed the payment of a fee to drivers.

Ecopass was introduced the 2nd of January 2008 and was enforced between 7:30am and 7:30pm during weekdays, excluding most of the month of August. The fee increased with the class of the vehicle according to European emission standards.[2] Vehicles above a certain class, motorbikes, and alternative fuel vehicles (electric, hybrid, methane, LPG) could circulate without paying.

Ecopass ended the 31st of December 2011. The 16th of January 2012 the administration introduced a new policy, called *Area C*. Unlike *Ecopass*, *Area C* is a congestion charge, meaning that all vehicles have to pay the same amount regardless

[2] Under the European emission standards, all vehicles are classified according to their emissions with an ascending order, from 'Euro 0' vehicles (no standards) to 'Euro 6' (tightest standards). New vehicles must comply with the most recent standard and more stringent standards are introduced on a regular basis. For a summary of European emission standards, see `http://transportpolicy.net/index.php?title=EU:_Light-duty:_Emissions` .

of their pollution class. In addition, some categories of vehicles cannot enter the city center. The hours and the area of application of *Area C* are the same of *Ecopass*.

Table 2.1 shows in detail the fees by vehicle class and fuel. Under *Ecopass*, diesel vehicles of the same class paid more per entrance, but newer, higher class vehicles were exempt. Under *Area C*, some classes of vehicles (mostly diesel) are not allowed to circulate.

Table 2.1: *Ecopass* and *Area C* fees

Vehicle class	**Ecopass**	**Area C**
Euro 0 gasoline	5€/day	No access
Euro 1 gasoline	2€/day	5€/day
Euro 2 gasoline	2€/day	5€/day
Euro 3 (or higher) gasoline	Free access	5€/day
Euro 0 diesel	10€/day	No access
Euro 1 diesel	5€/day	No access
Euro 2 diesel	5€/day	No access
Euro 3 diesel	5€/day	No access
Euro 4 diesel without filter	5€/day after June 2010	5€/day
Euro 4 diesel with filter	Free access	5€/day
Euro 5 (or higher) diesel	Free access	5€/day

Ecopass and *Area C* fees by pollution class. Motorbikes, motorcycles and alternative fuel vehicles are exempt from the charge. Residents in the city center benefit from discounted annual pass under *Ecopass* and 40 free accesses per year and 2€/day fee afterward under *Area C*.

Both policies are enforced with CCTV cameras at the 43 gates around the city center, with signs alerting drivers that they are entering the charged area. The cameras record the license plate of the vehicle and match it with vehicle registration data. The owner can pay in advance or within 24 hours of the entrance. Various methods of payment are available (tickets from stores, meters, free toll number, ATMs, internet, direct debit) including tickets for multiple entrances and (for *Area C*) an electronic toll-collection system already used for highways. The fee allows any number of entrances for that day. Standard fines for non-paying vehicles are 70€ for *Ecopass* and 82€ for *Area C*.

The city administration installed the CCTV and the signs informing drivers about the charge on the 15th of October 2007, before the official start of the policy. They performed a test of the system and gathered data on vehicle transit, but no fee was imposed.

The congestion charge is still ongoing, with small modifications. Due to an Italian court decision, *Area C* was unexpectedly suspended from the 25th of July to the 17th of September 2012. When reintroduced, *Area C* was slightly modified and on Thursday it ended at 6pm instead of 7:30pm.

Ecopass and *Area C* are not the only policies adopted by local authorities with the goal of reducing pollution concentration from vehicles. In the period and hours of application of *Ecopass* and *Area C*, the following regulations were in place: A regional circulation ban (area A1) of Euro 0 (gasoline and diesel) and Euro 1 (diesel), on weekdays from 7:30am to 7:30pm, from 15th of October to 15th of April of each year, starting in 2008.[3] A regional circulation ban (area A1) of Euro 2 (diesel), on weekdays from 7:30am to 7:30pm, from 15th of October to 15th of April of each year, starting in 2009. A circulation ban within the Province of Milan of Euro 3 (diesel) without particulate filter, from 8:30am to 6:00pm after 10 consecutive days

[3] Area A1 is a regional classification including certain areas generally in the provinces of Milan, Bergamo and Brescia. See the section Data for a more detailed explanation.

of PM10 average concentration levels exceeding 50 μg/m3 until PM10 concentration drops under that limit, starting in December 2011.

The provincial regulation on Euro 3 diesel vehicles without filter started almost at the same time as *Area C*, which is more stringent because it prohibits access to all Euro 3 diesel.

A few earlier studies examined the effectiveness of *Ecopass* in terms of pollution reduction and overall benefits. Rotaris et al. (2010) report a first assessment by the Milan Agency for Transportation, Environment, and Territory (AMMA) suggesting that in the first eleven months (January 2008 - November 2008) PM_{10} emissions at the source inside *Ecopass* area decreased by 23%, NOx by 17% and CO^2 by 14% (AMMA, 2008). These estimates rely on the comparison between the number and the type of vehicles entering in the city center after *Ecopass* was introduced and the vehicles which entered in October 2007. A follow-up study by Danielis et al. (2011) looked at the number of days exceeding PM_{10} limit and the average annual PM_{10} concentration. Both trends show a decrease across time. However, in both studies there is no measure of pollution concentration under the counterfactual. Percoco (2013) uses an event study design around the beginning of *Ecopass*, under the implicit assumption that no anticipation effect occurred. He finds a reduction in average daily concentration for various pollutants.

The majority of the studies focus on the suspension of the policy due to a court order between July and September 2012. The suspension was completely unexpected and, as such, it is interpreted as a natural experiment. Percoco (2013, 2014) and Gibson and Carnovale (2015) use an event study approach to estimate the effect of *Area C* on pollution and traffic. Due to differences in methodology, the results are not comparable, but the main finding is that the suspension of the policy increased vehicle circulation and pollution levels. There is solid evidence that entrances of motorbikes - exempt from the charge - increased when the policy was in place. Gibson and Carnovale (2015) find also that before the suspension traffic was more intense just outside the boundaries of *Area C*, suggesting a spatial substitution effect.

While previous studies examined some aspects of the vehicle charges introduced in Milan, they failed to use control samples or the analysis was limited to a very short period. Previous research on driving bans showed that long run effects might differ from short run effects (Gallego et al., 2013), and whether this conclusion is applicable also to vehicle charges is an open question.

In this paper I ask three research questions: first, whether there are any difference between the short and the long run effect of each vehicle charge introduced in Milan; second, how the vehicle pollution charge compares to the vehicle congestion charge in terms of NOx reduction; and third, if the effect of the two policies change across different seasons.

To answer these questions, the approach I am using is to find a good counterfactual group for the pollution concentration in Milan city center.

2.3 Pollution in Milan and other data

The focus of the paper is on nitrogen oxides (NOx). In Milan NOx regularly exceeded health limits in 2007 and in the previous years (ARPA, 2007, p. 38). NOx concentration has been measured for a longer period and in more areas than other pollutants, such as particulate matter and ozone.

Data on hourly NOx concentration from 2003 to 2013 come from the Regional Environmental Protection Agency of Lombardy (*Agenzia Regionale Protezione Ambiente* - ARPA), which maintains a network of pollution and meteorological stations in the whole region. Each pollution monitoring station is classified according to three indicators: location type (urban, suburban, rural); main source of pollution (traffic, industrial, background); and zoning (A1, A2, B, C1, C2).[4]

The first two standards follow the classification of the European Environmental Agency. The zoning standard is a regional classification which takes into account pollution concentration and meteorological and geological conditions, with A1 being the most urbanized and polluted zone, generally within the three major metropolitan areas of the region (Milan, Bergamo, Brescia); zoning is important because the additional regional environmental policies, like circulation bans of certain types of vehicles, are the same within each zone type.

There are two pollution monitoring station inside the city center, called *Verziere* and *Senato*. Both stations are classified as urban, traffic type and zone A1. I then

[4]The zoning classification changed between 2003 and 2013, but the monitors used in this paper were not affected.

select 15 stations outside the treatment area with the same classification and located in the most urbanized provinces of the region (Milan, Monza e Brianza, Bergamo, Brescia). Outside stations with different classification, located in different provinces, or with more than 50% of data missing for any year between 2003 and 2013, were dropped.

Figure 2.1 shows the boundaries of *Ecopass/Area C* and the location of the other 15 stations. Outside Milan there are two main groups of stations: those in the provinces of Milan and Monza e Brianza (MiMb), relatively close to Milan city center, and those in the provinces of Bergamo and Brescia (BgBs), relatively farther away.

Hourly data on precipitation and temperature come from ARPA as well. I use information on the position of the stations to match each polluting monitoring station to data from weather monitoring stations.[5]

It is common for weather data from monitors to have some missing observations. To address this problem, I use a slight modification of the algorithm of Auffhammer and Kellogg (2011). Given a weather station (primary station) matched with a polluting monitoring station, I take the 9 closest weather stations within a 50km distance and I regress the nonmissing hourly values of the primary station with the values of the other 9 stations, including year, hour, month and day of the week fixed effects, using the predicted values to complete the missing observations of the primary station. In case of precipitation, values are censored at zero. If there

[5]The maximum distance between a matched pollution and weather station is 15.28 km, while the median distance is 3.13 km.

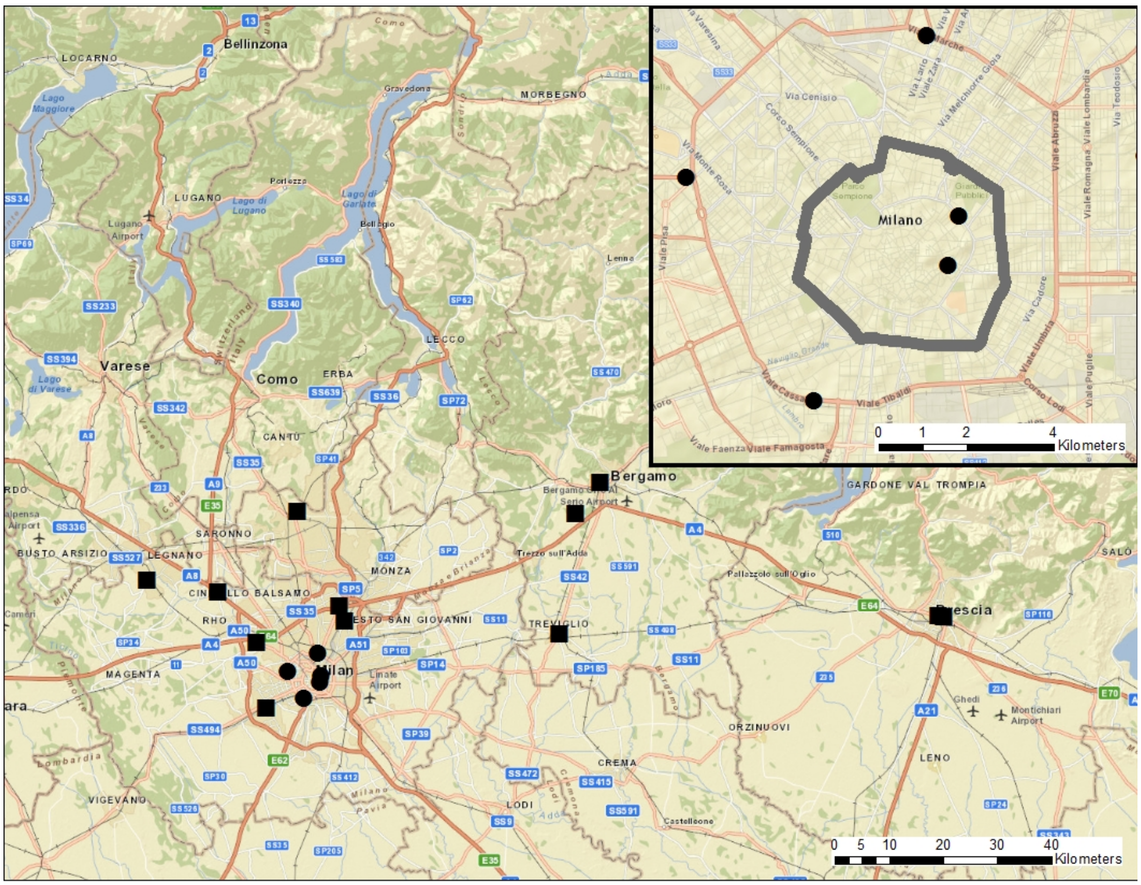

Figure 2.1: Map of pollution monitoring stations included in the data, Lombardy and Milan city (top right). Dots represent stations in Milan city, squares stations outside Milan. *Ecopass* / *Area C* boundaries in gray.

are missing values left (both primary station and any other station are missing the value), I repeat the process using the 8 closest stations, the 7 closest stations and so on.

If there are still missing values left (both primary station and its closest station are missing the value), I run the same regression dropping the closest station and using the other 8 closest stations, then the 7 closest stations and so on. In this way, a value is still missing if is missing from both the first and the second closest station (less than the 0.1% of the sample).

To test the algorithm, I select only the nonmissing observation of the primary station and the correspondent observations of the other stations. I then randomly drop about the 20% of the nonmissing values of the primary monitor and I compute the fitted values. The correlation between fitted and real values exceeds 95%. Among all the runs of this testing procedure, the distribution of the difference between real and fitted values is always centered at zero and generally has a very low variance.

Unless stated otherwise in the paper, the data are NOx hourly concentration in the period in which the policy was in force, i.e. Monday-Friday, from 8am to 7pm (8am-6pm on Thursdays from September 2012). I drop the month of August, because *Ecopass* and *Area C* were suspended for most of it, as well as official holidays and days in which the policy was suspended for any reason, and the period between the 23rd of December to the 6th of January of each year.[6] The testing period of

[6] The period close to Christmas and New Years' Eve is problematic for various reasons: first, *Ecopass* was often suspended during this period to allow shopping in the city center; second, the

the *Ecopass* system (15 October 2007 - 1 January 2008) is dropped from the sample as well.[7] Finally, I dropped all the observations from monitors outside the city center which have no correspondent non-missing value in the monitors inside the city center.

Table 2.2 and 2.3 shows summary statistics on the NOx concentration for five groups of monitoring stations: stations in the Provinces of Bergamo and Brescia; stations in the Provinces of Milan and Monza e Brianza but outside Milan; stations inside Milan but outside the city center; and the two stations in the city center (*Senato* and *Verziere*).

In the two years before the start of the *Ecopass* testing period (15 October 2007) the mean and median values of NOx in the city center are very close to Bergamo and Brescia, and lower than other surrounding areas. Also the city center has far lower extreme levels of NOx.

Looking at the first two years of the application of *Ecopass*, there is a general reduction in the mean, median, and maximum NOx concentration. In the following years, pollution concentration generally decreases at a slower pace or remains stable. Looking at these data alone, it is difficult to distinguish the effect of the pollution or the congestion charge from other factors, such as compositional changes in the vehicle fleet or economic conditions.

6th of January is an official holiday and several people take a long vacation during part of this period; third, fireworks during Years' Eve are known to increase the atmospheric concentration of various pollutants, among which NOx.

[7] The main concern is that the signs and the cameras at the boundaries might have triggered a behavioral response from drivers, even if fees or fines were not due.

Table 2.2: Summary statistics for NOx concentration

Group	n	mean	median	max
Before Treatment (Oct. 2005 - Oct. 2007)				
Senato (city center)	5,090	93.053	71	553
Verziere (city center)	5,090	86.005	60	584
Pollution charge - Ecopass (2008-2009)				
Senato (city center)	5,137	87.714	67	556
Verziere (city center)	5,137	67.274	49	426
Pollution charge - Ecopass (2010-2011)				
Senato (city center)	4,915	101.866	71	832
Verziere (city center)	4,915	69.093	50	618
Congestion charge - Area C (2012-2013)				
Senato (city center)	5,178	72.727	53	545
Verziere (city center)	5,178	64.577	47	556

Summary statistics for NOx concentration, monitors inside Milan city center. Values expressed in part per billion.

Figure 2.2 shows the plot of average daily NOx concentration for the *Verziere* station in the city center. For both of them is evident the strong seasonality of NOx concentration, sharply increasing during winter months. At the beginning of *Ecopass* and *Area C* (first and second dashed lines) there is a decrease in concentration with respect to previous years.

Figure 2.3 shows the plot of quarterly NOx concentration from 2005 to 2013 for stations in the Province of Bergamo and Brescia (BgBs) and in one of the treatment stations within the Milan city center (*Verziere*). In the pre-treatment years, NOx concentration in *Verziere* follows closely the values of Bergamo and Brescia, and then decreases with the beginning of the policy.

Table 2.3: Summary statistics for NOx concentration

Group	n	mean	median	max
Before Treatment (Oct. 2005 - Oct. 2007)				
Bergamo-Brescia, Provinces	23,910	83.174	59	865
Milan-Monza, Provinces	34,159	98.808	69	920
Milan, outside city center	14,706	107.570	76	833
Pollution charge - Ecopass (2008-2009)				
Bergamo-Brescia, Provinces	26,827	74.381	54	733
Milan-Monza, Provinces	33,260	78.714	55	707
Milan, outside city center	14,886	93.953	74	608
Pollution charge - Ecopass (2010-2011)				
Bergamo-Brescia, Provinces	26,445	77.988	54	738
Milan-Monza, Provinces	32,515	84.952	57	879
Milan, outside city center	14,201	93.686	68	738
Congestion charge - Area C (2012-2013)				
Bergamo-Brescia, Provinces	25,086	75.070	55	687
Milan-Monza, Provinces	33,358	79.371	52	820
Milan, outside city center	14,013	87.421	63	753

Summary statistics for NOx concentration, monitors outside Milan city center. Values expressed in part per billion.

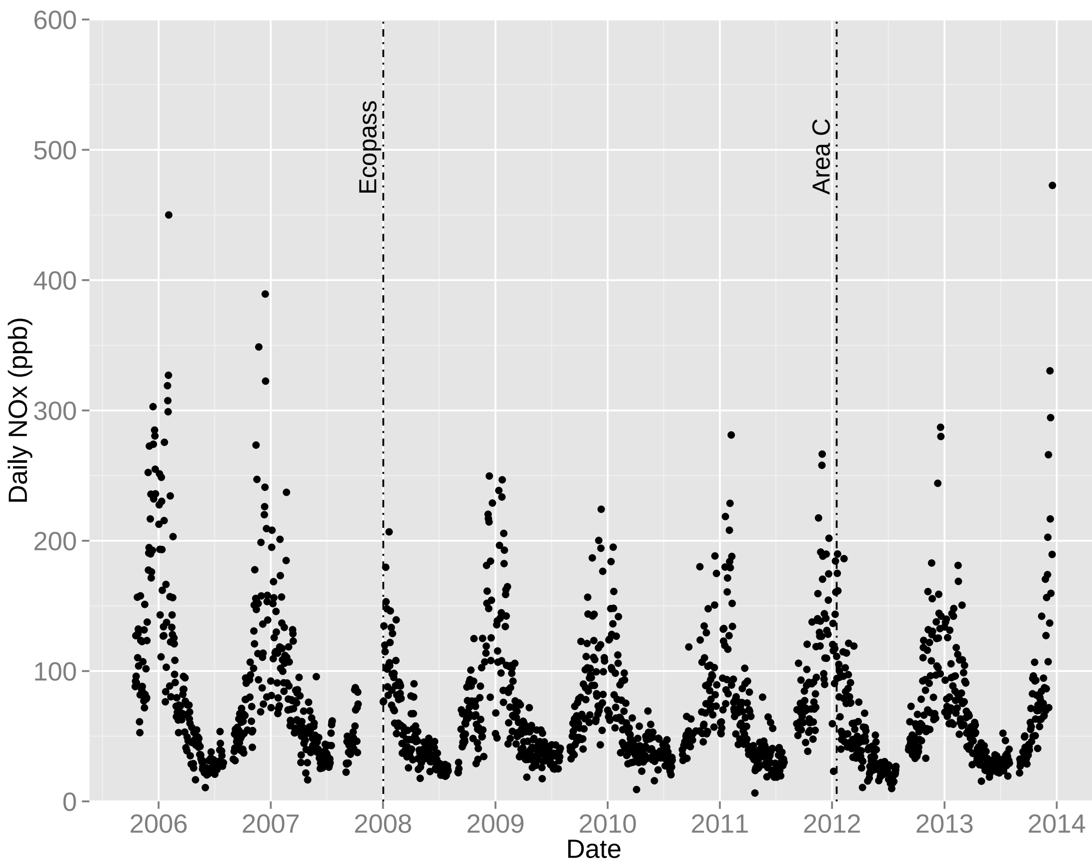

Figure 2.2: Average daily NOx concentration for *Verziere*. Dashed lines show start of *Ecopass* and *Area C*.

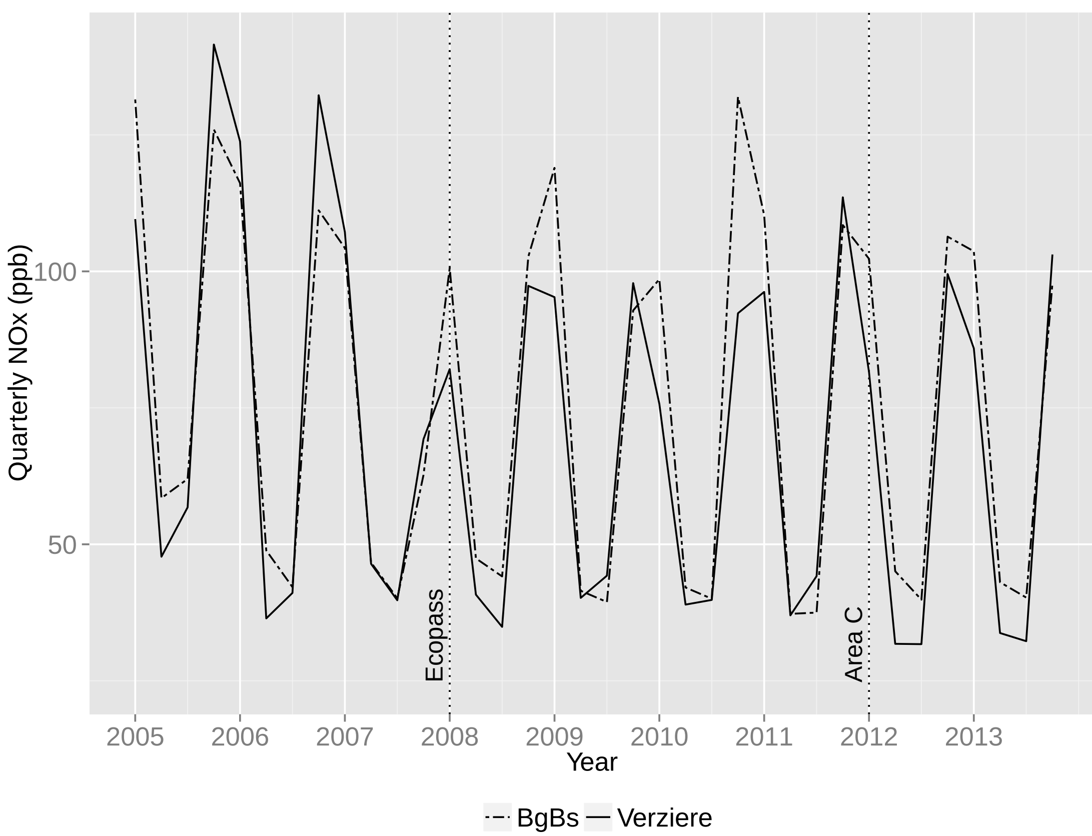

Figure 2.3: Quarterly NOx concentration by group

2.4 Model

In order to identify the effect of the pollution charge and the congestion charge on NOx concentration, I use a difference in difference approach, and estimate the regression equation

$$ln(y_{it}) = \beta_0 + \beta_1 W_{it} + \beta_2 \text{Charge}_t * \text{Center}_i + \alpha_i + \gamma_t + \epsilon_{it} \quad (2.1)$$

where $ln(y_{it})$ is the natural log of hourly NOx concentration for monitor i and time (expressed in hours) t, W_{it} is a set of weather variables, Charge_t is a dummy for the treatment period, Center_i is a dummy for the treated monitor, α_i is a set of monitor fixed effects and γ_t is a set of time fixed effects.[8]

One important aspect to consider is that hourly NOx concentration is strongly autocorrelated (Shi and Harrison, 1997, Abdel-Aziz and Frey, 2003). To account for that, I adopt a dynamic panel specification:

$$ln(y_{it}) = \beta_0 + \beta_1 W_{it} + \beta_2 \text{Charge}_t * \text{Center}_i + \beta_3 ln(y_{it-1}) + \alpha_i + \gamma_t + \epsilon_{it} \quad (2.2)$$

Dynamic panel models with fixed effects are well known to generate inconsistent estimates when the time dimension T is small, while the asymptotic bias tends

[8] Weather variables are second degree polynomials of hourly precipitation, lag of hourly precipitation, hourly temperature and lag of hourly temperature. The set of time dummies is composed by dummies for hour of the day, day of the week, month and year. The set of year controls in the treatment period is completely collinear with the treatment period dummy Charge_t.

to zero when T is large (Nickell, 1981). In my main sample, the number of time periods per monitor is above 10,000, so such bias does not appear as a concern.[9]

One of the fundamental identification assumptions for the difference in difference model is the common trend assumption. I test this assumption using the following model on the sample in the pretreatment period:

$$ln(y_{it}) = \beta_0 + \beta_1 W_{it} + \beta_2 ln(y_{it-1}) + \delta_1 T_t + \delta_2 (T_t * \text{Treatment}_i) + \alpha_i + \gamma_t + \epsilon_{it} \quad (2.3)$$

Where T_t is a day or week trend and Treatment_i is a set of dummies for the two monitors inside the city center (*Verziere* and *Senato*). The common trend hypothesis assumes that $\delta_2 = 0$. I run the test on three potential control groups: (i) monitors outside Milan in the provinces of Milan and Monza e Brianza; (ii) monitors in the provinces of Bergamo and Brescia; and (iii) monitors in both groups.

Table 2.4 and table 2.5 show the results of the test for trend assumption at the daily and weekly level respectively. I do not reject the null hypothesis of common trend for *Verziere* and the stations in the Provinces of Bergamo and Brescia, but I do reject the null hypothesis of common trend across *Verziere* (treatment monitor) and the stations in the Provinces of Milan and Monza e Brianza, outside the city of Milan. The hypothesis of common trend is rejected for *Senato* and any control

[9] Another related issue is whether the bias tend to zero when T is large but the number of units N is small, as in my sample. Kiviet (1995) derives a formulation of the bias in small samples, showing the presence of a $O(N^{-1}T^{-1})$ contribution. Monte Carlo simulations by Bun and Kiviet (2003) shows that the $O(T^{-1})$ contribution accounts for the vast majority of the bias, even when N is very small, suggesting that in my situation the bias should be very close to zero.

group. These conclusions are consistent with the evidence from graphical analysis and summary statistics. *Senato* appears to have a different trend than any potential control group.

Table 2.4: Common trend hypothesis test, daily level

	Day trend		
	(1)	(2)	(3)
	MiMb + BgBs	MiMb	BgBs
Common trend	-0.166	-0.062	-0.156
	(0.469)	(0.472)	(0.476)
Trend difference, Verziere	0.034^{**}	0.051^{***}	0.008
	(0.015)	(0.016)	(0.019)
Trend difference, Senato	0.085^{***}	0.102^{***}	0.061^{***}
	(0.022)	(0.024)	(0.023)
Observations	67920	44104	33922

Trend test results using model 2.3. Coefficients show the common daily trend between different control groups and the differences in trend of the two stations inside Milan city center for years 2005-2007. Different columns show results for different control groups (Stations outside Milan in Milan and Monza provinces, stations in Bergamo and Brescia provinces, the two groups together). All specifications include a 1-hour lag for log of NOx concentration, weather controls and station, year, month, day of the week and hour fixed effects. Standard errors in parenthesis, clustered by week-year. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

These results suggest that the difference in difference analysis should use *Verziere* as the only treatment station and Bergamo and Brescia as control group. One possible explanation is that the stations in the Milan and Monza provinces, but outside of Milan, are generally farther from city centers and closer to highways than the treatment stations, and therefore might experience different traffic composition

Table 2.5: Common week trend hypothesis test, weekly level

	Week trend		
	(1)	(2)	(3)
	MiMb + BgBs	MiMb	BgBs
Common trend	-1.163	-0.431	-1.091
	(3.280)	(3.305)	(3.334)
Trend difference, Verziere	0.239^{**}	0.361^{***}	0.055
	(0.105)	(0.111)	(0.133)
Trend difference, Senato	0.599^{***}	0.718^{***}	0.428^{***}
	(0.156)	(0.167)	(0.164)
Observations	67920	44104	33922

Trend test results using model 2.3. Coefficients show the common weekly trend between different control groups and the differences in trend of the two stations inside Milan city center for years 2005-2007. Different columns show results for different control groups (Stations outside Milan in Milan and Monza provinces, stations in Bergamo and Brescia provinces, the two groups together). All specifications include a 1-hour lag for log of NOx concentration, weather controls and station, year, month, day of the week and hour fixed effects. Standard errors in parenthesis, clustered by week-year. *p<0.1; **p<0.05; ***p<0.01

and trends. On the other hand, stations in the provinces of Bergamo and Brescia are located in more central areas. Among the two stations in the Milan city center, *Verziere* is the most central and thus NOx concentration is less likely to be affected by traffic right outside the boundaries of *Ecopass* and *Area C*; *Senato* is closer to the boundaries of the city center and in the proximity of one of the major traffic rings, which might explain why its pre-treatment trend is different than *Verziere* and the potential control group.

2.5 Results

Table 2.6 shows the results for model 2.2 of the effect of the pollution charge (*Ecopass*) on NOx concentration during the first two years of implementation (2008-2009). The effect is disaggregated by quarter-year. Because, for testing purposes, the city administration installed cameras monitoring vehicles entrances and signs showing the boundaries of *Ecopass* area from the 15th October 2007, before the actual beginning of *Ecopass*, I excluded the period from the 16th October 2007 to the 1st Janaury 2008 from my sample. The pre-treatment period is therefore from the 15th October 2005 to the 15th October 2007.

After controlling for seasonality, the preferred specification with weather controls and monitor fixed effects (column 4) shows evidence of a short term effect which vanishes in the long run. In particular, the largest reduction of NOx due to *Ecopass* occurred in the first quarter after its introduction (-8.6%). From the second quarter of 2009, there is no statistically significant decrease in NOx compared to pre-treatment levels. Within the year, the first quarter seems also to be the period in which the reduction is stronger, but it is not clear whether it is because the effectiveness of the policy changes across seasons or because the effect fades away in the long-run.

The effect of the congestion charge (*Area C*) is more difficult to address because the policy started almost immediately after the end of *Ecopass*. The two possible pre-treatment periods are the two years before *Area C* (2010-2011), when *Ecopass* was enforced, or the two years before *Ecopass*.

Table 2.6: Ecopass effect on NOx

	(1)	(2)	(3)	(4)
2008 Q1	0.007	-0.046^{***}	-0.105^{***}	-0.086^{***}
	(0.010)	(0.009)	(0.009)	(0.008)
2008 Q2	-0.039^{***}	-0.016^{**}	-0.021^{**}	-0.025^{***}
	(0.009)	(0.007)	(0.010)	(0.009)
2008 Q3	-0.048^{***}	-0.009	-0.052^{***}	-0.053^{***}
	(0.011)	(0.009)	(0.015)	(0.014)
2008 Q4	0.045^{***}	0.025^{**}	-0.046^{***}	-0.045^{***}
	(0.010)	(0.011)	(0.012)	(0.013)
2009 Q1	0.019	-0.048^{***}	-0.044^{***}	-0.039^{***}
	(0.013)	(0.010)	(0.013)	(0.012)
2009 Q2	-0.042^{***}	-0.002	0.006	0.010
	(0.009)	(0.009)	(0.010)	(0.010)
2009 Q3	-0.033^{***}	0.024^{***}	0.019^{*}	0.021^{**}
	(0.009)	(0.006)	(0.010)	(0.010)
2009 Q4	0.042^{***}	0.009	-0.010	-0.017
	(0.009)	(0.015)	(0.012)	(0.012)
Observations	60723	60721	60723	60721
Weather controls		X		X
Fixed effects			X	X

Results for model 2.2. Coefficients measure the relative change of hourly NOx concentration due to Ecopass in 2008-2009, disaggregated by quarter-year. Treatment station is *Verziere*, control stations are those in the Provinces of Bergamo and Brescia. Pre-treatment period is from 15th October 2005 to 15th October 2007. Data contains hourly NOx concentration in part per billion during Mon-Fri, from 8am to 7pm, excluding the month of August, holidays, days in which Ecopass was suspended, and observations with a correspondent missing value for the treatment station. Fixed effects are for station, year, month, day of the week, hour. Standard errors in parenthesis, clustered by week-year. *p<0.1; **p<0.05; ***p<0.01

Using as pre-treatment period the two years before *Area C* implies that model 2.2 will estimate the real effect only if there was no effect of *Ecopass* in those years. Otherwise, it will estimate the effect of *Area C* minus the long-term effect of *Ecopass*, and arguably it will be lower than the true effect.

Using as pre-treatment period the two years before *Ecopass* (from the 15th October 2005 to the 15th October 2007) implies that the model will identify the true effect of *Area C* only if the (conditional) difference in NOx between control and treatment monitors in the pre-treatment period is equal to their difference in the treatment period in the counterfactual world with no *Area C*.

Formally, consider the following basic difference in difference model, where y_{it} is the outcome at period $t = 0, 1$ for group $i = C, T$, X_{it} is a vector of control variables and $Y_{it} \equiv E(y_{it}|X_{it})$. If the treatment occurs at time 1 for group T, then $Y_{T1} = \alpha + \beta + \delta + \gamma_1$, $Y_{T0} = \alpha + \gamma_0$, $Y_{C1} = \alpha + \delta$ and $Y_{C0} = \alpha$.

Under the assumption that $\gamma_0 = \gamma_1 = \gamma$, the estimation of the treatment effect β is given by $(Y_{T1} - Y_{C1}) - (Y_{T0} - Y_{C0}) = (\beta + \gamma) - (\gamma) = \beta$. In my framework, $\beta \leq 0$ would be the effect of the *Area C* scheme.[10]

When using 2010-2011 as pre-treatment period $\gamma_0 = \gamma_1 + \eta$, that is the 'true' gap γ_1 plus the long term effect of *Ecopass* η. The model estimates $(Y_{T1} - Y_{C1}) - (Y_{T0} - Y_{C0}) = \beta + \gamma_1 - \gamma_1 - \eta = \beta - \eta$. Reasonably $\eta \leq 0$, so $\beta - \eta \geq \beta$ and the estimate is a upper bound of the true effect.

When using the years before *Ecopass* (Oct 2005- Oct 2007) as pre-treatment

[10]It is simple to generalize for a quarterly disaggregated effect by assuming $\gamma_{0q} = \gamma_{1q} = \gamma_q$, where q is the quarter.

period for *Area C*, there are more than four gap years between the pre-treatment and the treatment period. The common trend assumption might not hold for such long period, and it is possible that $\gamma_0 \neq \gamma_1$, for instance $\gamma_0 + \xi = \gamma_1$. Therefore, the model would estimate $(Y_{T1} - Y_{C1}) - (Y_{T0} - Y_{C0}) = \beta + \gamma_0 + \xi - \gamma_0 = \beta + \xi$.

Only if $\xi = 0$ (common trend assumption holds) the model estimates correctly β. Otherwise the estimate is biased, and the direction of the bias depends on whether ξ is positive or negative, which cannot be determined a priori.

Table 2.7 show the results for model 2.2 on the effect of *Area C* on NOx concentration using 2010-2011 as the pre-treatment period. This can be interpreted as the gap between the effect of *Area C* and the long run effect of *Ecopass*. Results show again that the highest reduction occurred in the first quarter of 2012. Statistically significant reduction in NOx occurred only in the first and third quarter of each year. That means that, in those quarters, *Area C* improved air quality with respect to *Ecopass* in its last years.

Table2.8 shows the results for model 2.2 on the effect of *Area C* on NOx concentration using the two years before *Ecopass* as pre-treatment period. The coefficients are all negative and significant. Differently from *Ecopass*, the effect of *Area C* on NOx lasts longer, even if the coefficients are slightly higher across time (after the first quarter, such difference is not statistically significant).

It is reasonable to assume that coefficients using data from 2010-2013 shown in table 2.7 are an upper bound of the true effect. Results of table 2.8 using two years before *Ecopass* as pre-treatment, are more difficult to interpret because it is not possible to guess a priori the direction of the bias, if any. However, considering that

Table 2.7: Area C effect on NOx, pre-treatment during Ecopass

	Pre-treatment 2010-2011			
	(1)	(2)	(3)	(4)
2012 Q1	-0.002	-0.054^{***}	-0.060^{***}	-0.051^{***}
	(0.013)	(0.014)	(0.013)	(0.012)
2012 Q2	-0.064^{***}	-0.041^{***}	-0.010	-0.018
	(0.010)	(0.008)	(0.014)	(0.014)
2012 Q3	-0.065^{***}	-0.024	-0.030	-0.038^{*}
	(0.015)	(0.018)	(0.020)	(0.020)
2012 Q4	0.030^{***}	0.023^{**}	-0.013	-0.009
	(0.011)	(0.010)	(0.011)	(0.010)
2013 Q1	0.017	-0.052^{***}	-0.023^{*}	-0.024^{**}
	(0.011)	(0.011)	(0.012)	(0.010)
2013 Q2	-0.065^{***}	-0.048^{***}	-0.0003	-0.009
	(0.010)	(0.009)	(0.012)	(0.011)
2013 Q3	-0.064^{***}	-0.018^{**}	-0.021	-0.025^{*}
	(0.010)	(0.008)	(0.013)	(0.013)
2013 Q4	0.031^{***}	0.029^{***}	-0.001	0.002
	(0.010)	(0.009)	(0.010)	(0.009)
Observations	61302	61302	61302	61302
Weather		X		X
Fixed effects			X	X

Results for model 2.2. Coefficients measure the relative change of hourly NOx concentration due to Area C in 2012-2013 compared to Ecopass long-term effect (columns 1-4) or no treatment (columns 5-8), disaggregated by quarter-year. Treatment station is *Verziere*, control stations are those in the Provinces of Bergamo and Brescia. Pre-treatment period is from 2010 to 2011 for columns 1-4, and from 15th October 2005 to 15th October 2007 for columns 5-8. Fixed effects are for station, year, month, day of the week, hour. Standard errors in parenthesis, clustered by week-year. *p<0.1; **p<0.05; ***p<0.01

Table 2.8: Area C effect on NOx, pre-treatment before Ecopass

	Pre-treatment 2005-2007			
	(1)	(2)	(3)	(4)
2012 Q1	-0.006	-0.065^{***}	-0.091^{***}	-0.080^{***}
	(0.014)	(0.014)	(0.012)	(0.012)
2012 Q2	-0.067^{***}	-0.044^{***}	-0.042^{***}	-0.043^{***}
	(0.010)	(0.008)	(0.013)	(0.013)
2012 Q3	-0.069^{***}	-0.025	-0.056^{***}	-0.057^{***}
	(0.015)	(0.016)	(0.021)	(0.021)
2012 Q4	0.027^{**}	0.010	-0.037^{***}	-0.038^{***}
	(0.011)	(0.010)	(0.011)	(0.011)
2013 Q1	0.014	-0.064^{***}	-0.054^{***}	-0.053^{***}
	(0.010)	(0.011)	(0.011)	(0.010)
2013 Q2	-0.068^{***}	-0.052^{***}	-0.033^{***}	-0.036^{***}
	(0.010)	(0.008)	(0.011)	(0.010)
2013 Q3	-0.068^{***}	-0.021^{***}	-0.046^{***}	-0.045^{***}
	(0.010)	(0.008)	(0.012)	(0.012)
2013 Q4	0.027^{***}	0.015^{*}	-0.026^{***}	-0.026^{***}
	(0.010)	(0.009)	(0.010)	(0.009)
Observations	58917	58915	58917	58915
Weather		X		X
Fixed effects			X	X

Results for model 2.2. Coefficients measure the relative change of hourly NOx concentration due to Area C in 2012-2013 compared to Ecopass long-term effect (columns 1-4) or no treatment (columns 5-8), disaggregated by quarter-year. Treatment station is *Verziere*, control stations are those in the Provinces of Bergamo and Brescia. Pre-treatment period is from 2010 to 2011 for columns 1-4, and from 15th October 2005 to 15th October 2007 for columns 5-8. Fixed effects are for station, year, month, day of the week, hour. Standard errors in parenthesis, clustered by week-year. *p<0.1; **p<0.05; ***p<0.01

the results suggest a small long run effect of *Ecopass* and that differences between coefficients of tables 2.7 and 2.8 using different pre-treatment periods range between 1.9 and 2.9 percentage points, it is plausible to consider coefficients of table 2.7 as a lower bound.

The major difference between *Ecopass* and *Area C* is that *Area C* seems to have a longer-lasting effect on NOx. For both policies, the first quarter is when the effect is the highest.

Including *Senato* station in the treatment group (table 2.9) does not change the main findings. The preferred specification shows a short term effect of *Ecopass* and a more persistent decrease in NOx for *Area C*. The only differences are that the effect during the first quarter is lower than the model with only *Verziere* and the coefficients for *Area C* using the 2010-2013 period (column 2) are lower and closer to those using the period before *Ecopass* (column 3).[11] Results of column (1) and (3) are quite similar to those using only *Verziere*.

The difference in difference approach identifies the true effect of the policies if the control group is a valid counterfactual of the treatment group. Another identification strategy that does not rely on a control group is the event study approach, which is by far the most common in the analysis of road pricing policies in Milan (Percoco, 2013, 2014, Gibson and Carnovale, 2015) and in other cities

[11]This suggest either that the long term effect of *Ecopass* at *Senato* is lower than at *Verziere*, narrowing the differences between the two specifications, or that *Senato* and the control group have no common trend.

Table 2.9: Area C and Ecopass effect on NOx, all treatment monitors

	(1) Ecopass	(2) Area C (2010-2013)	(3) Area C (2005-2007, 2012-2013)
First year Q1	-0.072^{***}	-0.042^{***}	-0.052^{***}
	(0.008)	(0.012)	(0.012)
First year Q2	-0.025^{***}	-0.026^{**}	-0.037^{***}
	(0.008)	(0.012)	(0.012)
First year Q3	-0.041^{***}	-0.043^{***}	-0.052^{***}
	(0.011)	(0.015)	(0.015)
First year Q4	-0.045^{***}	-0.043^{***}	-0.055^{***}
	(0.013)	(0.010)	(0.011)
Second year Q1	-0.014	-0.042^{***}	-0.053^{***}
	(0.012)	(0.010)	(0.010)
Second year Q2	0.022^{**}	-0.018^{*}	-0.030^{***}
	(0.009)	(0.009)	(0.009)
Second year Q3	0.027^{***}	-0.036^{***}	-0.045^{***}
	(0.009)	(0.012)	(0.012)
Second year Q4	-0.010	-0.030^{***}	-0.042^{***}
	(0.011)	(0.009)	(0.009)
Observations	70889	71356	69140

Results for model 2.2. Coefficients measure the relative change of hourly NOx concentration due to Ecopass and Area C, disaggregated by quarter-year. Treatment stations are *Senato* and *Verziere*, control stations are those in the Provinces of Bergamo and Brescia. Pre-treatment period is from 15th October 2005 to 15th October 2007 for column (1) and (3) and from 2010 to 2011 for column (2). All models control for lag of log NOx, weather, and station, year, month, day of the week, hour fixed effects. Standard errors in parenthesis, clustered by week-year.

*p<0.1; **p<0.05; ***p<0.01

(Davis, 2008, Lin et al., 2014).[12]

Because such an approach identifies the local treatment effect at a certain threshold, it does not provide insights on the short and the long term effects of a policy. However, the results of the event study should be comparable to those obtained with difference in difference in the same quarter of the same year.

In the case of Milan, the previous literature focused on a suspension of *Area C* due to a court decision between the 26th of July and the 16th of September 2012. Because such a decision came unexpectedly and was enforced immediately, there are no possible anticipation effects which might bias the results.

I use the following model:

$$ln(y_{it}) = \delta_0 + \delta_1 W_{it} + \delta_2 \text{Suspension}_t + \delta_3 ln(y_{it-1}) + \eta_t + \epsilon_{it} \tag{2.4}$$

in which Suspension_t is a dummy equal to 1 for the period of the suspension, and η_t is a set of time fixed effects.[13]

The sample is limited to a window that is symmetric around the suspension (3 June 2012 - 16 September 2012). I also run placebo tests using the same period for years 2011 and 2013.

Table 2.10 shows the result of the event study: the suspension of *Area C* caused an increase in NOx by 5.1% in *Verziere* and 10.1% in *Senato*. In particular, the

[12]Some authors use the term "regression discontinuity" instead of "event study".

[13]To allow more flexibility, fixed effects include hour of the day, day of the week, and the interaction terms between month and hour, month and day of the week, and day of the week and hour.

result for *Verziere* is very similar in absolute value to the effect in the third quarter (July-August-September) of 2012, -5.7%, estimated with a difference in difference model (table 2.8, column 4).[14] The placebo regressions using other years show no significant effect.

Table 2.10: Effect of Area C suspension

	(1)	(2)	(3)
Year	2012	2011 (placebo)	2013 (placebo)
Panel A: Verziere			
Suspension	0.051^{**}	-0.001	-0.004
	(0.025)	(0.030)	(0.017)
Observations	781	707	847
Panel B: Senato			
Suspension	0.101^{**}	-0.001	0.008
	(0.041)	(0.030)	(0.022)
Observations	782	706	847

Results for model 2.4. Coefficients in column (1) measure the percentage increase in NOx due to the suspension of Area C from the 26th of July 2012 to the 16th of September 2012. Columns (2) and (3) are placebo tests for the same days of the year in the years before and after the suspension. Data contains hourly NOx concentration in part per billion for *Verziere* (panel A) or *Senato* (Panel B) stations between the 3rd of June to the 16th of September, only during Mon-Fri, from 8am to 7pm, excluding holidays and observations with a correspondent missing value for the treatment station. Fixed effects are for month, day of the week, hour, and their interactions. Standard errors in parenthesis, clustered by week. *p<0.1; **p<0.05; ***p<0.01

[14] The two coefficients have opposite sign because the first is the effect of a suspension of *Area C*, the second is the effect of its introduction.

Subtracting from this value the gap between *Area C* and *Ecopass* estimated in table 2.7, column 4, the long term effect of *Ecopass* during the third quarter is about -1.3% change in NOx.

The results of the event study show that 1) the estimates of the effect of *Area C* are at least a plausible lower and upper bound of the true value, and 2) the long term effect of *Ecopass* is low.

To further test the robustness of the results, I ran model 2.2 on two placebo treatment periods before *Ecopass* and *Area C* were introduced. I considered two subsamples of my dataset, one including the years 2004 and 2005, the other including 2005 and 2006, and I considered 2005 and 2006 as the placebo treatment period. Otherwise, the two subsamples of my dataset are similar to the one used for the main analysis (considering only weekdays, 8am-7pm, and dropping official holidays and August).

Table 2.11 shows the results of the placebo tests indicating that there is no statistically significant decrease in NOx in any quarter. In some quarters the coefficients are positive and significant, but without a clear pattern in terms of seasonality.

The sample used so far included only the hours from 8am to 7pm, when the charge was enforced. I repeated the same analysis as in tables 2.6, 2.7 and 2.8 using only the two hours before and the two hours after *Ecopass* and *Area C* (6-7am and 8-9pm).

The two hours before and after the policy are the closest to the treatment hours. Under the assumption of no intertemporal spillover effect (people driving at different time of the day due to the policy), they can be used in a falsification test.

Table 2.11: Effect of placebo treatments

	(1) 2005 treatment	(2) 2006 treatment
Quarter 1	−0.016	0.032***
	(0.011)	(0.011)
Quarter 2	0.034***	−0.001
	(0.011)	(0.012)
Quarter 3	0.043***	0.019
	(0.015)	(0.013)
Quarter 4	0.016	0.033**
	(0.014)	(0.015)
Observations	29400	29271

Results for model 2.2 with placebo treatment. Coefficients measure the relative change of hourly NOx concentration assuming a placebo treatment, disaggregated by quarter-year. The treatment period is year 2005 for column (1) and year 2006 for column (2). Treatment station is *Verziere*, control stations are those in the Provinces of Bergamo and Brescia. Data contains hourly NOx concentration in part per billion during Mon-Fri, from 8am to 7pm, excluding the month of August, official holidays, and any observation with a correspondent missing value for the treatment station. Standard errors in parenthesis, clustered by week-year. *p<0.1; **p<0.05; ***p<0.01

Allowing for intertemporal spillovers, results show the extent of those spillovers.

Table 2.12 shows the results for the preferred specification with fixed effects and weather controls. Most of the coefficients are not significant, suggesting that the main results are robust and no clear evidence of systematic spillover effects.[15]

[15]Obviously that does not rule out completely the hypothesis that an hypothetical bias and a spillover effect of opposite sign and similar size are canceling out each other. On the light of other robustness checks, this seems unlikely.

Table 2.12: Results for 6-7am and 8-9pm

	(1) Ecopass	(2) Area C (2010-2013)	(3) Area C (2005-2007, 2012-2013)
First year Quarter 1	-0.007	-0.014	0.010
	(0.021)	(0.020)	(0.020)
First year Quarter 2	-0.005	-0.080^{***}	-0.057^{**}
	(0.014)	(0.027)	(0.027)
First year Quarter 3	0.005	-0.050	-0.039
	(0.032)	(0.039)	(0.039)
First year Quarter 4	0.035	-0.007	0.031
	(0.022)	(0.023)	(0.024)
Second year Quarter 1	0.026	-0.018	0.012
	(0.020)	(0.018)	(0.018)
Second year Quarter 2	0.021	-0.019	0.008
	(0.016)	(0.018)	(0.017)
Second year Quarter 3	0.023	0.028	0.042^{*}
	(0.026)	(0.025)	(0.025)
Second year Quarter 4	0.048^{**}	0.032^{*}	0.071^{***}
	(0.022)	(0.018)	(0.019)
Observations	20453	20660	19918

Results for model 2.2. Coefficients measure the relative change of hourly NOx concentration due to Ecopass and Area C, disaggregated by quarter-year. Pre-treatment period is from 15th October 2005 to 15th October 2007 for column (1) and (3) and from 2010 to 2011 for column (2). Treatment station is *Verziere*, control stations are those in the Provinces of Bergamo and Brescia. Data contains hourly NOx concentration in part per billion during Mon-Fri, considering only 6-7am and 8-9pm, excluding the month of August, official holidays, and any observation with a correspondent missing value for the treatment station. All the models control for lag of log NOx, weather, and station, year, month, day of the week, hour fixed effects. Standard errors in parenthesis, clustered by week-year. *p<0.1; **p<0.05; ***p<0.01

2.6 Discussion

The results in the previous section suggest that *Area C* has been more effective than *Ecopass* in maintaining the NOx reduction in the long run. In this section of the paper, I explore two possible explanations for this result: different characteristics of the policies and changes in the vehicle stock.

Data on vehicle entrances in the area affected by the policies are available on a regular basis for the period after the introduction of *Ecopass* and *Area C*, while before the introduction of *Ecopass* there are only few weeks of data from October and November 2007. These data on the pre-treatment period are problematic because this range is not necessarily representative of the whole sample period and because cameras and signs announcing the future program were already installed in the city center boundaries, these may have had an effect on vehicle traffic.

Table 2.13 shows the average entrances into the *Ecopass* area. The share of paying drivers is decreasing with time while average entrances per day are lower than the period before *Ecopass* but are increasing with time.

Table 2.13: Statistics on vehicle entrances into Ecopass area

	Average daily vehicle entrances			
Time	Paying	Exempt	Total	% paying
Oct-Nov 2007	38,079	52,501	90,580	41.80%
2008	16,322	55,407	71,729	22.75%
2009	12,255	62,842	75,097	16.31%
Jan-Jun 2010	11,569	64,545	76,144	15.19%

Source: Agenzia Mobilità, Ambiente e Territorio

The main conclusion is that the total number of entrances decreased after the introduction of *Ecopass* but started increasing again after one year. The share of paying vehicles decreased over time. This conclusion is consistent with the structure of the *Ecopass* fee, where new vehicles do not pay because they belong to the cleanest class (EURO 5).

For *Area C* there is more public information available on entrances: figure 2.4 shows a comparison between average monthly entrances in *Area C* between 2012 and 2013. Consistent with this paper's findings, the gap in 2012 and 2013 is large only in the first quarter, while in other months the difference is small. One of the reasons is that with the congestion charge no vehicle - not even the cleanest - is exempt from the fee. Because, under *Area C*, the dirtiest vehicles are banned from circulation, the number of entering vehicles might increase with time when these vehicles are replaced with new, cleaner vehicles subject to the fee but allowed to enter.

A complementary explanation of the poor long run effect of *Ecopass* is given by changes in the vehicle stock: even in a world without any charge, drivers replace their vehicles. Those vehicles are different from the new ones by a series of characteristics affecting NOx emissions.

I use data on vehicle fleet composition in the provinces of Milan and Monza e Brianza between 2005 and 2012, provided by the Automobile Club d'Italia, to provide some suggestive evidence on the fleet turnover behind the results.[16]

[16]Data at the provincial level provide more detailed information on vehicle stock than at the municipal level. Additionally, there is evidence that a sizable share of entrances in the city center comes from vehicles registered outside Milan: between 2005 and 2012, the total number of circulat-

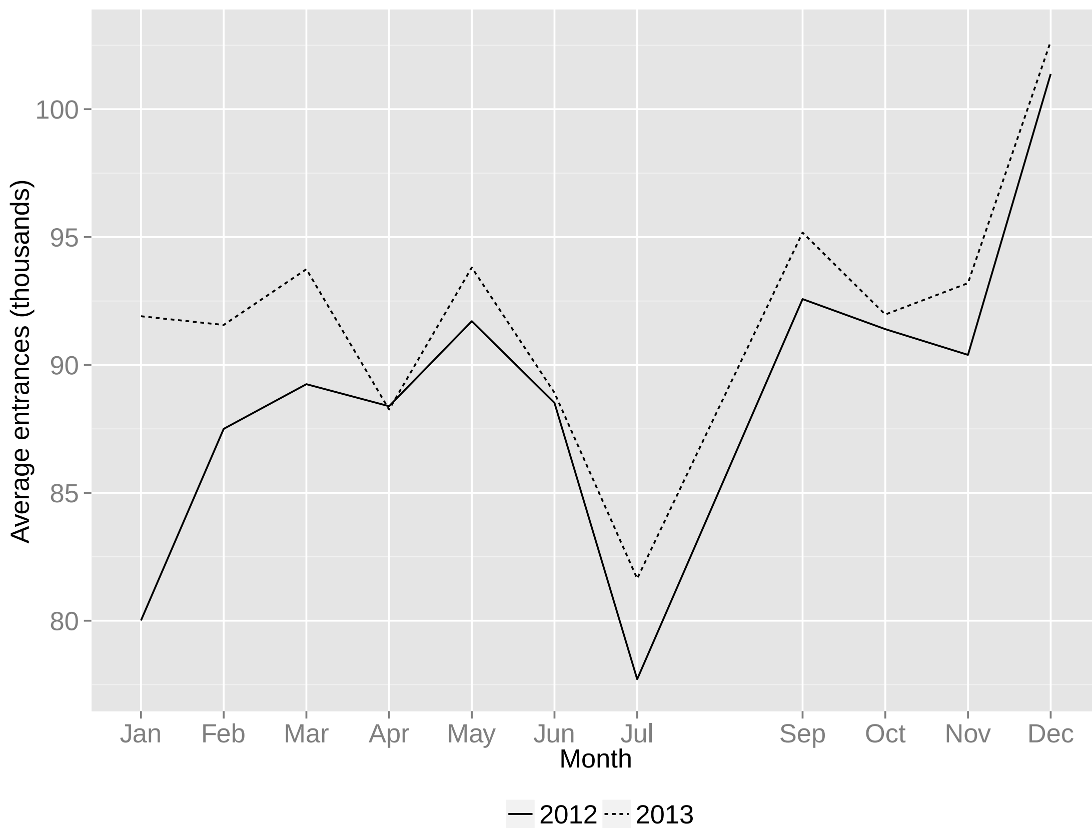

Figure 2.4: Monthly average entrances in *Area C*, 2012-2013. Source:Agenzia Mobilità, Ambiente e Territorio

Figure 2.5 show the evolution of the vehicle stock based on the amount of the charge for *Ecopass* and *Area C*. The trend looks pretty linear and shows a gradual replacement of those vehicles which had to pay the fee (*Ecopass*) or are banned from circulation (*Area C*).

Figure 2.6 shows the number of registered units in the provinces of Monza and Milan disaggregated by gasoline cars, diesel cars, and motorcycles. The share of diesel vehicles is increasing across time.

Diesel vehicles can mitigate the long run effect of *Ecopass* when drivers change their old gasoline vehicles, subject to the fee, with new, exempt diesel vehicles: people switch to diesel because its fuel price is lower than gasoline, but diesel vehicles have lower NOx standards than gasoline vehicles within the same pollution class. Furthermore, recent studies in atmospheric science show that there is a difference in vehicle emissions under controlled conditions - used to define vehicle pollution class - and emissions on road: actual NOx emissions by new, high vehicle class vehicles are closer to those attributable to older vehicles than what is suggested by official pollution standards (Carslaw et al., 2011, Chen and Borken-Kleefeld, 2014). Along with evidence on entrances, this imply that the initial decrease in pollution for *Ecopass* is mainly due to owners of old vehicles not driving into the city center to avoid the fee.[17]

ing vehicles registered in Milan city was between 715,413 and 736,897 units, the number of unique vehicles entered in the city center in 2008, 2012 and 2013 was respectively 1,306,201, 1,014,980 and 1,075,662. Even assuming that all vehicles registered in Milan entered at least once in the city center, more than one fourth of the entrances is due to vehicles registered outside Milan.

[17]For instance, in 2007 the share of vehicles registered in Milan city as Euro 4 and Euro 5 was

The main conclusion is that while *Ecopass* was targeted at the correct vehicle groups (older vehicles), it overlooked trends in new vehicle sales. Also, new diesel vehicles were treated as new gasoline vehicles and therefore there was no incentive for drivers to change their preferences towards gasoline vehicles. On the other hand, under *Area C* new vehicles are subject to the fee as well.

Another important characteristic of both policies is that motorcycles are exempt from the fee. In the provinces of Milan and Monza, the share of motorcycles was increasing even before the introduction of *Ecopass* (Figure 2.6). The ratio between motorcycles and cars is increasing across time: in the province of Milan and Monza in 2005 there were 143 motorcycles for every 1000 cars (167 in Milan city), since then the ratio increased constantly, getting to 177 motorcycles for every 1000 cars in 2012 (216 in Milan city). This suggests that over time more drivers were able to circumvent the policies, thus making them less effective in the long run. Other studies on *Area C* suggest that the policy increased motorcycle use (Gibson and Carnovale, 2015, Percoco, 2014) and likely the same phenomenon occurred during *Ecopass*.

A final consideration concerns the effect of *Ecopass* and *Area C* on the distribution of traffic flows within the city center. As with most of the other existing policies, both systems are focused on the number of vehicle entrances but not on the traffic density within the area.

Estimating the causal effects of the charges on driving intensity within the

29.4%, in 2008 it was 36.4%, in 2009 42.7% Similarly, in 2011 the share was 52.7% but in 2012 it was 55.9%.

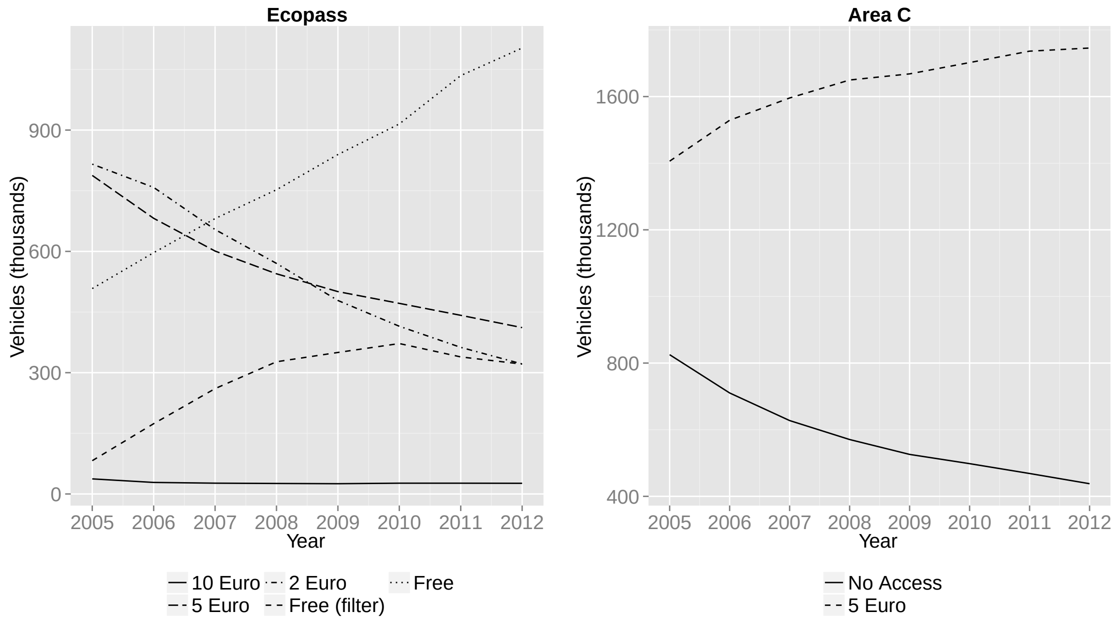

Figure 2.5: Passenger vehicle stock by *Ecopass* and *Area C* charge. Source: Automobile Club D'Italia

city center is a very difficult empirical question, and beyond the scope of the paper. However, there is some suggestive evidence that some variation in traffic flows might have occurred due to the policies. An official report from the *Ecopass* committee shows both positive and negative changes in traffic intensity in certain streets within the city center after the introduction of *Ecopass* (Commissione Ecopass, 2010, p.21). The proposed explanation is that with decreased congestion, exempt drivers and residents had an incentive to drive more in the inner city rings within *Ecopass*, instead of the outer city rings outside *Ecopass*. While it is not clear if and how much traffic flow distribution affected the concentration of NOx, it is useful to consider such dynamics when looking at the results of the paper.

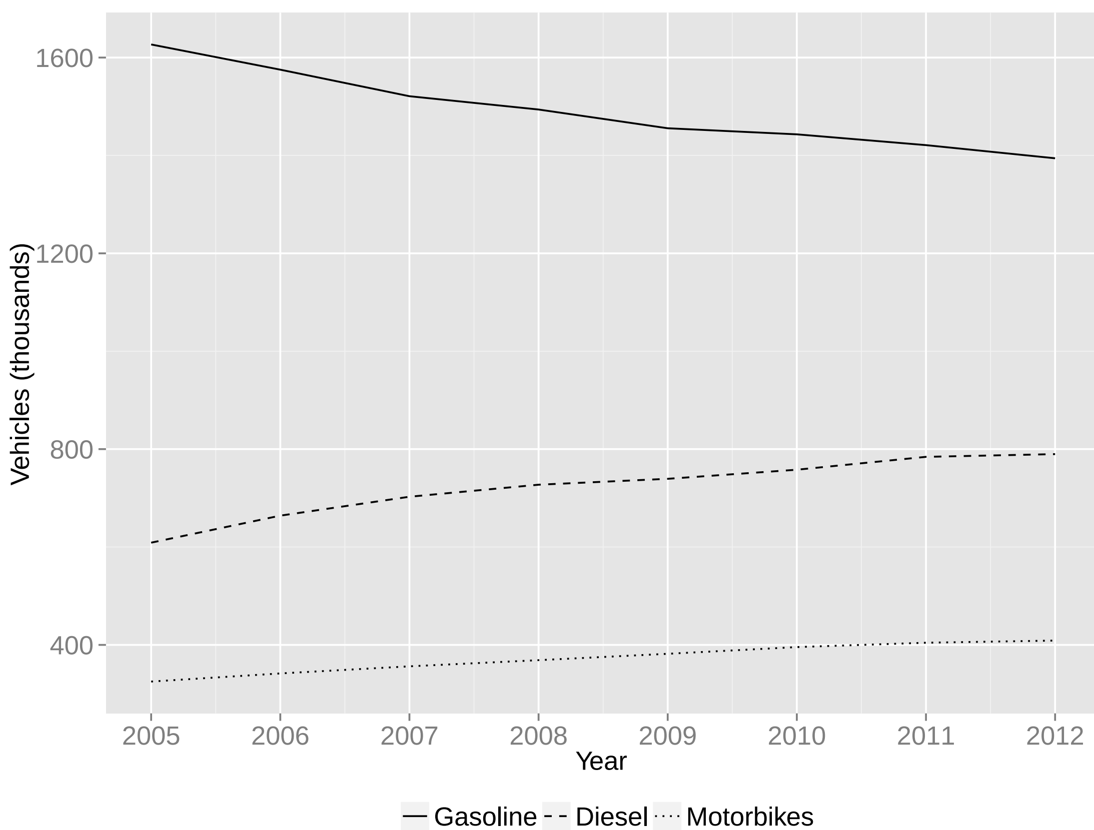

Figure 2.6: Passenger vehicle stock by fuel, plus motocycles. Source: Automobile Club D'Italia

2.7 Conclusion

This paper analyzes the short and long-term impact of two different policies - a vehicle pollution charge (*Ecopass*) and a vehicle congestion charge (*Area C*) - on NOx concentration in the Milan city center. Using a difference in difference specification I estimate a change in NOx concentration of -8.6% under Ecopass and between -5.1% and -8.0% under Area C in the first quarter of their introduction. I also find that, in case of *Ecopass*, the long-term effect is much weaker than the short-term, while the difference between short and long term is less pronounced for *Area C*. A comparison with an event study specification using an exogenous suspension of *Area C* shows an effect of similar magnitude. In term of seasonal effects, I see some variation between different quarters of the year, generally non statistically significant, and a much stronger effect at the very beginning of the two policies.

Results are consistent with previous literature finding a decreasing effectiveness of driving bans. Vehicle charges suffer of the same problem, but certain policy designs seem to perform better than others. Because the effect of both policies varies across quarters and years, even a well identified effect in a specific month or quarter might not be representative of the overall performance of the policy. There is a potential complementarity between the methodology used in this paper and identification strategies, like event studies, looking at the treatment effect around a certain threshold.

The main lesson is that very often policy interventions that tax or limit vehicle circulation do not consider drivers response to the policy or that the vehicle stock changes with time. In the well-known case of Mexico City, drivers were keeping older vehicle to bypass the ban. In the case of *Ecopass*, at the beginning owners of polluting vehicles entered the city center less often due to the fee, but once they adopted new, exempt vehicles they started circulating again. Across time the vehicle stock had a higher share of diesel vehicles, which emit more NOx, and because *Ecopass* did not discriminate between new diesel vehicles and new gasoline vehicles, the pollution charge was even less effective in the long run. On the other hand, under *Area C* getting a new vehicle does not allow to circulate freely in most cases, and excluding the first quarter, the long-run effect is closer to the short-run effect than the case of *Ecopass*.

A broader and more speculative interpretation is that in case of feasibility or political economy considerations preventing the regulator from making modifications to the charge, a policy targeting specific types of vehicles - in the case of *Ecopass* older vehicles - might not have an impact on pollution concentration in the long run. On the other hand, the long run effect of policies charging every vehicle the same amount might be more predictable.

The analysis has some limitations. The identification of the effect of *Area C* is less clean than the case of *Ecopass*, because *Area C* started almost immediately after *Ecopass*. It is possible to identify the additional effect of *Area C*, but not the total effect. To give a better sense of the estimates, I propose two set of results based on two different pre-treatment periods which, under plausible assumptions, represent the lower and the upper bound effect.

This paper does not provide any considerations in terms of welfare change due to the two policies. NOx reduction is only a part of the total potential effects from *Ecopass* and *Area C*, including change in other pollutants, change in congestion, change in vehicle stock, change in the number of accidents, and spatial and temporal spillover effects. In part, this question has been already addressed by Danielis et al. (2011), who looked at the costs and benefits for the first three years of *Ecopass*. The net benefits remain positive, but unchanged in the long run.

Chapter 3: Taxing vehicle miles traveled: The Traffic Choices Experiment in the Puget Sound region

3.1 Introduction

Driving creates a number of externalities, including congestion, emissions of conventional air pollutants and greenhouse gases, road wear and tear, and accidents. In principle, these externalities could be corrected by imposing a tax on drivers. In the past century, motor fuel taxes were the most common fiscal instrument of this type. Starting from the early 2000s, local and national governments around the world introduced a set of measures broadly defined as "road pricing," in which drivers are charged according to how much, when and where they drive.[1]

Within road pricing schemes, taxes on vehicle miles traveled (VMT) are one of the most recent measures. A VMT tax charges drivers for the miles actually traveled with their vehicle, either at a fixed rate per mile or at a variable rate according to various criteria, like household or vehicle characteristics, type of road, day of the week, and hour of the day.

[1] Some of the most widespread interventions of this type are congestion charges, adopted in Europe and Latin America and under consideration in various US cities, and dynamic road pricing, adopted in several US highways.

The motivation for the introduction of a VMT tax is twofold. First, such a tax helps mitigate major vehicle externalities like emissions, accidents, congestion, road wear and tear. Second, it guarantees a reliable source of revenue to local administrations for infrastructure construction and maintenance.

Currently, no VMT tax has been adopted on a large scale for passenger and light-duty vehicles, but various US states, including Oregon and California, have introduced or considered pilot programs.[2] In addition to that, a series of reports by state and federal government institutions in the US have discussed at length the rationale and the feasibility of a VMT tax (Council of State Governments, 2010, National Surface Transportation Infrastructure Financing Commission, 2009, Washington State Transportation Commission, 2014, Virginia Department of Transportation, 2008, Maryland Department of Transportation, 2011, Congressional Budget Office, 2011).

In this paper, I examine the effectiveness of the "Traffic Choices" VMT tax trial, implemented by the Puget Sound Regional Council and the US Federal Highway Administration in the Seattle, WA, metropolitan area between 2005 and 2006. Prior to the beginning of the trial, the driving of the 276 participant households was monitored through GPS devices installed in their cars. The information collected included miles traveled, time of travel, type of road and travel purpose. After three

[2]Starting on July 2015, the state of Oregon introduced a program on a volunteer basis with a flat VMT tax of 1.5 cents per mile and a reimbursement of the gasoline tax, open to 5,000 participants and based on GPS and odometer systems. California approved the adoption of a pilot VMT tax program starting from 2017.

months of data collection, participant households received a monetary endowment based on their past travel behavior and were charged a toll on the number of miles driven, ranging from 5 cents to 50 cents per mile. The toll rate varied depending on the hour of the day, the day of the week and the type of road. The toll cost was subtracted to the endowment and after seven months of trial period the participant could keep the amount that was left.

The Traffic Choices Experiment is one of the very few occasions in which a VMT tax has ever been introduced so far. There is a vast literature studying the effect of such a tax on driving, its comparison with alternative policies, and its distributional effects (West, 2004, Parry and Small, 2005, Safirova et al., 2007, Parry and Timilsina, 2010). However, these results are generally based on simulation exercises using survey data. There is little empirical evidence on how drivers would respond to an actual VMT tax.

Another important characteristic of the trial is that its average toll rate per mile is much higher than what current and past pilot programs and experiments have adopted, and it is very close to what economic literature suggests to be the optimal uniform VMT tax (Parry and Small, 2005, Safirova et al., 2007).

I use the data from the trial to answer three main research questions. First, what is the impact of the toll on the number of miles driven, and how long does the effect last? Second, does this effect vary with the type of road and travel purpose? Third, is there evidence of heterogeneous effects across different participant household characteristics?

To answer these questions, I exploit the variation in toll rates across hours, days and types of road. Because there is a baseline period with no toll but all participants are subject to the toll in the trial period, I take confounding factors like seasonality into account using data on traffic volume from traffic monitoring stations, which should not be affected by the participants' behavior. I use the treatment effects for each household and information on their characteristics to understand how factors like number of cars, number of children and income and past travel behavior affect the response to the toll.

Results show a reduction of miles driven by -6.83% for a 10 cents per mile toll during the first week of the trial period. However, the average treatment effect is extremely short lived and disappears after one week, with the exception of miles driven on highways. The initial response seems to be caused by a reduction in miles driven during weekends and in travels from work to home locations.

I suggest two possible behavioral mechanisms to explain the short-term response. In the first mechanism, the introduction of the VMT tax acted as a "cue" for drivers to pay attention to their miles driven. However, as previous literature shows, the effect of those "cues" tend to vanish pretty rapidly unless the "cues" are resumed (Allcott and Rogers, 2014). The second explanation is that the beginning of the trial serves as a learning period in which participants try alternative transportation modes, routes and times, weighting cost savings from paying a lower toll against the costs and the inconvenience of seeking an alternate route. If the alternative travel plan does not compensate for the cost of rerouting, participants would switch back to their original behavior.

The short-term effect of the toll also warn about the risks of drawing long and medium term conclusions based on immediate effects of policy interventions.

There is evidence of a strong heterogeneity in the effect of the toll system. In particular, households with high cost of time, like households with children, and households with high income tend to respond less to the toll, while households who are driving more and at more expensive toll times or on highways are more likely to reduce their miles driven. Such heterogeneous effects last even after the first week of introduction of the toll.

The paper is organized as follows. Section 3.2 reviews the literature and actual VMT tax programs. Section 3.3 presents the Traffic Choices trial. Section 3.4 describes the data. Section 3.5 explains the methods. Section 3.6 presents the results and Section 3.7 concludes.

3.2 Literature review

The economic and transportation engineering literature have studied VMT taxes from a theoretical and empirical standpoint, especially in the context of welfare and distributional consequences and their comparison with other policy instruments. Parry and Small (2005) present a calibration exercise in which they estimate the optimal gasoline tax and VMT tax for the US and the UK, finding an optimal VMT tax of 14 cents/per mile in the US and 15.5 cents per mile in the UK. In the US, switching from the optimal fuel tax to the optimal VMT tax would increase revenue by 2.5 times, and would increase welfare gains by nearly four times.[3] An analysis for Mexico City by Parry and Timilsina (2010) suggests an optimal VMT tax of 20.3 cents per mile which would increase welfare by \$110.8 per capita and reduce auto mileage by 24.8%. Safirova et al. (2007) use a general equilibrium model calibrated to the Washington DC metropolitan area to compare the welfare effects of various types of road pricing, estimating an optimal rate of 14.59 cents per mile and a reduction of miles traveled by 26.2%. In terms of welfare, a fixed rate VMT tax performs better than congestion charges and freeway tolls, and produces almost the same welfare gains than a comprehensive variable time-of-day pricing. A VMT tax would also be welfare enhancing compared to a gas tax in reducing accident

[3] Note that while the rest of the literature focus on transportation externalities in estimating the optimal VMT tax, in Parry and Small (2005) the estimate includes three components: an adjusted Pigouvian tax, a Ramsey tax and a congestion feedback. The Pigouvian tax accounts from 57% of the optimal tax rate.

externalities (Parry, 2004).

Some studies have proposed differentiated VMT taxes by vehicle and/or driver, where the rate is equal to the per-mile external costs (Parry, 2004). Such tax would be equivalent to a first-best tax on emissions if all the driver and vehicle specific determinants of emissions are known (Fullerton and West, 2002). In certain cases, a uniform or differentiated VMT tax can be introduced under the form of per mile premiums through Pay-As-You-Drive insurance schemes (Parry, 2005, Ferreira and Minikel, 2012).

With the introduction of alternative fuel vehicles, VMT taxes have to potential to control congestion and emissions externalities in a better way than fuel taxes. The role of VMT taxes for ethanol and gasoline blends and electric vehicles has been studied by Khanna et al. (2008) and Holland et al. (2015), showing that VMT taxes differentiated by fuel type or in addition with other fuel-based instruments can improve welfare compared to the current situation.

A widely studied topic in this literature is the distributional impact of a VMT tax, especially when compared to a gasoline tax. West (2004) and West (2005) use a discrete choice model with travel and consumer expenditure data to represent vehicle choice and mileage under a gas tax or a VMT tax, a tax on vehicle size or subsidies to new vehicles. The main finding is that under a uniform VMT tax high-income households tend to have a lower response to a VMT tax than low income households, but because poor households are less likely to own a vehicle, a VMT tax is regressive only for higher income households. The VMT tax is less regressive than taxes on size or subsidies for new vehicles. Evidence of the regressivity of a

VMT tax is found also by Sana et al. (2010) and Zhang and Lu (2012).

Various papers compare the distributional effects of a VMT tax to those of a gasoline tax. In most cases, the VMT tax is found to be equally or less regressive than a gasoline tax (Robitaille et al., 2011, Weatherford, 2011, Larsen et al., 2012, Paz et al., 2014). When the VMT tax is found to be more regressive than a gasoline tax, the difference is very small (McMullen et al., 2010).

Because the introduction of VMT taxes is relatively recent, most of the current literature simulate the effect of a VMT tax using travel survey data like the US National Household Travel Survey. However, there are some examples of studies analyzing the effect of treatments very similar to a VMT tax, generally in the form of discounts on insurance premiums. Agerholm et al. (2008), Bolderdijk et al. (2011), and Hultkrantz and Lindberg (2011) look at the impact of a Pay-As-You-Speed scheme, in which participants receive a compensation if they decrease the time driving while exceeding speed limits. The incentive scheme pushed drivers to reduce their driving speed, but a reduction is found also among participants who were informed whenever they exceeded speed limits but received no monetary incentives. Reese and Pash-Brimmer (2009) and Bolderdijk et al. (2011) performed experiments in which participants were compensated whenever they decrease miles driven by a given amount. Results are mixed, showing a reduction in miles driven between 3.2% and 5.7% in one case (Reese and Pash-Brimmer, 2009), but no reduction in another (Bolderdijk et al., 2011).

So far, only Abou-Zeid et al. (2008) studied the effect of an actual VMT tax on miles driven.[4] They use data on trips of 130 participant households in a baseline period without any tax and a trial period in which each household was assigned in a group with a different VMT tax rate, including a control group with no tax. Using a difference in difference approach, they found a decrease in miles driven by treated households, but the results are not statistically significant, probably as a result of the small sample size.

As the review of the literature suggests, evidence on how drivers would respond to VMT taxes is mixed at best. While several studies have shown the potential welfare benefits of switching in part or completely to a per-mile charge, the analysis of an actual VMT tax scheme would allow to look at treatment effects during different hours of the day or on different roads, the evolution of the response across time and presence of heterogeneity in the change in miles driven. In order to do so, I am looking at an experiment with differentiated mile pricing by hour, day and type of road in the Puget Sound region, WA, in 2005.

[4] Also Hanley and Kuhl (2011) implemented an experimental VMT tax, but they focus on the technical feasibility and public acceptance of the scheme, and not on the impact on miles driven.

3.3 The Traffic Choices Experiment

The Traffic Choices experiment was organized by the Puget Sound Regional Council, a metropolitan planning organization based in Seattle, WA, and the Federal Highway Administration. They installed On-Board Units (OBU) on the vehicles of 276 volunteer households living in the Puget Sound area (Figure 3.1). These devices monitored the trips of each vehicle using GPS technology throughout the duration of the experiment. The experiment started on April 2005 and ended on February 2006: in the first phase, between April 2005 and June 2005 (baseline period), trip characteristics of the participating vehicles were recorded without imposing any kind of toll[5].

The trial period started in July 2005. Households received a travel budget, a monetary endowment in a dedicated account that was used to pay the tolls. The budgets assigned to each household were substantial, much higher than similar experiments, ranging from \$791 (the minimum amount possible) to \$4,116. The amount assigned for each household was based on past travel history and the degree of uncertainty on predictions of future miles traveled. The main concern was to make sure that the endowment was large enough to last until the end of the trial, so the amounts were always higher than what strict predictions based on past travel history would have suggested.

[5]The beginning of the monitoring period was supposed to start in early 2005, but due to a technical problem most of participants have substantial breaks in data collection before April 2005. In this paper I ignore the data gathered in the previous months.

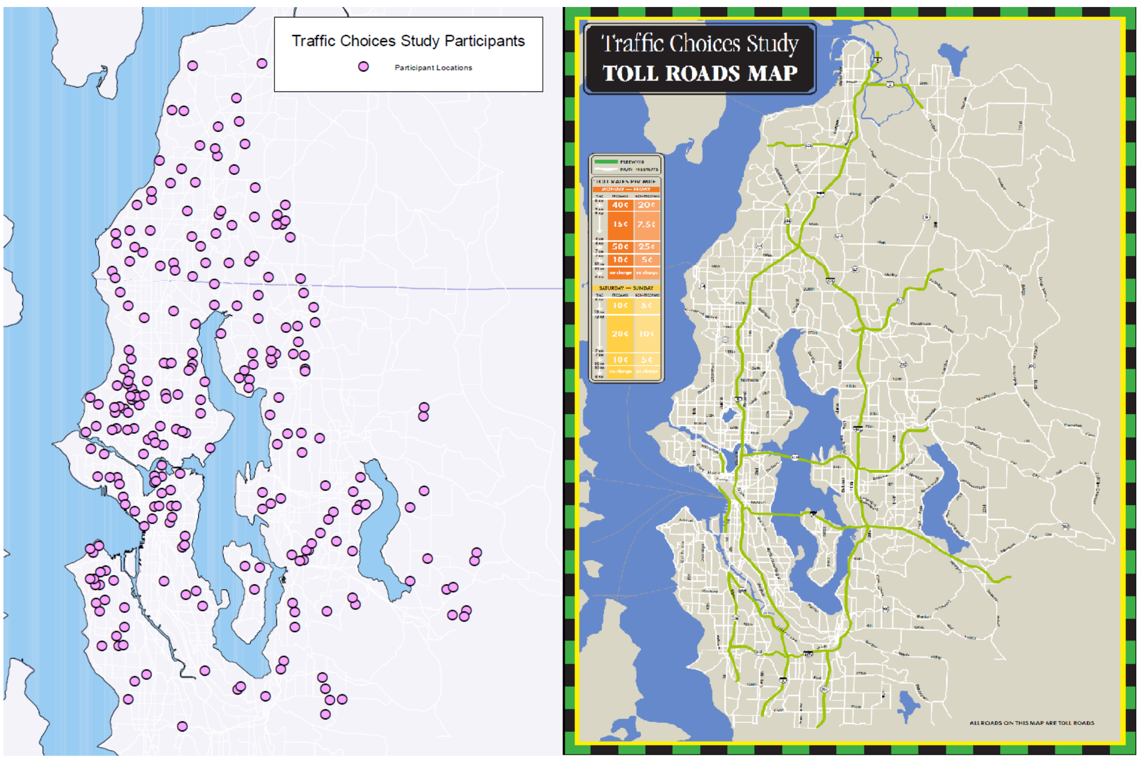

Figure 3.1: Map of tolled roads and participants' location in the Traffic Choice experiment. Source: Puget Sound Regional Council (2008)

Throughout the trial period, households were charged a road toll. Such toll was a tax per mile driven based on the day of the week, hour of the day, and type of road (Table 3.1). The main principle was that the charge was higher at hours with higher traffic volume. Highways charges were twice as much as normal roads. There was no charge for miles driven in downtown Seattle, due to the technical difficulties measuring the exact miles traveled.

Mon-Fri			
	Highway	Non-highway	Downtown Seattle
10pm-6am	$0	$0	$0
6am-9am	$0.40	$0.20	$0
9am-4pm	$0.15	$0.075	$0
4pm-7pm	$0.50	$0.25	$0
7pm-10pm	$0.10	$0.05	$0
Sat-Sun			
	Highway	Non-highway	Downtown Seattle
10pm-6am	$0	$0	$0
6am-10am	$0.10	$0.05	$0
10am-7pm	$0.20	$0.10	$0
7pm-10am	$0.10	$0.05	$0

Table 3.1: Traffic Choices toll rates per mile.

The toll scheme had much higher rates per mile and much more variation across time and road type than current pilot programs and trials done in the past.[6] Average rates were also closer to the optimal uniform rate of 14 cents per mile suggested by Parry and Small (2005) and Safirova et al. (2007).

[6]The Oregon pilot program has a uniform toll of 1.5 cents per mile. The trial described in Hanley and Kuhl (2011) had rates ranging between 2.19 and 0.33 cents per mile. Rates in Abou-Zeid et al. (2008) were between 5 and 25 cents per mile, but each household was subject to either a unique rate or two different rates.

At the end of the trial, participating households kept the money left in their travel budget. Negative balances were not allowed and every households received a minimum compensation of $150 at the end of the experiment. However, only a very small percentage of the household spent all their travel budget.

While participating households were informed of the scope of the trial since the enrollment phase, they were not aware of the key details. In particular, they did not know that the travel budget and the toll rates were based on their previous travel behavior or any other element they could have strategically manipulated. Participants were informed that the baseline period was a technical trial rather than a period in which their travel decisions were actually part of the experiment. The participants were not aware of the structure of the toll or the amount of their travel budget until June 2005, one month before the starting of the trial period.

Participants were recruited through phone calls during Fall 2004. Initially, 196,451 phone numbers were contacted. Among the 43,465 household reached who were eligible for the experiment, 776 were shortlisted to participate to the experiment. Eventually, 276 households were enrolled at the beginning of the operations (Table 3.2). A comparison with the 2001 National Household Travel Survey (NHTS), limited to households in the Seattle metropolitan area with at least one vehicle and one working family member, shows strong similarities between the two samples for most of the observable characteristics.[7]

[7]Respectively for the Traffic Choices sample used in the analysis and the NHTS sample, the share of drivers employed partial or full time is 75.89% and 76.78%, the average number of vehicles per household is 1.61 and 1.73, the share of household owning their house is 80% and 76%, the average number of children per household is 0.54 and 0.78, the median education is 16 years and

Table 3.2: Data on recruitment and participation

Category	N. Households
Total enriched call list	196,451
Phone numbers called	126,796
Working, non-business phone numbers	99,267
Household member reached	53,229
Household with eligible characteristics	43,465
Household preliminary enrollment	776
Household installed all toll meters	307
Household fully enrolled at start of operations	276
Household awarded non-zero compensation	239
Household with compensation > \$150	228

Source: Puget Sound Regional Council (2008) and additional information provided by the authors of the trial

Due to budget constraints and technical requirements, there was a limit on the number of households who could be enrolled in the trial. Therefore the focus was on a subset of households, based on a list of recruitment goals about the ideal sample composition. Table 3.3 shows the list of recruitment criteria: in particular, they showed preferences towards households who used their vehicle to commute to work in peak hours, with no current experience with carpooling and with no plans of changing job, location or vehicle during the experiment. These criteria were used as a guideline to select the final list, based on observable household characteristics correlated with the criteria.[8] The selection process occurred right

associate degree, the median household income per year in 2001 dollars is 63,339 and 55,000-59,000, the average age is 45.95 and 40.90, and the share of female drivers is 61.31% and 52.10%. Some differences are explained with the fact that the Seattle metropolitan area in the NHTS is larger than the area of the trial.

[8]For instance, it was assumed that households with lower income and higher number of workers were more likely to carpool in response to a road toll.

after the households were contacted for the first time.

After the households accepted to enroll into the experiment, all their vehicles were equipped with OBUs starting from November 2004. During the baseline period OBUs displayed only the name of the road, while during the trial period they displayed the name of the road, the cost per mile of the road and the cumulative cost of the trip (Figure 3.2). While the installation of all the OBUs was mostly complete by February 2005, a technical issue required all the vehicles to be recalled in March 2005 for a software update. This created a break in the data collection, and it was decided to extend the baseline period until the end of June 2005.

Figure 3.2: On Board Unit. Source: Puget Sound Regional Council (2008)

The OBUs provided information on the position of the vehicles at any given time. The OBUs measured also the time in which the engine was turned on and turned off. Using this information, the system was able to match in real time the position of the vehicle with the road, calculate the number of miles driven, identifying the type of road and assigning the correct toll rate.

Once the toll was introduced, participants were able to monitor their driving costs by accessing an online account in a dedicated web page. The account showed

Table 3.3: Recruitment criteria

Variable	Desired share of sample
Number of households with 3 vehicles	5%
Number of households with 2 vehicles	40%
Number of households with 1 vehicle	55%
Proportion of households already carpooling	0%
Proportion of households with transit accessibility	50%
Proportion of households with at least one worker commuting in peak period and direction on congested facilities within study area	100%
Proportion of households with second worker commuting in off-peak period and direction or outside of study area	Max 20%
Proportion of households with installation-compliant vehicles	100%
Proportion of households with plans to purchase additional vehicles in study period	0%
Proportion of households with plans to move in study period	0%
Proportion of households with likelihood to change employment status in study period	0%
Final maximum number of OBUs installed	500

Source: Puget Sound Regional Council (2008) information on past travels like time, name of segment of tolled road, miles driven, tax per mile and total cost. Participants could also access their account balance showing how much travel budget they had left. At the beginning of each month participants received a monthly invoice through email, showing the monthly costs, the previous and the current balance, and a summary of the trip records. The balances could be visualized both at the vehicle and at the household level. Participants had also access to a help desk phone number in case of any malfunctioning issues of the OBUs.

3.4 Data

Most of the data used in the analysis come from the National Renewable Energy Laboratory (NREL) - Transportation Secure Data Center. It contains a wide range of information on travel behavior and participating household characteristics. The final dataset used in the analysis includes 210 households. In particular, I removed from the sample households who left the experiment before the end, households with adjusted final compensation due to gaps in data collection, households not included in the original recruitment round, household who changed their vehicle during the experiment and households with a baseline period length lower than three months (April-June 2005), for a total of 66 households excluded.[9]

While the home address of the participants was known, explicit information on the purpose of a trip or the identity of the driver for a trip was not available. However, it was possible to identify two particular subgroups of destinations: "work destinations" and "home destinations." Work destinations included any frequent location not corresponding to home.[10] Those locations were matched manually with spatial data on land use and employment to distinguish work locations from second homes.

Trips were aggregated in *tours*, defined as a group of trips starting and ending either at home or work locations: thus, they distinguished between home-work,

[9] Appendix B.1 contains more information on the data and the data cleaning process.

[10] That implies that work destinations might correspond to workplaces, schools or volunteer activities.

work-home, home-home and work-work tours.

Table 3.4 shows information about drivers characteristics, collected during the recruitment phase. Most of the drivers are employed and commute to work most of the week using their own car without carpooling. A small percentage of them use public transit or carpooling before the beginning of the experiment.

Table 3.4: Summary statistics, drivers

Variable	Mean	Median	Std. Dev.
Age	45.95	47.00	13.27
Years of education	15.35	16	4.00
Share females	61.31%		
N. times commuting per week	6.03	5	9.56
Share commuting alone at least once a week	85.41%		
Share commuting with transit at least once a week	4.46%		
Share commuting with carpooling as passenger at least once a week	5.95%		
Share commuting with carpooling as driver at least once a week	12.5%		
Share full time employed	66.37%		
Share part time employed	9.52%		
Share students	5.36%		
Share homemaker	9.82%		
Share retired	6.55%		
Share unemployed	2.24%		

Table 3.5 shows information about participating households. The vast majority of the households have the same number of drivers and vehicles. Household income varies considerably across households, ranging from \$10,000 to \$300,000. Most of the households believes that road tolls should cover a considerable part of infrastructure spending, and at the same time they are fairly concerned about

privacy implications of road tolls.

Table 3.5: Summary statistics, households

Variable	Mean	Median	Std. Dev.
Number of drivers	1.60	2	0.66
Number of vehicles	1.61	2	0.67
Number of kids	0.54	0	0.864
Share households with kids	32.38%		
Age household head	47.97	47	12.21
Household gross income 2004	75500	67560	44444
Percentage paid by tolls	39.25	40	30.81
Concern for privacy (1-7)	3.98	4	1.94
Share home owners	80.00%		
Share renting	17.61%		
Household endowment	1959	1745	996.46
Vehicle endowment	1302	1167	551.68
Final household compensation	801.9	719.3	510.61
Compensation as share of endowment	0.4311	0.4178	0.1921

Respondents are the heads of household. *Percentage paid by tolls* is the share of road infrastructure budget revenue that should be provided by pay as you drive fees like road tolls or gas taxes versus general taxation. *Concern for privacy* is the degree of concern about the privacy implications of a toll system involving the collection of individual vehicle road use data, with 1="Not concerned at all" and 7="Very concerned".

The average initial endowment per vehicle corresponds to $1302, while at the household level the average initial endowment is $1942. The minimum endowment for each vehicle is around $791. Household compensation, calculated by subtracting toll costs to the initial endowment, is considerable: the average household got $797.5 at the end of the experiment (roughly $100 per month), corresponding on average to 43.19% of the initial endowment. The minimum final compensation was of $150.[11]

[11] Endowment and compensation data are not contained in the NREL dataset and were provided

Figure 3.3 shows the relationship between endowment, final compensation, income and driving behavior in the baseline period. Final compensation is strongly correlated with initial endowment (Panel A). However, households with a lower initial endowment received a larger share of it as compensation (Panel B). Household with higher income generally received a higher initial endowment (Panel C). Finally, the initial endowment is strongly correlated with the hypothetical toll costs in the baseline period if the toll was introduced, but the correlation is not perfect (Panel D). This occurs because initial endowment was based on driving behavior before the trial, but for some households there were travel data available even before the official baseline period (April-June 2005).

Figure 3.4 offers some information on the evolution of participants' driving behavior before and after the introduction of the toll. The average toll amount per trip is generally around 75 and 85 cents (Figure 3.4.A). There is a first sharp decline in average toll in the month immediately after the beginning of the trial, then a rebounding and another decline from September to December. As we approach to the end of the trial, average tolls start to rise again.

Average tolled miles (normal roads and highways) per trip follow a similar pattern (Figure 3.4.B): a decline in July, a bounce back in August and another decline until December. Because the average tolled miles per trip are between 3.9 and 4.3 miles, the average toll per mile is roughly around 20 cents.

Figure 3.4.C shows the evolution of average highway miles, which are taxed twice as much as normal roads. The pattern follows pretty well figure 3.4.A and

directly by the authors of the trial.

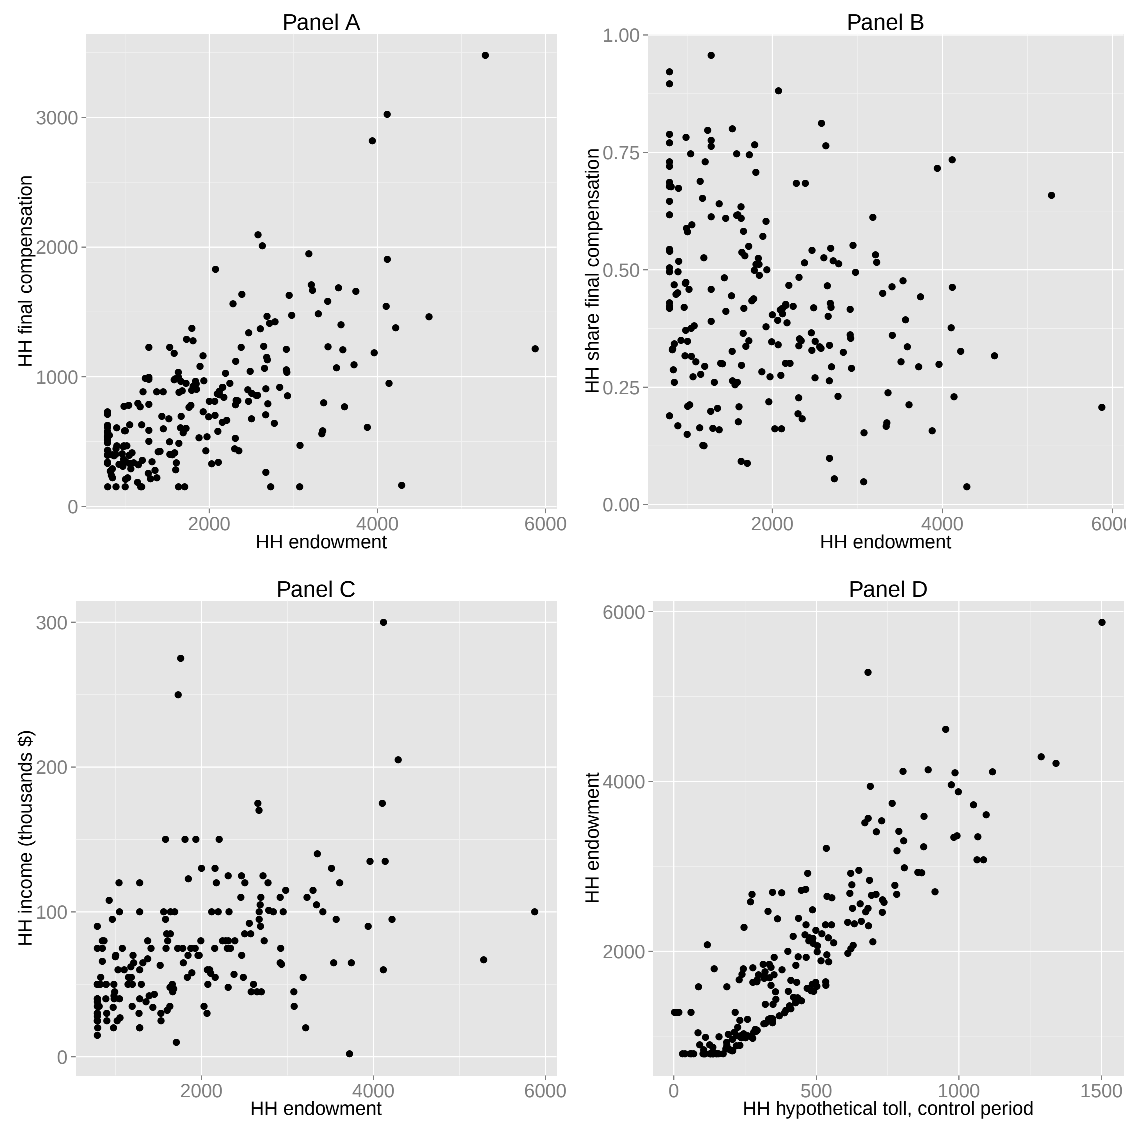

Figure 3.3: **Panel A**: Household initial endowment and final compensation; **Panel B**: Household initial endowment and compensation in share of endowment; **Panel C**: Household initial endowment and 2004 annual income; **Panel D**: Household hypothetical toll cost during the baseline period and initial endowment

figure 3.4.B, but in a more pronounced way. On average, highway miles seem to account for half of the tolled miles per trip, but there is considerable variance in the number of highway miles driven per trip.

Average trip duration looks quite stable across months, and the variation is not particularly large in magnitude (Figure 3.4.D).

Looking at the tour level, I focus on two variables that might change in response to the toll: participants might change the number of trips for each tour during high or low toll hours, and they might change the dwelling time between trips to strategically avoid driving in the high toll hours. Figure 3.4.E shows the evolution of the number of trips per tour. While there seems to be a negative correlation between number of trips and tolled miles driven in certain months, it does not seem to be a consistent pattern. Dwelling time within tour - i.e. time spent without driving - is particularly high in June, but changes little in most of the other months.

Another way to look at the evolution of driving behavior before and after the introduction of the toll is by observing the number of circulating participating vehicles in a given time. Figure 3.5 looks at the average number of participating vehicles on the road in a given minute. Averages are calculated for the three months before the trial (April-June, solid line) and the three months after the beginning of the trial (July-September, dashed line). Figure 3.5.A shows average circulation during weekdays (Monday-Friday). The period in which the toll is the highest, between 6am and 9am and 4pm and 7pm, is also the period in which circulation is higher. However, after the introduction of the toll circulation is decreasing. By

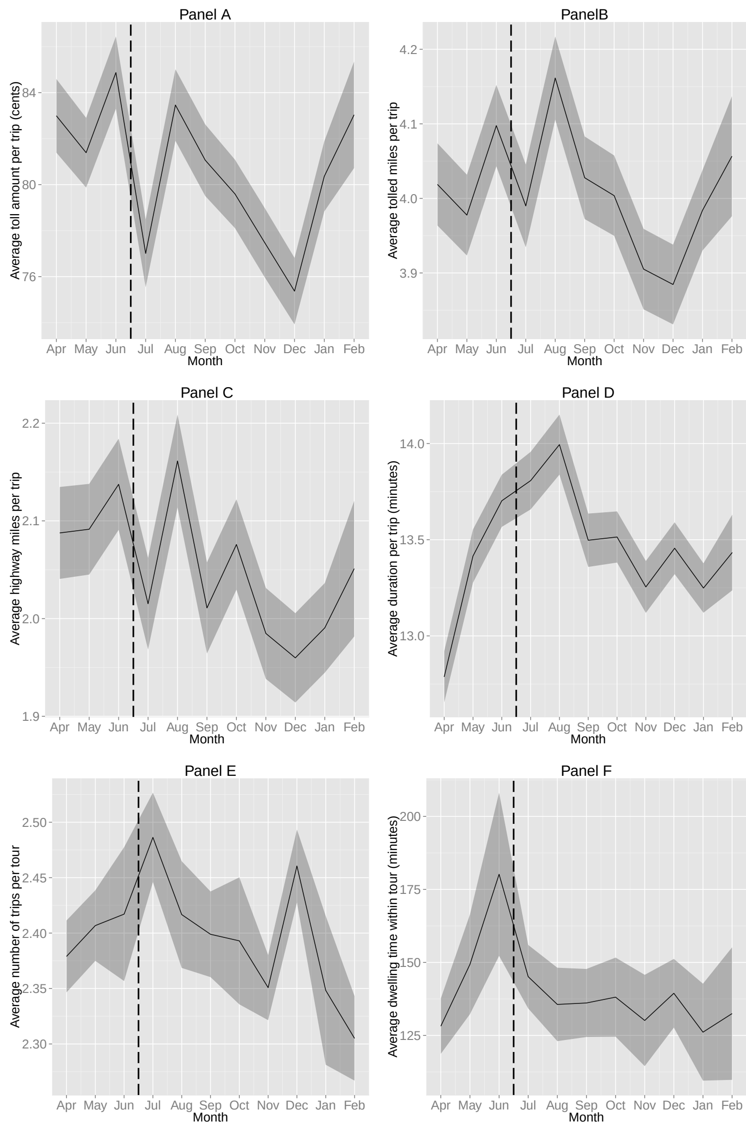

Figure 3.4: Monthly summary statistics at trip and tour level. Vertical dashed line shows the start of the trial period. Confidence intervals at 95%. **Panel A**: Average toll amount per trip (hypothetical toll if before trial); **Panel B**: Average miles tolled per trip; **Panel C**: Average highway miles per trip; **Panel D**: Average trip duration; **Panel E**: Average number of trips per tour; **Panel F**: Average dwelling time within a tour.

contrast, in the hours when tolls are lower, between 9am and 4pm and 7pm and 10pm, circulation increases. Figure 3.5.B shows the same graph during weekends (Saturday-Sunday): here circulation seems to increase slightly in the low toll hours (6am-10am and 7pm-10pm) but also in the high toll hours (10am-7pm).

The other panels of figure 3.5 show average circulation during weekdays divided by tour purpose (home-home, work-work, home-work, work-home).[12] Home to home tours seem to occur more often in the afternoon (Figure 3.5.C). After the beginning of the trial, circulation increases with the exception of the morning time. Work to work tours are concentrated in the mid afternoon (Figure 3.5.D), but there no systematic change in circulation before and after the beginning of the trial. Home-work (Figure 3.5.E) and work-home (Figure 3.5.F) are concentrated respectively between the 6am-9am and 4pm-7pm time frame, where the toll rate is the highest. In both cases, there is a reduction in peak circulation during the trial.

Displaying data at the trip level might miss important substitution patterns in response to the toll. For instance, people might travel fewer miles, but more frequently, or stopping driving during certain hours, certain days or in certain types of roads. To address this issue, I aggregate miles driven in tolled roads per hour at the weekly level, disaggregated by toll time band corresponding to a different toll rate and type of road.[13] Figure 3.6 plots the weekly tolled miles disaggregated by type of road and/or time, dropping the week when the trial starts. Figure 3.6.A and

[12]I focus on weekdays because during weekends tours belong mainly to the home-home type.

[13]For instance, a toll band is the average miles driven per hour between 6am-9pm during weekdays. See Appendix B.1 for further explaination on data cleaning and week-level aggregation.

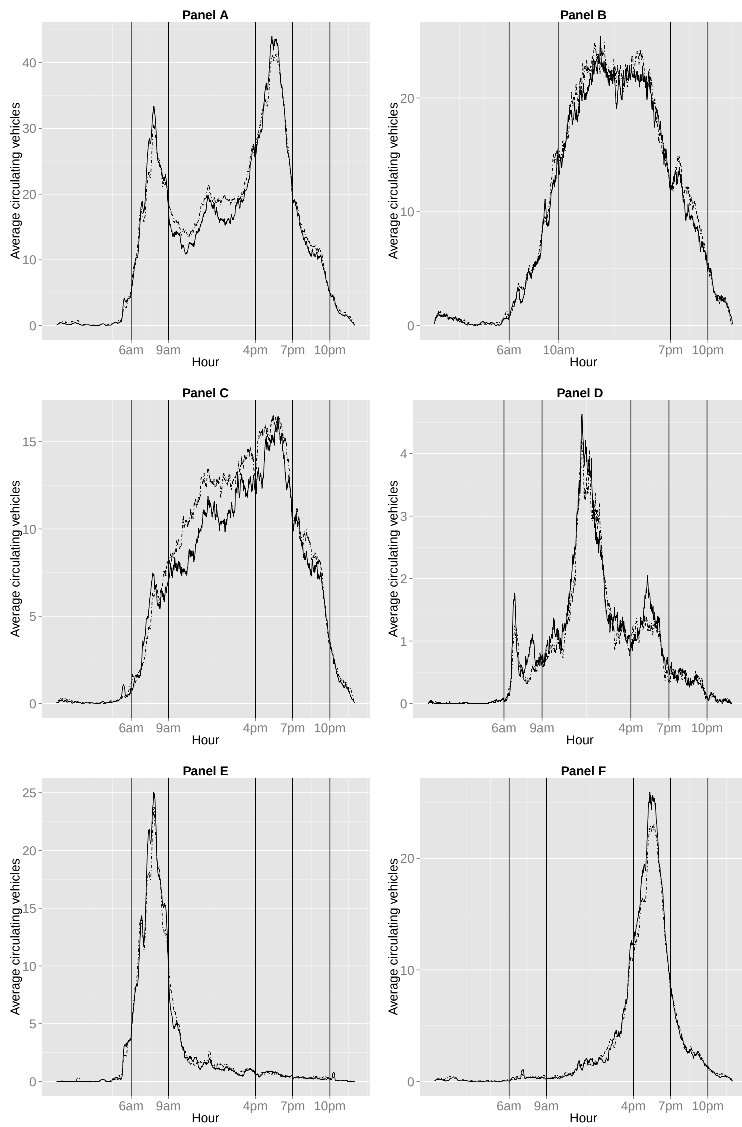

Figure 3.5: Average daily number of circulating vehicles per minute. The solid line is the average for the three months before trial (Apr-Jun 2005), the dashed line is the average for the three months after the start of the trial (Jul-Sep 2005). Vertical lines shows the boundaries for the different tax rates. **Panel A**: Monday-Friday; **Panel B**: Saturday-Sunday; **Panel C**: Monday-Friday, home-home tours; **Panel D**: Monday-Friday, work-work tours; **Panel E**: Monday-Friday, home-work tours; **Panel F**: Monday-Friday, work-home tours.

3.6.B show average weekly tolled miles in highways and normal roads respectively. In both cases there is a drop in miles driven in the first full week after the introduction of the toll. After that, miles driven tend to return to the previous level, especially in the case of “normal” roads.

At certain times of the day, tolls are much higher than average: figures 3.6.C and 3.6.D look at miles driven in highways and tolled normal roads between Monday and Friday, between 6am and 9am, and between 4pm and 7pm. During those times, the toll for normal roads was 20 cents/mile and 25 cents/mile respectively (the rate doubled for highways). I still observe a temporary drop in average miles driven at the beginning of the trial. For highways, the drop seems to persist longer, although it is not statistically different than the values in the baseline period.

I also look at weekdays (Monday-Friday) during times in which tolls are lower, in particular between 9am and 4pm and 7pm and 10pm, when the toll for normal roads was respectively 7.5 cents/mile and 5 cents/mile (the rate doubled for highways). Panel 3.6.E shows the results for highways and panel 3.6.F shows the results for normal roads. The reduction in miles driven at the beginning of the trial is much less pronounced and disappears completely when looking at normal roads.

One of the concerns is that driving behavior and traffic have a strong seasonal component (Memmott and Young, 2008). The difference in miles driven could be caused by the introduction of the tax, by some unrelated seasonal changes in driving behavior, or something else. In the absence of a control group, which would allow us to disentangle the latter two factors, I separate the two effects by looking at traffic flows data throughout the Seattle Area during the time of the experiment: I argue

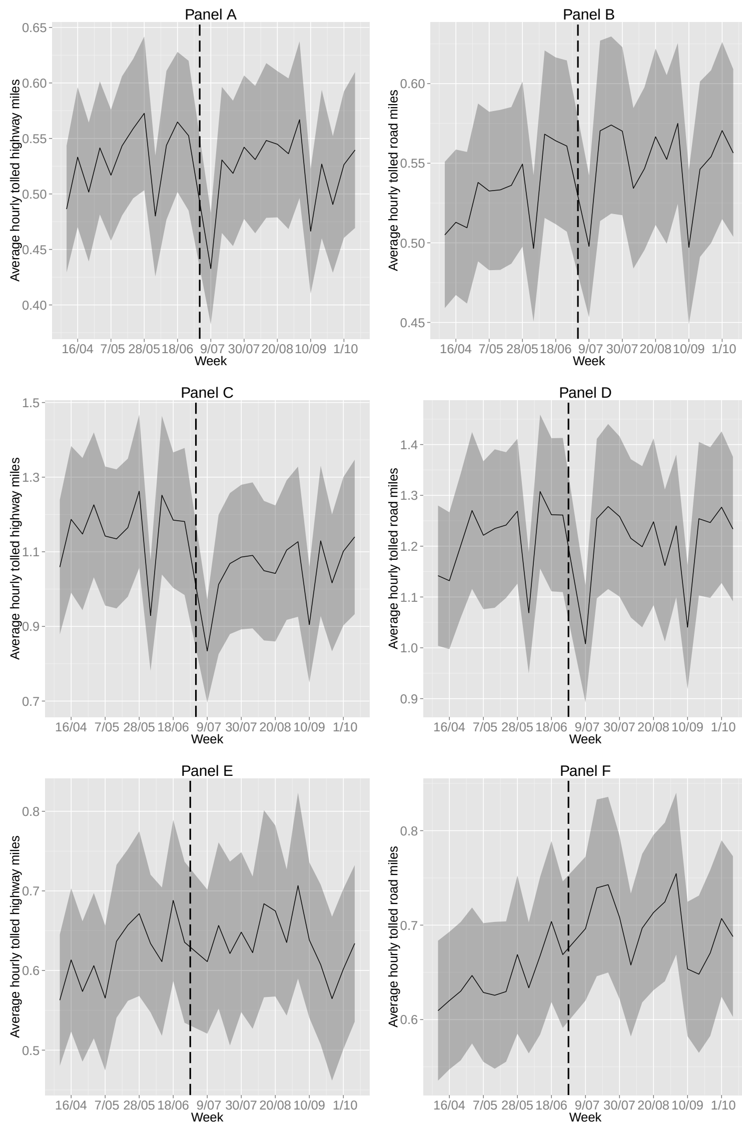

Figure 3.6: Average hourly tolled miles driven per week. Vertical dashed line shows the start of the trial period, the week in between has been dropped. Confidence intervals at 95%. **Panel A**: Average miles on highways; **Panel B**: Average miles on normal roads; **Panel C**: Average miles on highways during high toll hours (Monday-Friday 6am-9am, 4pm-7pm); **Panel D**: Average miles on normal roads during high toll hours (Monday-Friday 6am-9am, 4pm-7pm); **Panel E**: Average miles on highways during low toll hours (Monday-Friday 9am-4pm, 7pm-10pm); **Panel F**: Average miles on normal roads during low toll hours (Monday-Friday 9am-4pm, 7pm-10pm).

that, due to the small number of participants in the Traffic Choice experiment, its impact on actual traffic is virtually zero.

Figure 3.7 shows the weekly average traffic volume per hour obtained by aggregating Department of Transportation data from thirteen highway traffic monitoring stations in King County, WA, where the Traffic Choice experiment took place. Two facts immediately noticeable are the drop in traffic volume corresponding to the first week of July, when the trial period started, and the higher than average traffic flow during the months of July and August and in the last week of June.

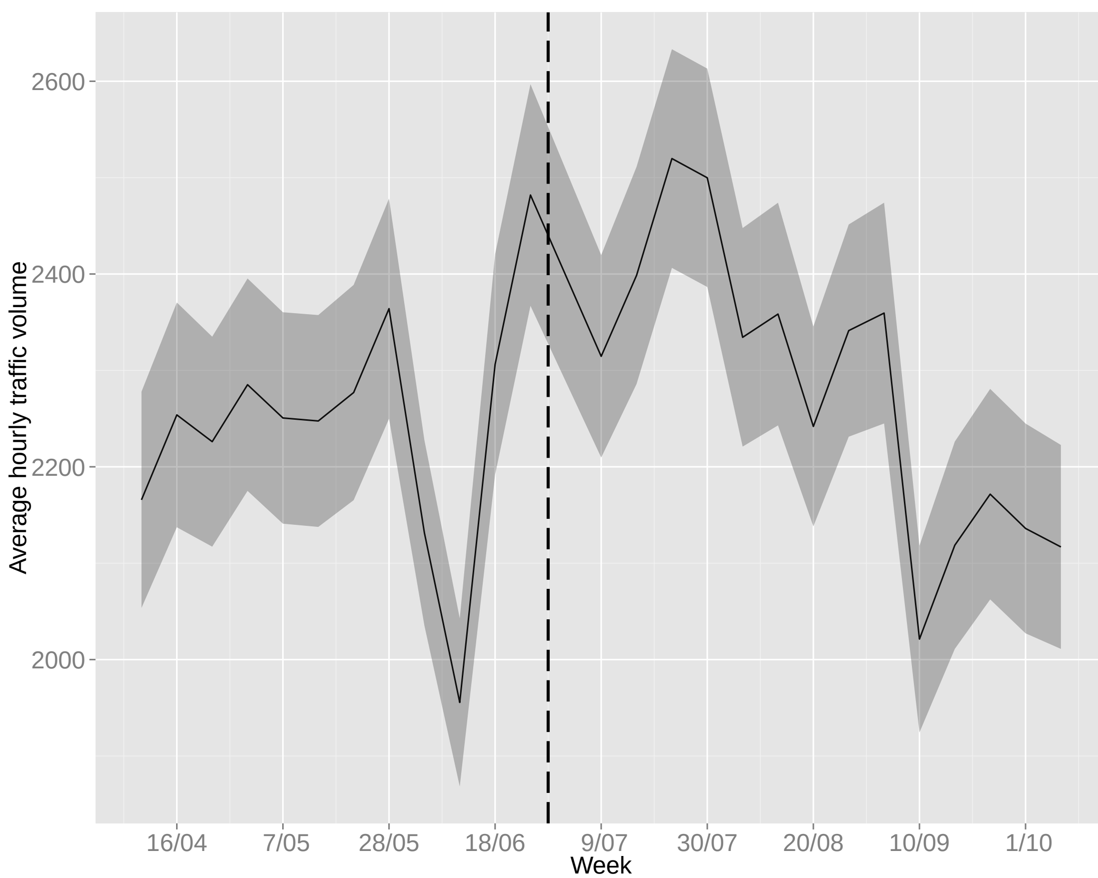

Figure 3.7: Average hourly highway traffic volume in King County, WA. Data from 13 traffic monitoring stations.

Another aspect to notice is that while average area-wide traffic volume and average participants' miles driven are highly correlated (0.913), traffic volume explains very little of the variation in miles driven compared to toll band differences and road type. When regressing average miles driven disaggregated by week, toll bands and type of road, toll bands-road type fixed effects alone give an R-square of 0.9782.

3.5 Empirical strategy

The main goal of the analysis is to assess the impact of the introduction of the VMT tax on tolled miles driven by the participants. To identify the effect, I exploit the variation in toll per mile before and during the trial and across different toll bands and types of roads. Such variation can be considerable, as during the trial period the toll could be as low as 0 cents per mile and as high as 50 cents per mile. I ask three research questions. How does the effect of the toll evolve over time? How does differ across different dimensions like type of road and travel purpose? I am also interested in any heterogeneity in response between different types of households, with particular emphasis on three characteristics: household income, driving behavior before the trial, and the size of the initial endowment.

Due to budget constraints, the Traffic Choice experiment did not include a control group of participants in the sample, implying that the actual response from the toll might be indistinguishable from changes in driving behavior due to seasonal variation or holidays. To address this problem, I use information on highway traffic gathered from thirteen monitoring stations in King County, WA, where the trial took place. Such data are particularly useful because the number of households enrolled in the trial is too low to have a noticeable impact on actual traffic. I assume that, for each toll band in each type of road, there is a constant, linear relationship between average traffic volume and average miles driven by the participants. The identification strategy relies in estimating that coefficient using the baseline period, and assuming that the relationship between traffic volume and participants' miles remains the same during the trial.

I wish to emphasize that any effect on miles driven caused by the introduction of the toll does not impact the general equilibrium. The household sample is too small to have any meaningful impact on overall traffic.[14] While this limits the capability of providing insights on what would happen if the toll was applied on a larger scale, it gives a precise outlook on how people changed their driving behavior as a direct consequence of the introduction of a toll, which would be much harder to do in a general equilibrium framework.

To estimate the impact of the VMT tax on tolled roads, I use the following specification:

$$m_{hrbt} = \alpha + \delta_{rb} + \eta_h + \beta_t \tau_{rbt} + \gamma_{rb} T_{bt} + \epsilon_{hrbt} \tag{3.1}$$

where m_{hrbt} is the total miles driven for household h on road type r in toll band b at time t, τ_{rbt} is the toll rate expressed in 10 cents units, T_{bt} is the average traffic volume in the Seattle Area in toll band b, and δ_{rb} and η_h are the toll band-road type and household fixed effects. Road types can be either normal roads or highways. For this specification, β_t represents the average change in miles driven for \$0.10/mile toll in a given week (or day) t in the trial period. I divide the coefficient by the average miles driven during the baseline period to have an estimate of the

[14]The main mechanism of an hypothetical general equilibrium effect would be that due to the toll drivers might drive less in certain roads and hours and more in others. The change in traffic patters would create an additional incentive in changing driving behavior. Also, in the long run drivers would incorporate toll costs in their vehicle purchasing decisions.

percentage change in miles.[15] When the time unit is the day, I use fixed effects at the toll band-road type-day of the week level.

Some of the changes might occur at the extensive margins, that is, deciding not to drive in a certain toll band and in a certain type of road due to the introduction of the toll. I use d_{hrbw} as an indicator variable on whether the household drove a positive number of miles during a week in a certain toll band and type of road, obtaining the following linear probability model:

$$d_{hrbw} = \alpha + \delta_{rb} + \eta_h + \beta_t \, \tau_{rbt} + \gamma_{rb} \, T_{bt} + \epsilon_{hrbt} \tag{3.2}$$

where β_t in this case represents the average change in probability of driving in a certain toll band in a certain type of road for a \$0.10/mile toll.

In the specification by week, my time period includes 12 weeks before and 14 weeks after the beginning of the trial (up to the first week of October), in the specification by day, I include 30 days before and 30 days after the start of the trial.[16] The week when the trial started is not included in the sample due to the possibility of anticipation effects from participants. In the specification by day, I drop the two days before and after the beginning of the trial.

In addition to the main analysis, I estimate equations 3.1 and 3.2 disaggregated

[15]I do not use a log-linear specification because for some toll bands and type of roads there is a non-negligible share of observations with zero miles driven. However, results using a log-linear specification and dropping the observations equal to zero look very similar to the main specification, generally with a slightly lower semi-elasticity coefficient.

[16]Results look very similar when using different time windows.

by road type (highway and normal roads) and travel purpose (home-home, work-work, home-work, work-home). In a given toll-band, highway toll is twice as much as the toll for normal roads, so drivers have a stronger monetary incentive to reduce miles driven on highways. However, alternate routes might be impractical, too long or too time consuming, and it is possible that the response to the toll for highways is different than the one for normal roads.

Disaggregating the effect by travel purpose allows me to distinguish between two broader categories of tours: tours from and to known destinations (especially home-work and work-home tours), in most cases following a familiar route, and tours stopping at relatively new locations (especially home-home tours). On the one hand, tours from and to work locations may imply little flexibility in changing routes or times. On the other hand, familiarity with the route can help in adopting low-toll alternate routes.

My second goal is to study how the ability to respond to the toll is influenced by household characteristics. I estimate separately the "treatment" effect separately for each household for each week after the introduction of the toll:

$$m_{hrbt} = \alpha_{hrb} + \beta_{ht}\,\tau_{rbt} + \gamma_{hrb}\,T_{bt} + \epsilon_{hrbt}. \tag{3.3}$$

Next I transform the vector of treatment effects β_{ht} to percentage change in miles with respect to the average household's miles driven during the baseline period, and regress it on various household characteristics, including number of children in the household, number of vehicles, an indicator for home ownership, age of household head, household income in 2004 and initial endowment. I use a similar specification for measuring the effect on the extensive margin, with a dummy equal to one if a household drove a positive number of miles in a certain week-toll band-road type as dependent variable.

In the context of the trial, the mechanism through which initial endowment affects the response to the toll is not clear a priori. In fact, there are two possible reasons why household with high endowments might react differently to the toll than households with low endowments. The first is that, conditional on their driving behavior, their travel budget increased. The second is that households who drive more tend to receive a higher endowment, and households driving more might react differently to the toll.

To distinguish between the two effects, I exploit a specific characteristic of the relationship between endowment and driving behavior. As seen in figure 3.3.D, hypothetical toll costs during the baseline period are strongly correlated with endowment, but not perfectly. While participating households received an endowment based on past travel history, the endowment does not follow a strict mathematical formula and there is considerable variation in amounts given, even if miles driven are the same.

Besides the presence of a minimum endowment level, one of the reasons why such variation occurred is due to the different availability of baseline data across participants. The minimum amount of data is three months in most cases (April-June 2005), but several households provide up to six months of data. The different dates of installation of the OBUs and the presence of a technical problem with the OBUs in March 2005 that required a recall of all participating vehicles for a software update created a different range of baseline data used to estimate the initial endowment. Moreover, days very close to the beginning of the trial period arguably were not taken into account in the calculation.

In general, endowments were pretty generous relative to the actual driving behavior to prevent the risk that participants depleted their travel budget before the end of the experiment.[17] This was a more important aspect than making sure the endowments were exactly proportional to driving behavior.

I exploit this variation in endowment to see how the response to the toll changes with the endowment, conditional on driving behavior. To separate the variation in endowment not attributable to driving behavior I estimate the following equation:

$$\text{Endowment}_h = \alpha + \gamma_{rb} m_{hrb} + \xi_h \tag{3.4}$$

where the dependent variable is the household endowment, and m_{hrb} is the number of total miles driven for household h in toll band b on road type r. The residual $\hat{\xi}_h$ represents the variation in endowment not explained by driving behavior.

I use the residual $\hat{\xi}_h$ from the endowment equation to identify the effect of a change in endowment on toll response conditional on driving behavior, and the total hypothetical toll cost during the baseline period if the toll were in place to identify the effect of differences in driving behavior on toll response. The model is:

[17] In the data the ratio between the monthly endowment and the monthly hypothetical toll cost in the baseline period is always higher than 1, and very often higher than 2.

$$c_{ht} = \alpha + \gamma_t + \beta_1 \,\mathrm{Kids}_h + \beta_2 \,\mathrm{Veh}_h + \beta_3 \,\mathrm{Home}_h + \beta_4 \,\mathrm{Age}_h + \beta_5 \,\mathrm{Inc}_h + \beta_6 \,\mathrm{HToll}_h + \beta_7 \,\hat{\xi}_h + \epsilon_h \tag{3.5}$$

where c_{ht} is the change in miles driven for a 10 cents/mile toll estimated in equation 3.3 for household h during trial week t, divided by average miles driven by the household during the baseline period. On the right-hand side γ_t represents trial week fixed effects, Kids_h is a set of dummies for households with one, two, and three or more children, Veh_h is a set of dummies for households with two and three or more vehicles, Home_h is a dummy equal to one if the household lives in its own house, Age_h is the age of the household head, Inc_h is a set of dummies for quartiles of household total income in 2004, HToll_h is a set of dummies for quartiles of hypothetical toll, the amount the household would have paid if the toll was in place in the baseline period, and ξ_h is a set of dummies for quartiles of the residual from equation 3.4.

I estimate a similar equation for the effect on the extensive margin, using the same regressors as equation 3.5 and as dependent variable the average change in probability of driving in a toll band and type of road.

3.6 Results

Figure 3.8.A shows the results for equation (3.1) at the week level, using 12 weeks in the baseline period and 14 weeks from the trial period (until the beginning of October): in the first full week of its introduction, each 10 cents/mile toll decreases miles driven by 6.83%. Considering that the toll ranges between 0 cents and 50 cents per mile, the effect on miles driven can be substantial in the hours of the day and the road type in which the toll is high. For example, this would correspond to a decrease of 17.08% for a 25 cents toll per mile, and a decrease of 34.15% for a 50 cents toll per mile (the highest toll rate possible in the trial).

After the first week of the trial and until the last week of July, the effect on miles driven is still negative (about -2% in the second and third week) but is further decreasing with time and not longer significant at the 5% level.

Figure 3.8.B shows the results for equation 3.1 at the day level, using 30 days in the baseline period and 30 days in the trial period and dropping the two days before and after the beginning of the trial.[18] In particular, the reduction in miles continues after the 4th of July for all the rest of the week (with Thursday not statistically significant). In the following weeks there are only occasional statistically significant declines in miles driven during weekdays.

In the main equation, the dependent variable is the number of miles driven. That is because many observations are equal to zero due to the disaggregation

[18]Including the two days after the beginning of the trial shows that the two coefficients are negative and statistically significant.

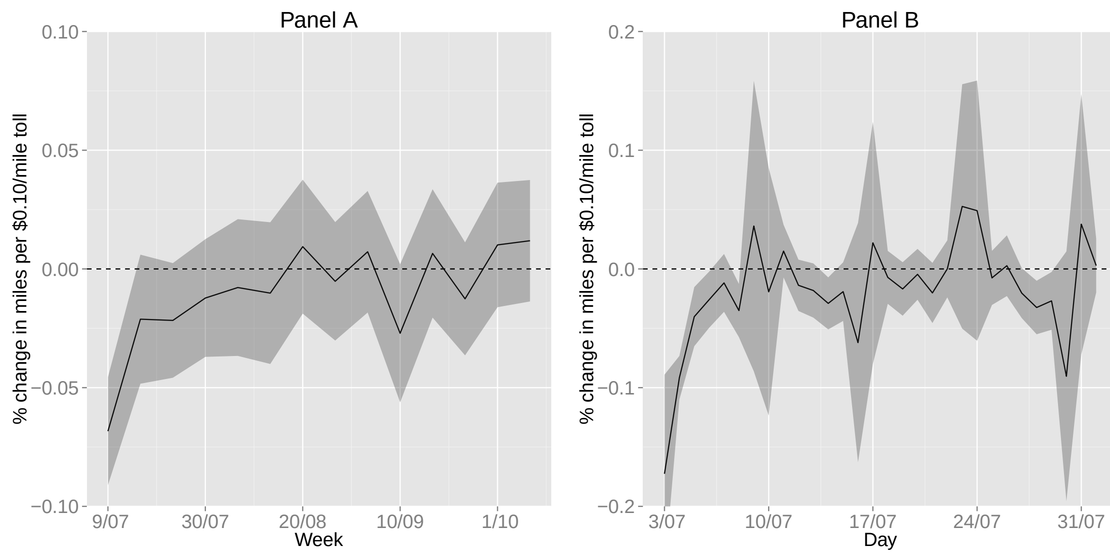

Figure 3.8: Effect in percentage terms on tolled miles driven for $0.10/miles toll. 95% confidence intervals from standard errors clustered at household level. **Panel A**: Results by week, 12 weeks in baseline period, 14 weeks in trial period. The week in which the trial started has been dropped. **Panel B**: Results by day, using 30 days in the baseline period and 30 days in the trial period. The two days before and after the start of the trial have been dropped. Tick marks represent Sundays.

of miles driven by time, day and road type. As a robustness check, I run also a specification with the log of miles driven as dependent variable, dropping the observations with zero mileage. While this specification does not take into account when drivers decide not to drive at all in a given time, day and type of road due to the introduction of the VMT tax, the results shown in table 3.9 are in line with those of the main specification.

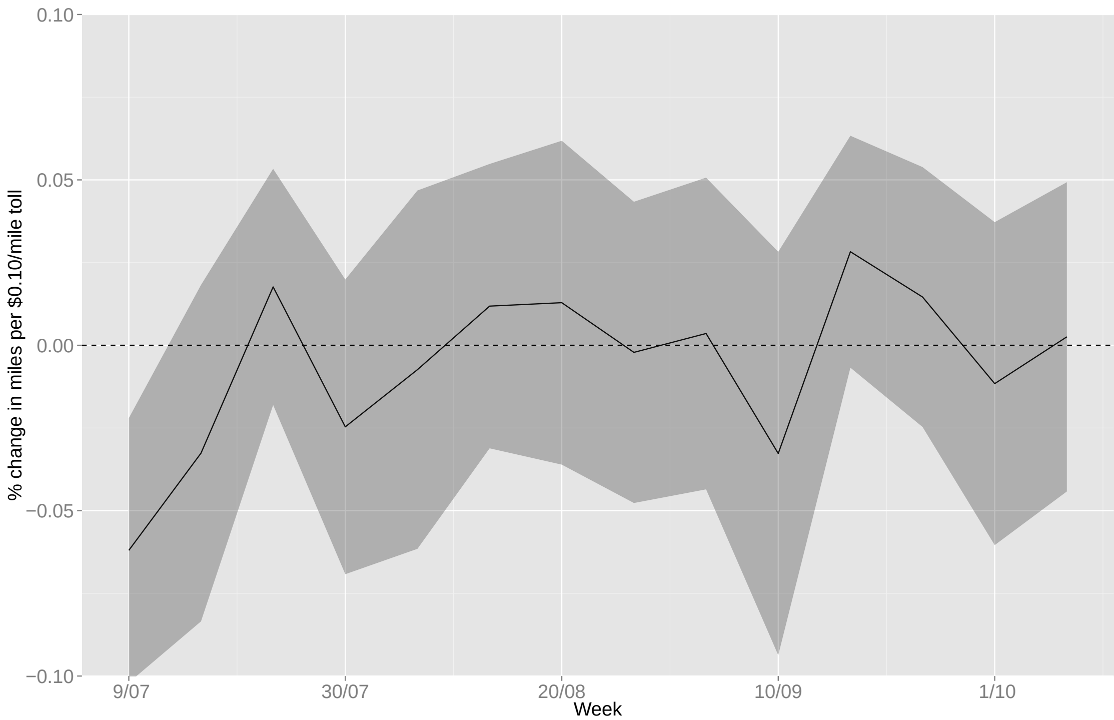

Figure 3.9: Effect in percentage terms on tolled miles driven for \$0.10/miles toll, log equation at the week level. 95% confidence intervals from standard errors clustered at household level.

Figure 3.10 looks at the change in probability per 10 cents/mile toll of driving in a certain toll band-road type for the month of July. At the week level, the effect looks pretty flat across time (a decrease of about 1%) and borderline significant at the 5% level. At the day level, the pattern is very similar to the result for miles driven, with a strong initial response during the first few days and a quick decline afterward.

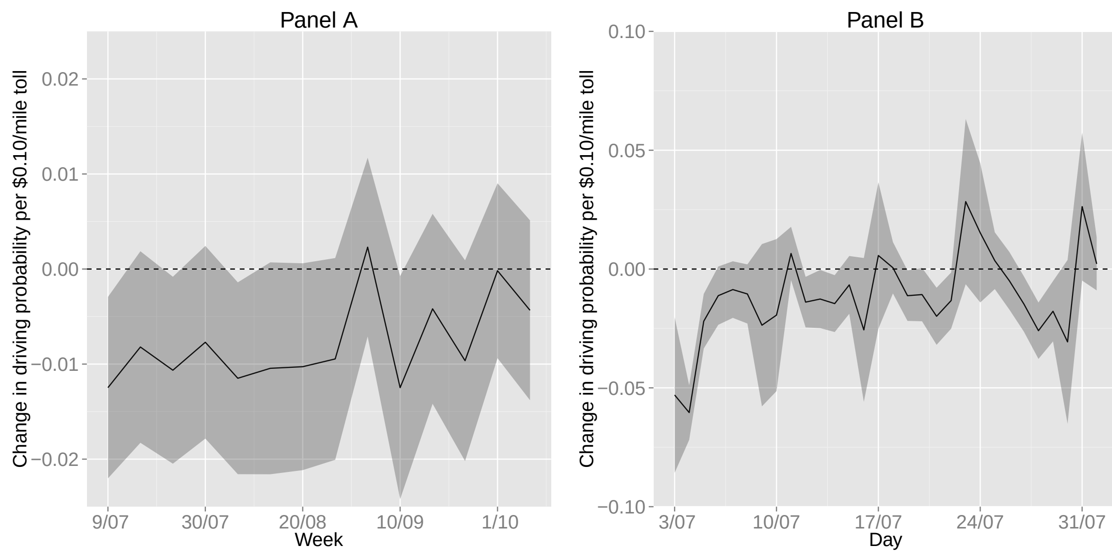

Figure 3.10: Effect in percentage terms on driving probability for \$0.10/miles toll. 95% confidence intervals from standard errors clustered at household level. **Panel A**: Results by week, 12 weeks in baseline period, 14 weeks in trial period. The week in which the trial started has been dropped. **Panel B**: Results by day, using 30 days in the baseline period and 30 days in the trial period. The two days before and after the start of the trial have been dropped. Tick marks represent Sundays.

To check that the results are not caused by drivers who have abnormal driving habits, like driving very frequently or very rarely, I trimmed the top and bottom households in terms of hypothetical toll in the baseline period - how much they would have paid if miles were taxes in the baseline period. Results are reported in figure 3.11 when dropping the top and bottom 1% (panel A), and the top and bottom 10% (panel B). Results show that even in case of drastic trimming of 20% of the households, results remain very robust.

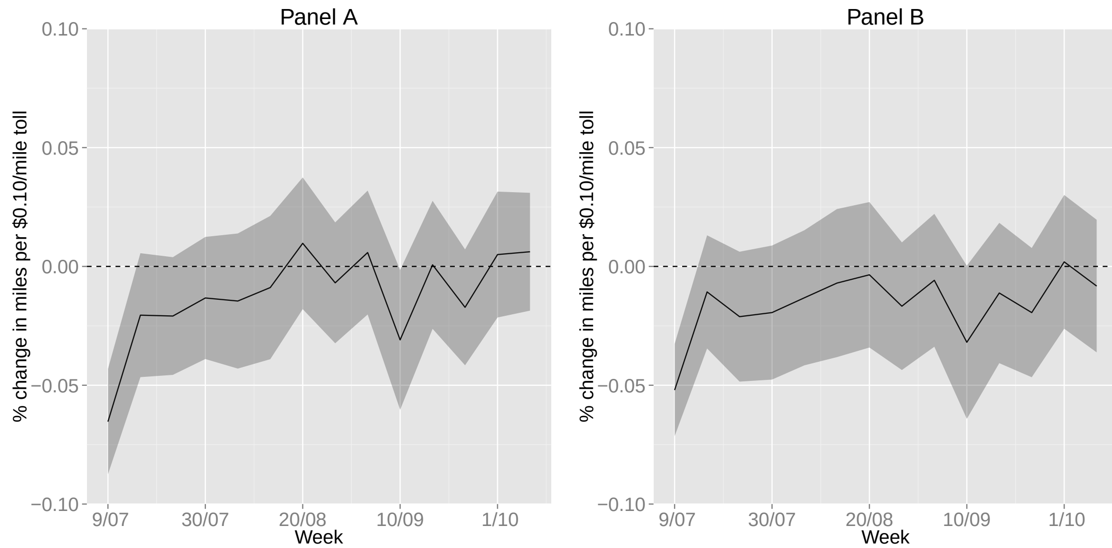

Figure 3.11: Effect in percentage terms on tolled miles driven for \$0.10/miles toll. 95% confidence intervals from standard errors clustered at household level. **Panel A**: Results by week, removing top and bottom 1% drivers in terms of hypothetical toll. **Panel B**: Results by week, removing top and bottom 10% drivers in terms of hypothetical toll.

A back of the envelope calculation on the effect of the VMT tax on carbon emissions using available data on vehicle fuel economy (average miles per gallon for urban driving) and official EPA CO_2 conversion rates shows that on average the carbon emissions from participating vehicles are 107.50 kg per week (based on average weekly mileage of 265.65 miles in the baseline period). The average change in CO_2 emissions per vehicle was then -7.34 kg in the first week and -2.28 kg in the second week. Using a social cost of carbon of \$25 per ton, as done by Parry and Small (2005), implies that the average dollar value per vehicle of this change was 18.3 cents in the first week and 5.6 cents in the second week. This calculation however does not take into account heterogeneity in response to the tax by different vehicles.

Finally, I compare the VMT tax elasticities derived from the results with the elasticities assumed or derived by the previous literature, generally derived from the milage response to gasoline price or other operating costs per mile. In the first week of the introduction of the VMT tax, the tax elasticity was 0.112, and between 0.044 and 0.011 afterwards. As a comparison, the tax elasticity assumed by Parry and Small (2005) is 0.22, 0.30 in Parry and Timilsina (2010) and between 1.46 and 0.84 in West (2004). With the caveat that those papers assume that the VMT tax would replace the fuel tax, my results show a lower tax elasticity than the previous literature.

The extremely short term effect of the toll has two possible explanations. First of all, the start of the trial period might have represented a strong signal to participants to monitor their travel behavior and take appropriate steps to reduce toll costs. However, the effect of such signal is extremely short-lived. There is some evidence in the literature on why that could be the case. In their analysis of changes in household energy consumption following a personalized home energy report with energy conservation tips, Allcott and Rogers (2014) use the concept of "cues" to explain the drop in energy consumption immediately after receiving a home energy report and a "backsliding" after few days. In this framework, notifications about energy consumption act as "cues" changing the marginal utility of consumption, with a limited effect in time. In the Traffic Choice experiment, the "cues" were perhaps not strong enough after the initial period to remind participants to change their travel behavior, and therefore the effect would be extremely limited in time. Finally, the fact that tolls were charged automatically without any necessary action from the drivers might have mitigated the visibility of the trial.[19]

Another explanation is that at the beginning of the trial period participants were not aware of the exact trade off in utility between alternative routes or driving times. During the initial trial period they might have modified their travel behavior, finding out that the savings in tolls were not enough to compensate for the inconvenience of taking a different route or driving at different times.

Figure 3.12 represents the result of equation 3.1 disaggregated by type of road at the week level. Panel A shows that the effect on miles driven on highways lasted at least three weeks after the beginning of the trial, with the largest impact, a reduction of 6.78% per 10 cents/mile toll, still occurring during the first week. On the other hand, the reduction of 7.05% on normal roads lasts only the first week (Panel B), and after that the effect is not significant, with a coefficient near zero or slightly positive. These results are consistent with the idea that at the beginning of trial drivers tried to avoid driving on highways as the toll was twice as much as the toll on normal roads.

The effect only on the extensive margins is less variable, with a 2%-1% decline both in highway and normal road miles, borderline significant at the 5% level (figure 3.13).

For the disaggregation by tour purpose I distinguish between home-home tours during weekend and during weekdays, and I consider work-work, home-work and work-home tours only during weekdays.[20] As seen in figure 3.5, these three last

[19]See for instance Finkelstein (2009) on the salience of automated versus manual toll roads.

[20]The number of miles driven from and to work locations during weekends is negligible.

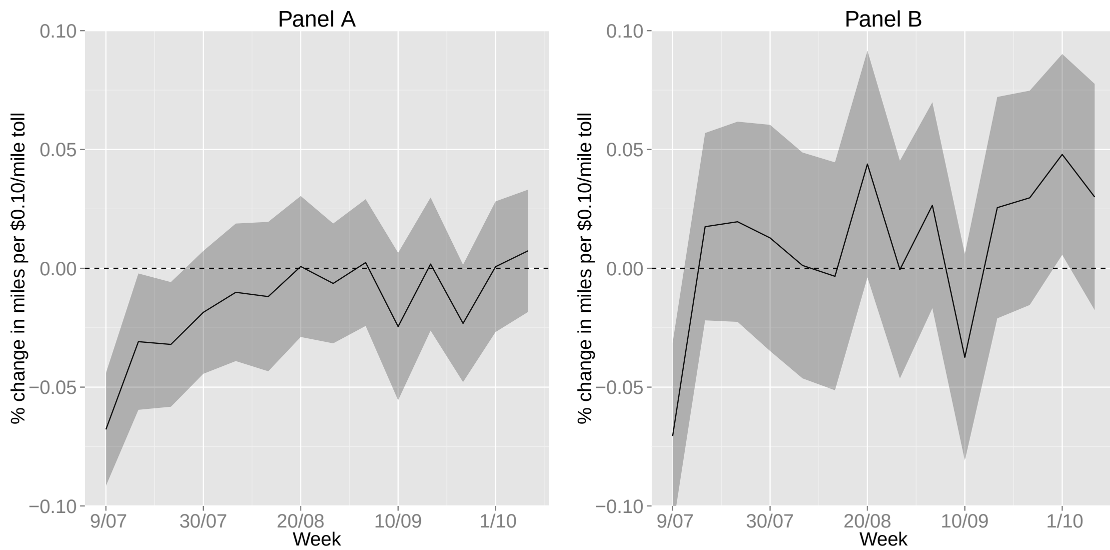

Figure 3.12: Effect in percentage terms on tolled miles for \$0.10/miles toll disaggregated by miles driven in highways and normal roads. 95% confidence intervals from standard errors clustered at household level. **Panel A**: Highways. **Panel B**: Normal roads.

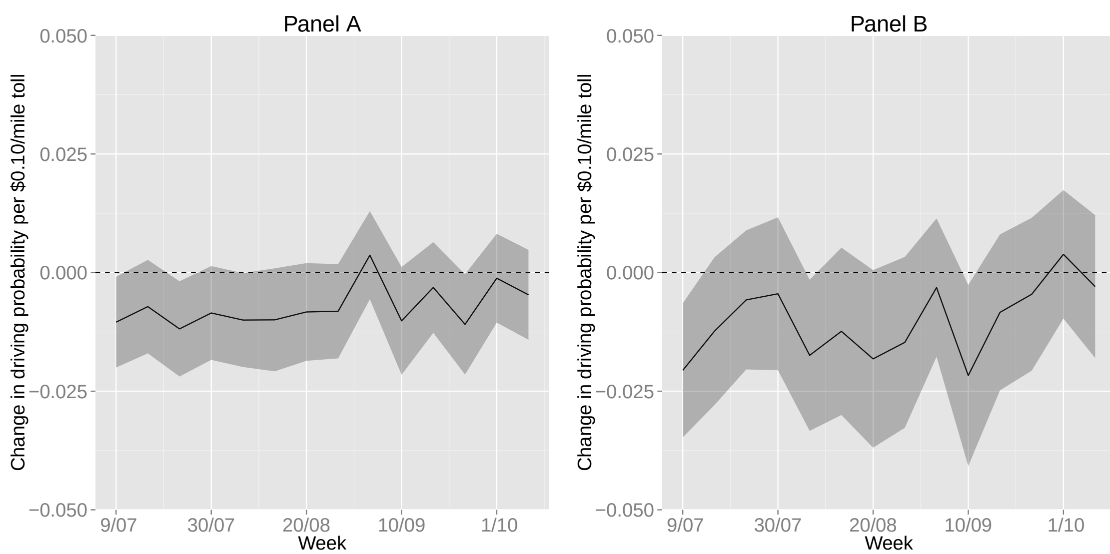

Figure 3.13: Effect on driving probability for \$0.10/miles toll disaggregated by miles driven in highways and normal roads. 95% confidence intervals from standard errors clustered at household level. **Panel A**: Highways. **Panel B**: Normal roads.

types of tours tend to be concentrated respectively in the early afternoon, in the morning and in the evening. That causes a drastic reduction in toll variation within tour purpose. To address this problem I also have a specification including all miles driven for any tour from or to work destinations during weekdays.

Figure 3.14 shows the effect of the toll on miles driven for various types of tours, limited to the first four weeks of trial: home to home tours during weekends (Panel A), home to home tours during weekdays (Panel B), work to work tours during weekdays (Panel C), home to work tours during weekdays (Panel D), work to home tours during weekdays (Panel E) and all the work-work, home-work, work-home tours during weekdays (Panel F).

Results show that in all cases the effects are not statistically significant after the first week. However, the largest percentage changes in miles driven occur for home-home tours in the weekends, and work-home tours in the weekdays. The fact that during weekdays there is no effect on home-home tours while there is one on tours from or to work destinations suggests that familiarity with the route allowed households to reduce their miles driven. Another explanation is most of work-home tours occur between 4pm and 7pm, when the toll is highest. However, also home-work tolls tend to occur during one of the highest toll rates, and yet I observe little reduction in miles driven during that period.

Figure 3.15 illustrates the distribution of coefficients from equation 3.3. While the distribution is centered slightly below zero, the variance is quite substantial, especially for the effect on miles driven. The most intuitive explanation for that is the presence of households driving relatively little, for which even small changes in

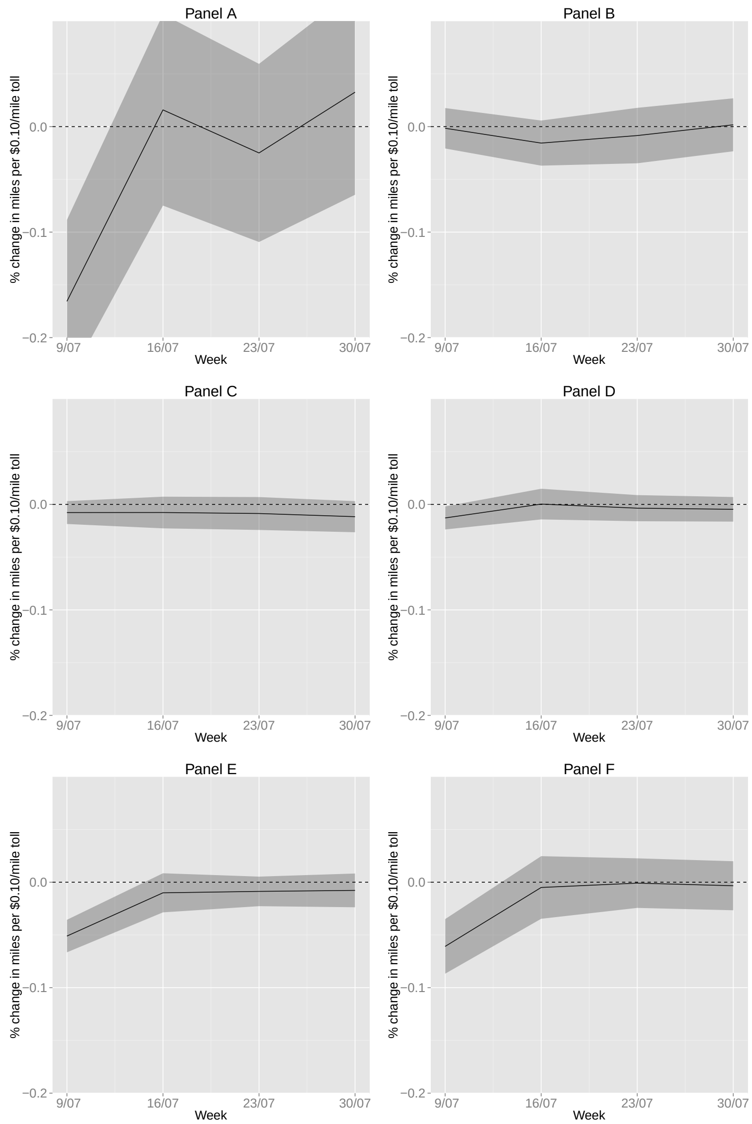

Figure 3.14: Percentage change in hourly miles driven on tolled roads for \$0.10/miles toll, disaggregated by tour purpose. 95% confidence intervals from standard errors clustered at household level. **Panel A**: Home-Home tours, Saturday-Sunday. **Panel B**: Home-Home tours, Monday-Friday. **Panel C**: Work-Work tours, Monday-Friday. **Panel D**: Home-Work tours, Monday-Friday. **Panel E**: Work-Home tours, Monday-Friday. **Panel E**: Work-Work, Home-Work and Work-Home tours, Monday-Friday.

miles driven are translated as high changes in percentage terms.

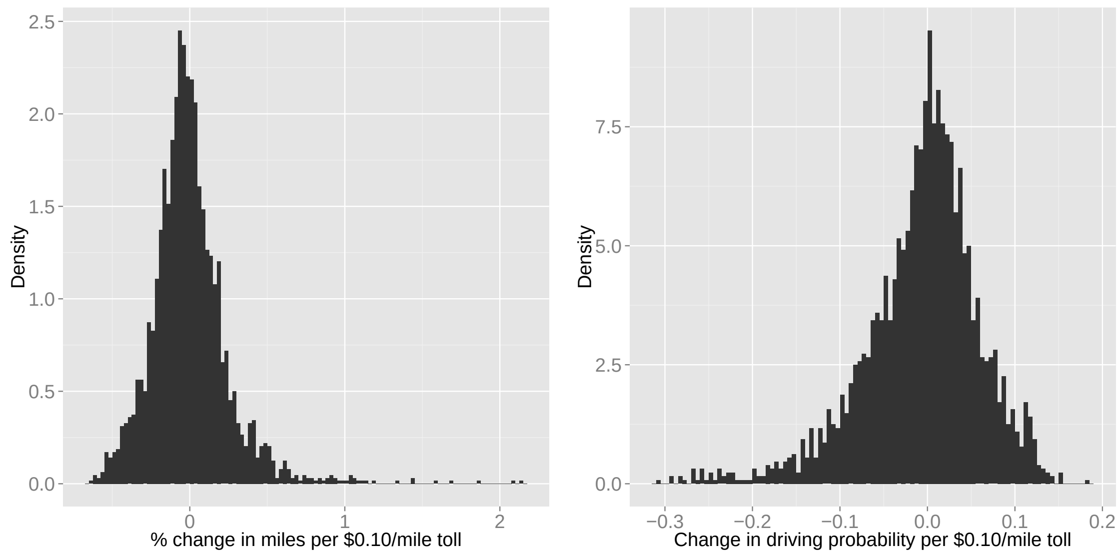

Figure 3.15: Distributions of coefficients from equation 3.3. Values at the top and bottom 1% for each week have been trimmed.

To check this hypothesis, I plot the relationship between average miles driven per toll band-road type in the baseline period (Figure 3.16). Households driving very little in tolled roads show much higher variance in the effects in the trial weeks.

Tables 3.6 and 3.7 show the results for equation 3.5 with percentage change in miles driven for 10 cents/mile toll as dependent variable. The first column shows the results including all the weeks in the sample. Due to the presence of some extreme values in the dependent variable, the second column shows the results after trimming the top 1% and the bottom 1% for each week. The third column exclude also all the households driving on average less than one mile per band-road type in the baseline. Table 3.7 shows the same results excluding the the first trial week from the sample.

Table 3.6: Heterogenous effects, miles driven

	(1)	(2)	(3)
N. children 1	0.556***	0.046***	0.049***
	(0.140)	(0.018)	(0.017)
N. children 2+	0.171***	0.010	0.011
	(0.045)	(0.013)	(0.013)
N. vehicles 2	−0.218***	−0.025*	−0.031**
	(0.050)	(0.015)	(0.015)
N. vehicles 3+	−0.331***	−0.014	−0.022
	(0.085)	(0.024)	(0.024)
Own home	−0.209***	−0.025	−0.031*
	(0.057)	(0.017)	(0.017)
Age HH head Q2	0.367***	0.137***	0.140***
	(0.076)	(0.036)	(0.036)
Age HH head Q3	0.082	−0.044	−0.046
	(0.077)	(0.051)	(0.050)
Age HH head Q4	0.081**	0.013	0.012
	(0.040)	(0.023)	(0.023)
Income Q2	−0.138***	−0.009	−0.013
	(0.040)	(0.016)	(0.016)
Income Q3	−0.024	0.025	0.025
	(0.034)	(0.015)	(0.015)
Income Q4	−0.004	0.084***	0.082***
	(0.039)	(0.019)	(0.018)
Hypothetical toll Q2	−0.135***	0.0002	−0.004
	(0.041)	(0.019)	(0.019)
Hypothetical toll Q3	−0.244***	−0.042**	−0.047***
	(0.052)	(0.016)	(0.016)
Hypothetical toll Q4	−0.192***	−0.079***	−0.081***
	(0.037)	(0.019)	(0.019)
Residual endowment Q2	0.009	−0.013	−0.012
	(0.021)	(0.016)	(0.016)
Residual endowment Q3	−0.135***	−0.039**	−0.038**
	(0.038)	(0.016)	(0.016)
Residual endowment Q4	0.300***	0.028	0.041
	(0.076)	(0.029)	(0.029)
Observations	2618	2562	2531

Results for equation 3.5. Column (1): All weeks. Column (2): Trimming bottom and top 1% for all trial weeks. Column (3): Trimming and removing households driving on average less than 1 mile per band-road type. White-Huber robust standard errors in parenthesis. P-values: *p<0.1; **p<0.05; ***p<0.01

Table 3.7: Heterogenous effects, miles driven, no first week

	(1)	(2)	(3)
N. children 1	0.590***	0.050***	0.053***
	(0.150)	(0.019)	(0.018)
N. children 2+	0.186***	0.014	0.016
	(0.048)	(0.014)	(0.014)
N. vehicles 2	−0.236***	−0.033**	−0.039**
	(0.054)	(0.016)	(0.016)
N. vehicles 3+	−0.356***	−0.020	−0.028
	(0.091)	(0.025)	(0.025)
Own home	−0.221***	−0.026	−0.032*
	(0.061)	(0.018)	(0.018)
Age HH head Q2	0.380***	0.137***	0.140***
	(0.080)	(0.036)	(0.036)
Age HH head Q3	0.097	−0.038	−0.040
	(0.083)	(0.054)	(0.054)
Age HH head Q4	0.082*	0.009	0.008
	(0.043)	(0.024)	(0.024)
Income Q2	−0.146***	−0.008	−0.012
	(0.043)	(0.018)	(0.018)
Income Q3	−0.022	0.029*	0.029*
	(0.036)	(0.016)	(0.016)
Income Q4	−0.003	0.090***	0.087***
	(0.041)	(0.019)	(0.019)
Hypothetical toll Q2	−0.150***	−0.007	−0.011
	(0.044)	(0.020)	(0.020)
Hypothetical toll Q3	−0.260***	−0.045***	−0.050***
	(0.056)	(0.018)	(0.017)
Hypothetical toll Q4	−0.198***	−0.080***	−0.082***
	(0.040)	(0.020)	(0.020)
Residual endowment Q2	0.007	−0.015	−0.015
	(0.022)	(0.017)	(0.017)
Residual endowment Q3	−0.141***	−0.038**	−0.037**
	(0.041)	(0.016)	(0.016)
Residual endowment Q4	0.327***	0.041	0.054*
	(0.082)	(0.030)	(0.030)
Observations	2431	2379	2351

Results for equation 3.5. Excluding first trial week. Column (1): All weeks. Column (2): Trimming bottom and top 1% for all trial weeks. Column (3): Trimming and removing households driving on average less than 1 mile per band-road type. White-Huber robust standard errors in parenthesis. P-values: *p<0.1; **p<0.05; ***p<0.01

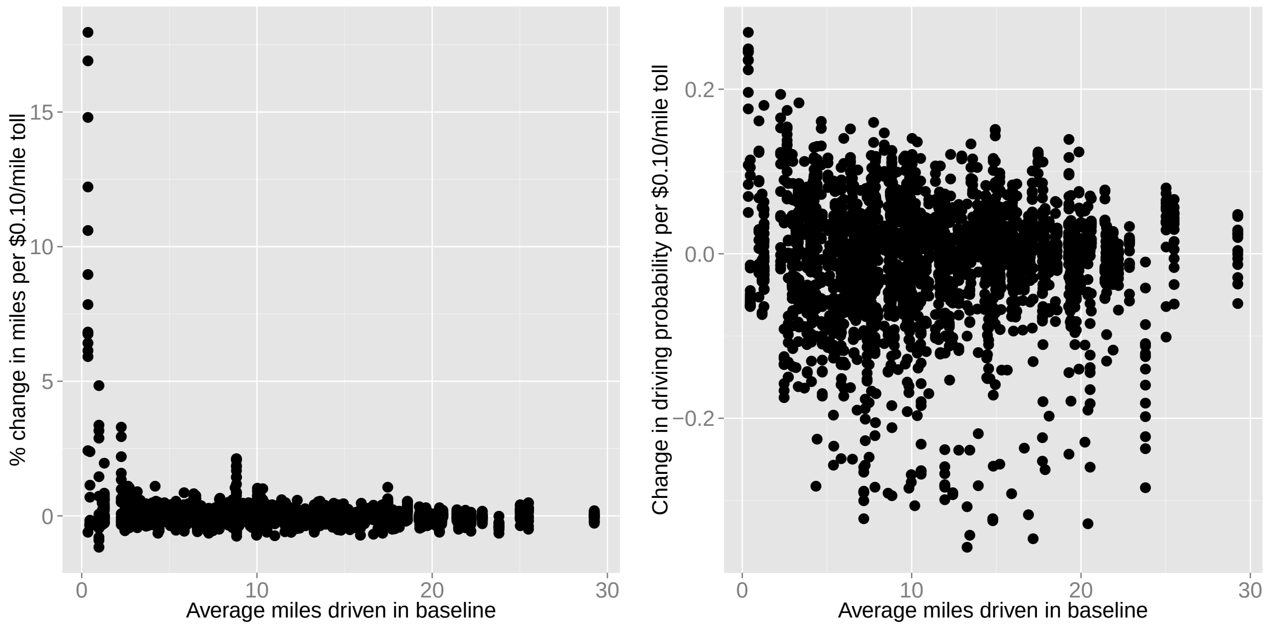

Figure 3.16: Relationship between coefficients from equation 3.3 and average miles driven per toll band-road type in the baseline period.

After removing the most extreme values (columns 2, 3 of tables 3.6 and 3.7) there is a considerable heterogeneity in the effect of the toll. Households with children react less to the toll than households with no children, probably due to the reduced degree of flexibility in their daily routine. Households with multiple vehicles have the possibility of carpooling in response to the toll: in fact they tend to reduce slightly their miles compared to households with only one vehicle, but the coefficient is significant only for households with two vehicles. Households in which the household head is between 38 and 47 years old (second quartile) tend to react less to the toll than younger and older household head.

Differences in household income, hypothetical toll in the baseline period and residual endowment have a strong effect in influencing the response to the toll: there is a difference of 8 percentage points in response to the toll between households on the bottom income quartile (below $45,000) and those on the top income quartile (over $100,000). This result is consistent with the findings by the previous literature showing a declining VMT tax elasticity with income.

Driving behavior itself has an important role in the response to the toll: households with higher hypothetical total toll costs during the baseline period tend to reduce the number of miles driven more than households with lower hypothetical toll costs. An explanation is that participants driving more or driving in more expensive hours and road types have more opportunities to change their driving behavior and saving more.

Finally, I observe a non-linear relationship between the reduction in miles driven and residual household endowment (the change in endowment not correlated with past driving behavior). Participants decrease the number of miles driven when the endowment increases, up to a certain point. After that, they increase the number of miles again. In Appendix B.2 I illustrate some summary statistics about the residual household endowment.

Similar considerations can be made for the heterogeneity in the probability of driving in a certain toll band-road type bundle. Tables 3.8 and 3.9 show analogous results for the effect on the extensive margins only. Also in this case, the number of children, high household income and high residual endowment reduce the response to the toll. The rest of the variables have relatively small and usually insignificant effects.

An important observation following the considerable heterogeneity in the effect of VMT tax on driving behavior is that while on average the effect of the VMT tax disappears after few days, a subset of participants might have indeed continued responding to the trial even after the first week. This consideration is coherent with past studies finding variation in VMT response between income groups, urban and rural drivers and households with and without children (West, 2004, 2005, Weatherford, 2011).

Table 3.8: Heterogenous effects, probability of driving

	(1)	(2)	(3)
N. children 1	0.028***	0.014***	0.013***
	(0.005)	(0.004)	(0.004)
N. children 2+	0.012***	0.008**	0.008**
	(0.004)	(0.004)	(0.004)
N. vehicles 2	−0.003	−0.002	−0.001
	(0.005)	(0.004)	(0.005)
N. vehicles 3+	0.011	0.015**	0.017***
	(0.007)	(0.006)	(0.006)
Own home	−0.009**	−0.009**	−0.009**
	(0.004)	(0.004)	(0.004)
Age HH head Q2	0.025**	0.019*	0.018*
	(0.011)	(0.011)	(0.011)
Age HH head Q3	−0.020*	−0.020*	−0.020*
	(0.012)	(0.012)	(0.012)
Age HH head Q4	0.008	0.010	0.009
	(0.007)	(0.006)	(0.006)
Income Q2	0.001	0.002	0.002
	(0.005)	(0.004)	(0.004)
Income Q3	0.012**	0.012***	0.011***
	(0.005)	(0.004)	(0.004)
Income Q4	0.017***	0.019***	0.019***
	(0.005)	(0.005)	(0.005)
Hypothetical toll Q2	−0.012***	−0.005	−0.004
	(0.005)	(0.004)	(0.004)
Hypothetical toll Q3	−0.013**	−0.001	−0.0001
	(0.005)	(0.005)	(0.005)
Hypothetical toll Q4	−0.019***	−0.008	−0.008
	(0.006)	(0.005)	(0.005)
Residual endowment Q2	−0.004	−0.003	−0.002
	(0.004)	(0.004)	(0.004)
Residual endowment Q3	−0.009*	−0.007	−0.007
	(0.005)	(0.005)	(0.005)
Residual endowment Q4	0.006	−0.002	−0.003
	(0.009)	(0.008)	(0.008)
Observations	2618	2562	2531

Results for equation 3.5. Column (1): All sample. Column (2): Trimming bottom and top 1% for all trial weeks. Column (3): Trimming and removing households driving on average less than 1 mile per band-road type. White-Huber robust standard errors in parenthesis. P-values: *p<0.1; **p<0.05; ***p<0.01

Table 3.9: Heterogenous effects, probability of driving, no first week

	(1)	(2)	(3)
N. children 1	0.029***	0.015***	0.014***
	(0.005)	(0.004)	(0.004)
N. children 2+	0.013***	0.008**	0.008**
	(0.004)	(0.004)	(0.004)
N. vehicles 2	−0.006	−0.004	−0.003
	(0.005)	(0.005)	(0.005)
N. vehicles 3+	0.006	0.013*	0.014**
	(0.008)	(0.007)	(0.007)
Own home	−0.008*	−0.008*	−0.008**
	(0.005)	(0.004)	(0.004)
Age HH head Q2	0.023**	0.017*	0.016
	(0.010)	(0.010)	(0.010)
Age HH head Q3	−0.017	−0.019	−0.019
	(0.013)	(0.012)	(0.012)
Age HH head Q4	0.008	0.010	0.009
	(0.007)	(0.007)	(0.007)
Income Q2	0.002	0.003	0.003
	(0.005)	(0.004)	(0.004)
Income Q3	0.013***	0.013***	0.012***
	(0.005)	(0.004)	(0.004)
Income Q4	0.018***	0.020***	0.020***
	(0.005)	(0.005)	(0.005)
Hypothetical toll Q2	−0.015***	−0.007	−0.006
	(0.005)	(0.005)	(0.005)
Hypothetical toll Q3	−0.014**	−0.002	−0.001
	(0.005)	(0.005)	(0.005)
Hypothetical toll Q4	−0.017***	−0.006	−0.006
	(0.006)	(0.005)	(0.005)
Residual endowment Q2	−0.004	−0.003	−0.003
	(0.005)	(0.004)	(0.004)
Residual endowment Q3	−0.007	−0.005	−0.005
	(0.005)	(0.005)	(0.005)
Residual endowment Q4	0.010	0.001	−0.0005
	(0.009)	(0.008)	(0.008)
Observations	2431	2379	2351

Results for equation 3.5. Excluding first trial week. Column (1): All sample. Column (2): Trimming bottom and top 1% for all trial weeks. Column (3): Trimming and removing households driving on average less than 1 mile per band-road type. White-Huber robust standard errors in parenthesis. P-values: *p<0.1; **p<0.05; ***p<0.01

3.7 Conclusions

In this paper I examine the impact on driving behavior of a VMT tax with a differentiated rate by hour, day and type of road, introduced as an experiment in the Puget Sound Region, WA, in 2005. Participating households were monitored during a baseline period, then received a monetary endowment in a dedicated account which was used to pay the toll during the trial period. Participants received the money left in the account at the end of the experiment. There were a total of 276 participating households (210 in the sample used for the analysis), which rules out any large scale effect from the trial.

While there exists a large and well established literature simulating the impact of a VMT tax on drivers' behavior and evaluating the welfare and distributional consequences of switching from a fuel tax to a VMT tax, this is the first study looking at the actual impact of a toll on miles driven, the evolution of the response across time, different types of road and different travel purposes, and heterogeneous effects.

I use the variation in toll rates across different toll bands and types of road to identify the average effect of the toll on miles driven and the probability of driving in a certain toll band in a certain road type. I control for seasonal effects by using information from traffic monitoring stations in the area of the trial, which, I argue, should not be influenced by the toll.

Results shows an extremely short-lived average effect on miles driven of about -6.83% for a 10 cents per mile toll. After the first trial week the effect disappears except for miles driven in highway roads, which are charged more than normal roads. Results seem to be driven by weekend travels, and weekday travels from work to home locations, possibly due to higher flexibility in choosing alternative, cheaper routes. A possible explanation of the short term effect is either that participants stopped paying attention to the toll, or that they tried alternative routes and driving times during the first trial week and realized that toll savings were not enough to compensate for the discomfort of changing their travel plans.

The difference in the effect immediately after the start of the trial and in later periods suggests that caution is needed when evaluating the impact of certain policies using a regression discontinuity strategy with time as a running variable (also called event study): while this methodology provides credible identification of the immediate impact in case of an unexpected policy, the focus on a very short time window implies that one may not be able to estimate possible changes in treatment effects even few weeks after the start of the treatment. In this paper, predictions on long and medium-term effects of a policy based on immediately after its introduction would be misleading.

Finally, I find considerable heterogeneity in drivers' response to the toll. In particular, household income and number of children are less responsive, consistent with the notion that those households have high value of time. Households with high total hypothetical toll costs during the baseline period decrease miles driven. Some of these results are consistent with previous findings from simulations of distributional effects of VMT taxes. These results hold even after the first week of introduction of the toll, and such heterogeneous response might hold even in the medium and long run.

The short duration of the average effect on miles driven suggests that VMT taxes might be a useful instrument for generating tax revenue. At the same time, there is some uncertainty on the effectiveness of VMT taxes in their current form to mitigate externalities from driving. In the case of non-revenue neutral environmental taxes, there is a trade-off between reducing externalities and generating revenue. As a result, the effectiveness of a VMT tax depends crucially on the reason why it was introduced.

There are a number of limitations to my study, which are due to the nature of the Traffic Choices Experiment. I am not able to assess what would be the effect of a full scale adoption of a VMT tax at the city or state level. The findings of this paper should be considered only as partial equilibrium effects about the direct response of participating drivers to the toll.

Another limitation is that participation to the Traffic Choice program was voluntary and subject to meeting certain recruitment criteria. In particular, due to such requirements participants in the experiment are frequent drivers and more likely to drive to work during peak hours than a completely random sample. Relatively little can be said about very occasional drivers, but it must be added that frequent drivers contribute more to congestion and other vehicle externalities, and they would likely be the most relevant source of revenue for a VMT tax.

Appendix A: Appendix Chapter 1

A.1 Fuel economy and mileage estimates

The NTS used for the estimation of the mileage equation does not have information on fuel economy, only on CO_2 emission rates. We use instead an European Environmental Agency database which contains information on passenger cars from 2000 to 2010 to estimate the relationship between fuel economy and emission rates.

We regress fuel economy with CO_2 emissions rates separately for diesel and gasoline vehicles, without constant terms. Table A.1 shows the virtually perfect correlation between fuel economy and emissions rate. We use the regression coefficients to estimate the fuel economy from the NTS dataset.

Once obtained the fuel economy for the vehicles included in the NTS dataset, we used the appropriate 12-month moving average of fuel costs (based on the date of the survey and the geographic location) to calculate the fuel cost in British Pounds per 100 km.[1]

We then use the NTS data to regress the log yearly mileage (in km) over

[1] The geographic subdivisions used are the Government Official Regions: North East, North West and Merseyside, Yorkshire and Humberside, East Midlands, West Midlands, Eastern, Greater London, South East, South West, Wales, Scotland. Northern Irland is not included in the database.

Table A.1: Results from emission rates regression

Variables	Gasoline g CO_2/km	Diesel g CO_2/km
Fuel economy	23.77***	26.49***
(l/ 100 km)	(0.0300)	(0.0075)
R-squared	0.9995	0.9986
Observations	29,080	17,364

Dependent variable is vehicle CO_2 emission rates in grams per km. Regressor is the fuel economy in liters per 100 km. Results separated by fuel type. Data from the European Environmental Agency.

vehicle age in years, fuel cost in real 2005 British Pounds per 100 km, and engine size in cc. Omitting the time and the vehicle subscript for simplicity, we obtain:

$$\ln \text{Miles} = \alpha + \beta \text{AGE} + \gamma \text{FUELCOST} + \delta \text{ENGINESIZE} + \epsilon \qquad \text{(A.1)}$$

We run separate regressions for diesel and gasoline vehicles. Because the NTS is not a panel dataset and each vehicle is observed only once, we exploit the variation in age and mileage of the different vehicles surveyed.

The results of the mileage regression are shown in table A.2. As expected, mileage decreases with fuel cost and vehicle age, and increases with engine size. We use the results from the mileage regression to estimate the km driven by each vehicle in the Polk dataset during the lifetime of the car. According to the Society of Motor Manufacturers and Traders, the average scrappage age in the UK in a given year is between 13 and 14.5 years. We do not have disaggregated information by vehicle characteristics, so we assume that all vehicles in our dataset have a lifetime of 14 years.[2]

Table A.2: Results from mileage regression

Variables	Gasoline Log km/year	Diesel Log km/year
Age	-0.0231^{***}	-0.0253^{***}
	(0.0200)	(0.0031)
Fuel cost/100 km	-0.0199^{***}	-0.0390^{***}
	(0.0061)	(0.0078)
Engine size (cc)	0.0003^{***}	0.0002^{***}
	(0.0000)	(0.0000)
Observations	21,776	7,750

Results from equation A.1. Dependent variable is the log of km driven per year. Results separated by fuel type. Data from the UK National Travel Survey.

We use the following formula to calculate the lifetime mileage M:

$$M = \int_0^{14} \alpha + \beta \text{AGE} + \gamma \text{FUELCOST} + \delta \text{ENGINESIZE} \, d\text{AGE} \tag{A.2}$$

which becomes

$$\begin{aligned} M =& \text{Exp}\left(\frac{\alpha + 14\beta + \gamma \text{FUELCOST} + \delta \text{ENGINESIZE}}{\beta}\right) \\ &- \text{Exp}\left(\frac{\alpha + \gamma \text{FUELCOST} + \delta \text{ENGINESIZE}}{\beta}\right) \end{aligned} \tag{A.3}$$

Figure A.1 shows the distribution of total mileage under the VED and the carbon tax. The double hump shape is due to using different regression coefficient to predict mileage for diesel and gasoline vehicles. Under the carbon tax, the whole distribution shifts towards left, showing a decrease in km driven.

[2]Source: http://www.smmt.co.uk/2014-sustainability/environmental-performance/end-life-vehicles/

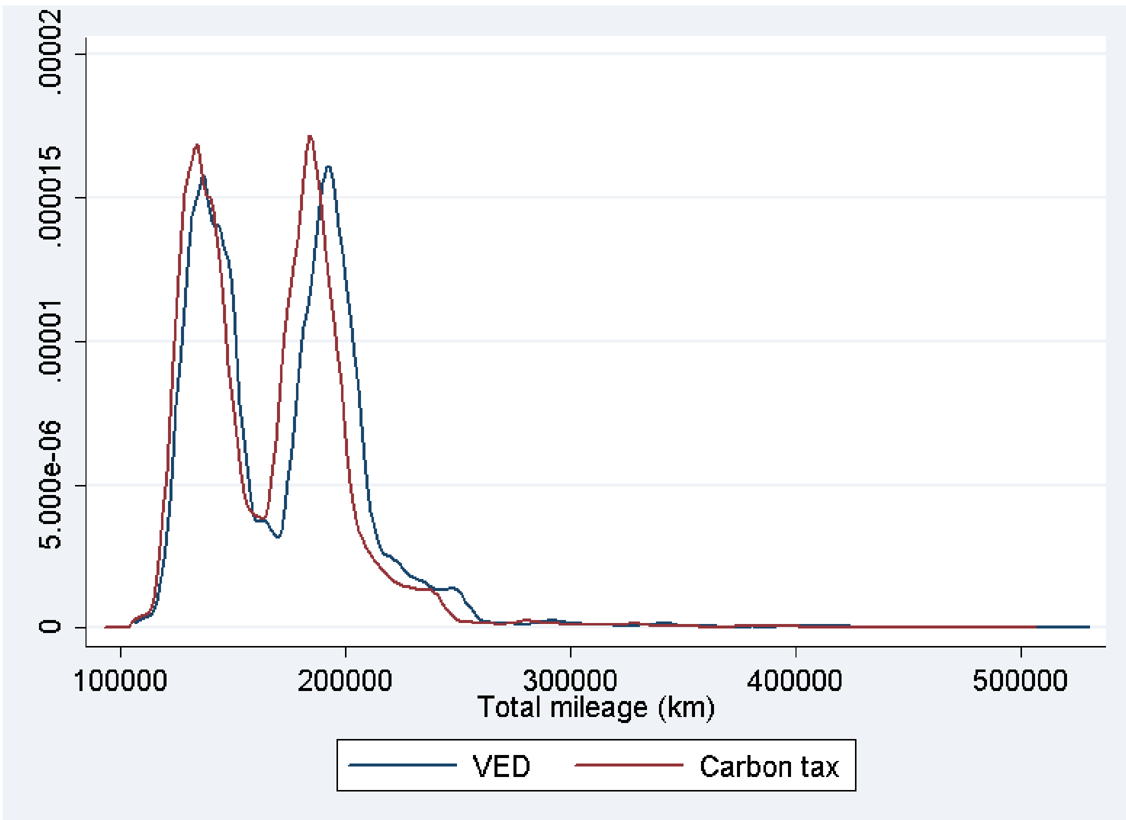

Figure A.1: Distribution of mileage driven during vehicle lifetime under VED and carbon tax. Results not weighted by vehicle sales.

A.2 News articles and web searches on the VED

Changes in VED occurred regularly at the beginning of each budget period and from 2008 the government disclosed future changes in the VED in their budget documents. If people are informed in advance about potential changes, they can react accordingly. For instance, they can buy a more polluting vehicle before the new rates are introduced. We want to know how much the general public is aware of changes in the VED. To do so, we rely on two sets of data. First, the number of newspaper articles about the VED, and second, an index of interest over time through Google searches. We want to see if the peaks of articles and search interest occur before the changes are implemented. That would mean that the general public is aware that changes in the VED are due shortly.

The newspaper articles data come from LexisNexis, and includes 156 publications in the UK. Among those outlets, we searched for articles including the words VED or vehicle excise duty. We considered articles from 2006 to 2010, as before that period the exact date of the article is not always specified.

Figure A.2 shows the distribution of the news articles across time. Peaks in VED newspaper coverage occur generally right after the VED changesin March 2006, March 2007, March 2008 and April 2010. In May 2009, when the changes were modest, we do not observe peaks. We do not observe peaks in the month before changes in rates took place. The peaks in news coverage between May 2008 and July 2008 are caused by protests against scheduled increases in the VED which hit existing vehicle owners, instead of just new vehicles. Eventually, these planned increases were scrapped in November 2008, which generated another news peak.

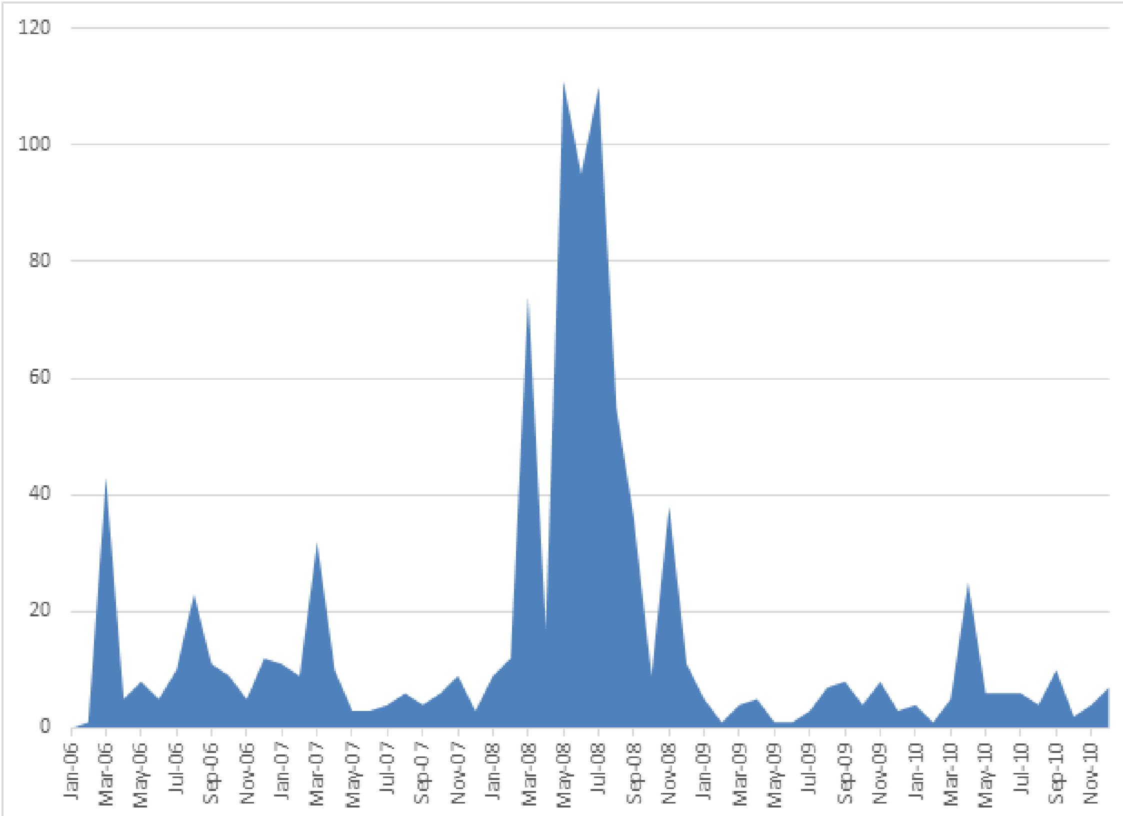

Figure A.2: Number of newspaper article about the VED per month between 2006 and 2010. Source: LexisNexis.

Importantly, we do not observe in the headlines and in the article contents information or speculation about future rates, with the partial exception of 2008. Similarly, we do not observe articles warning the readers about imminent changes in the VED. The majority of the articles in our dataset inform the general public about current rates. We then look at the Google Trends index for web searches about the VED. Google provides an index for all the searches related to the VED (topic), regardless the exact words searched. Figure A.3 shows a measure of relative interest to the VED between 2005 and 2010. Earlier data, especially 2005, is less reliable but we include it for completeness. The graph shows that changes in interest occurred the months in which changes in VED rates (2006 and 2007) occurred and in summer 2008 as seen for newspaper coverage. In general, the measure of interest is flat across time.

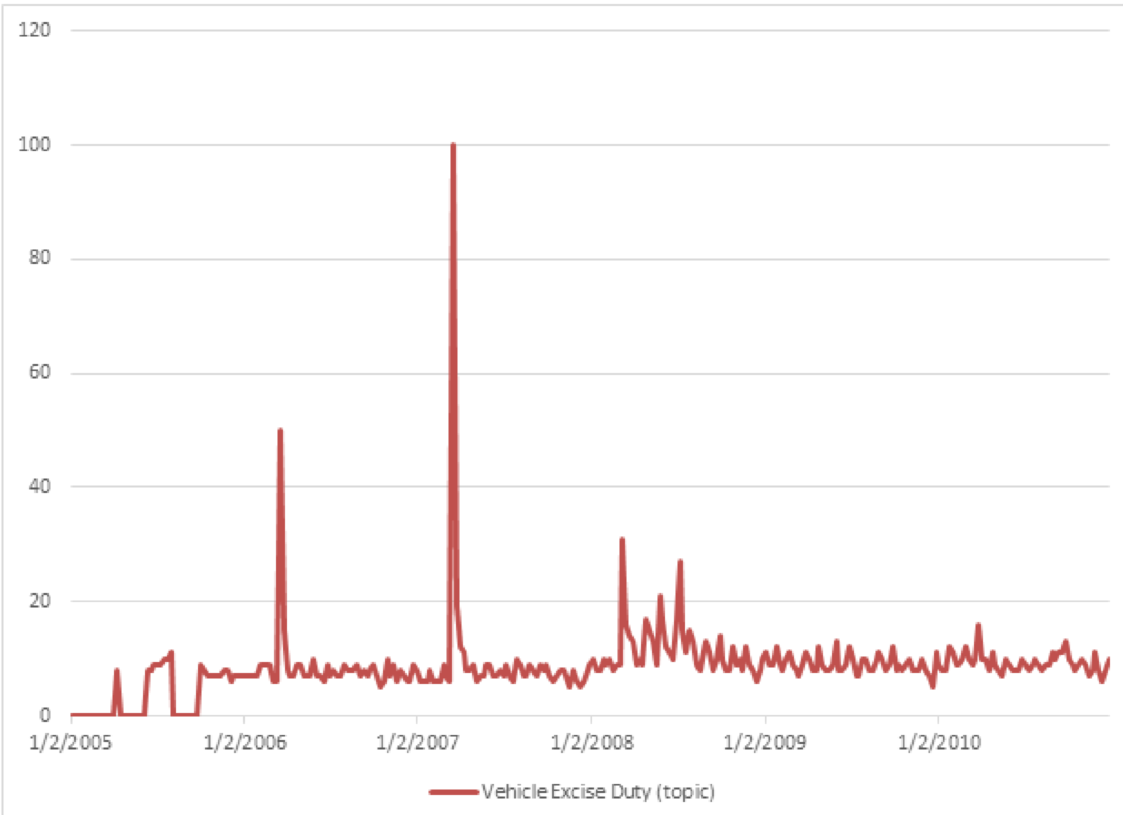

Figure A.3: Google Trends index about searches on the VED as a topic. Source: Google.

A.3 Calculation of the proportional tax and carbon tax rates

The nominal tax rate for the tax proportional to the carbon emission rates and for the carbon tax are calculated so that the total real revenue from the vehicles sold and registered between April 2005 and October 2010 is equal to the VEDs. The revenue takes into account the whole vehicle lifetime, assumed to be 14 years (SMMT data), so the VED revenue from a single model-trim-variant i during its lifetime is given by:

$$\text{Revenue}_i = \sum_{0}^{13} \frac{\text{VED}_i}{\frac{\text{HICP}_i}{100}\,(2.8t/100)} \tag{A.4}$$

Where VED_0 is the nominal VED rate for that vehicle at the time of purchase, $\frac{\text{HICP}_0}{100}$ is the Harmonised Index of Consumer Prices at the time of purchase, and t is the age of the vehicle.

We make some assumptions on the VED rates and consumer price index following the first registration year: 1) Each year the HICP increases by 2.8 points, which is the average yearly increase between 2005 and 2015. 2) We assume that the VED rates do not change over time, with the exception of the period April 2010 October 2010, where at the moment of the registration the second year rate was set different than the first year. The total revenue from the VED is simply the sum of the revenues from the predicted number of new registrations. The total revenue from proportional tax is given by

$$\text{Revenue} = \sum_{i}^{N} \tau_p \delta_i \text{CO2}_i \text{REG}_i \tag{A.5}$$

where

$$\delta_i = \sum_{t=0}^{13} \frac{1}{\frac{\text{HICP}_i}{100}\left[1 + (2.8t/100)\right]} \tag{A.6}$$

while τ_p is the proportional nominal tax rate in Pounds per g CO_2/km, CO2_i is the carbon emission rate in grams per km and REG_i is the number of predicted registrations from model-trim-variant i. In a similar fashion, the total revenue from the carbon tax is given by

$$\text{Revenue} = \sum_{i}^{N} \sum_{t=0}^{13} \frac{\tau_C \delta_{it} \text{CO2}_i M_{it} \text{REG}_i}{1,000,000} \tag{A.7}$$

where τ_C is the nominal carbon tax rate in Pounds per ton CO_2 and M_{it} is the mileage in km for model-trim-variant i at vehicle age t. The mileage for each year of the vehicle is predicted with the methodology explained in Section A.1 of this appendix. To calculate the tax rate to use for the proportional tax and the carbon tax, we use a simple algorithm with the following steps: 1) Select a tax rate from a range of possible rates for the proportional or the carbon tax 2) Predict lifetime mileage using the model explained in Section A.1. In the case of the carbon tax, the equivalent cost per 100 km is added to the driving costs. 3) Predict registrations for each model-trim-variant and each period (normalized by model-period) under the VED or one of the alternative policies using the main model 4) Calculate the total revenue from the VED and from the proportional or the carbon tax 5) Keep the tax rate only if the absolute value of the difference between the two revenues is within the 0.5% of the revenue from the VED. The revenue from the VED is estimated to be roughly 14.2 billion Pounds. Figures A.4 and A.5 shows the estimated total revenue for a range of values of the proportional tax and the carbon tax rates, respectively. The equal revenue tax rates are 0.83 Pounds per g CO_2/km for the proportional tax and 72 Pounds per ton CO_2 for the carbon tax.

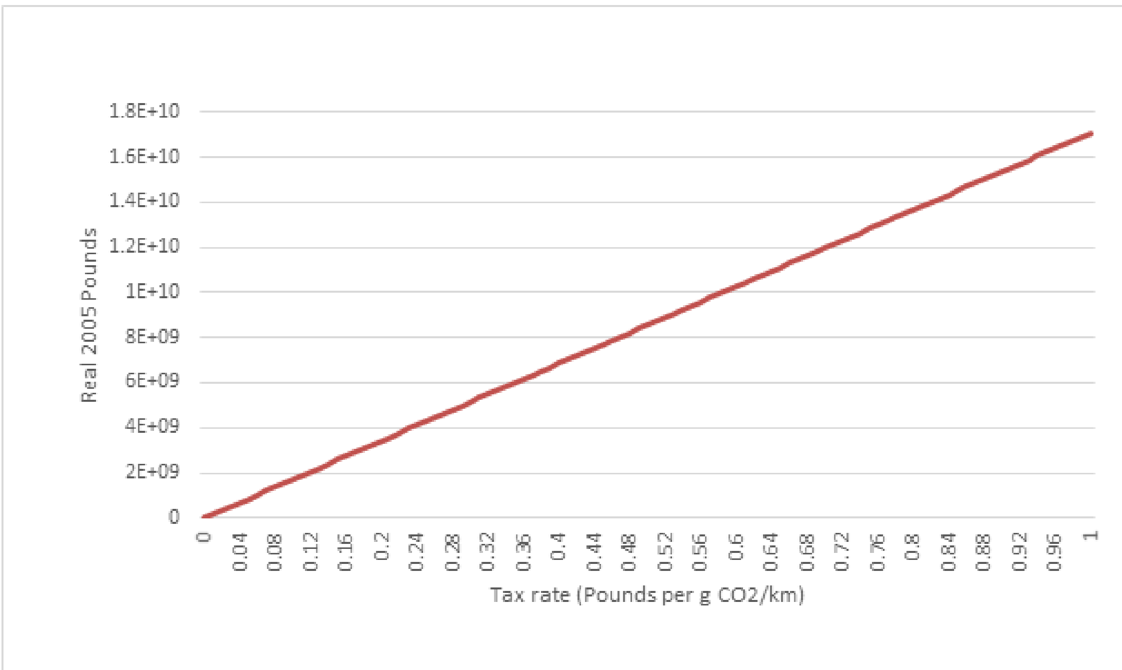

Figure A.4: Lifetime predicted aggregate real revenue with a registration tax proportional to CO_2 emission rates.

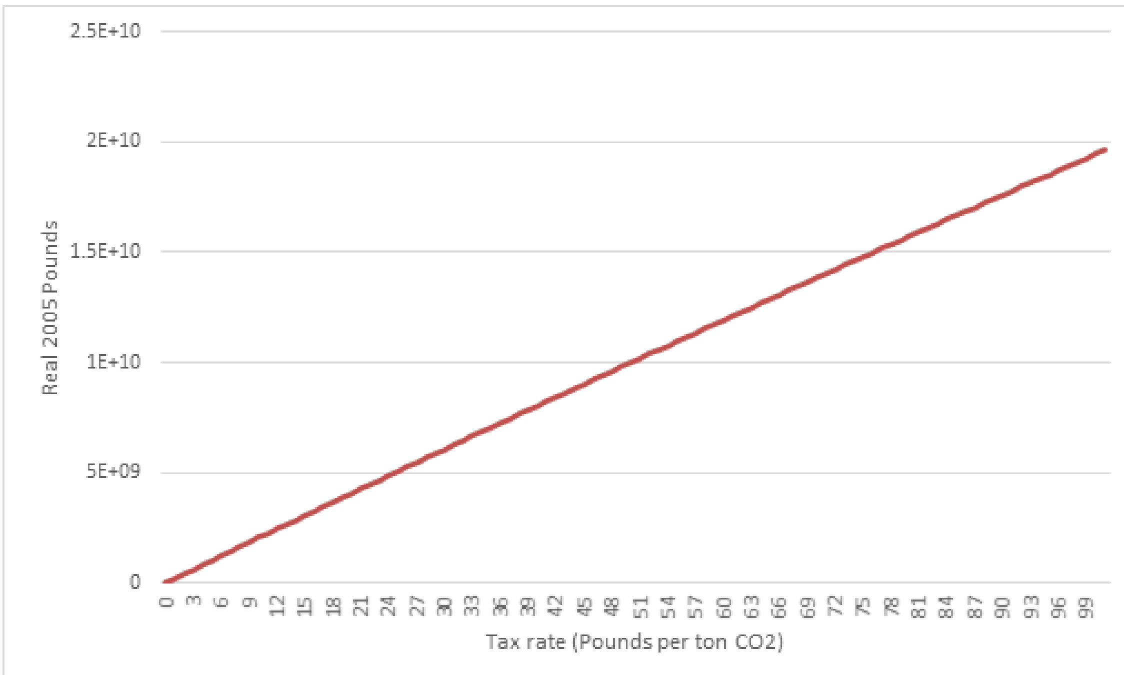

Figure A.5: Lifetime predicted aggregate real revenue with a carbon tax.

Appendix B: Appendix Chapter 3

B.1 Data cleaning and further data information

The original data in the Traffic Choice database contains information on 300,129 tours and 776,272 trips. I consider a subset of those tours and trips: those which belong to the official period of the experiment (From April 2005 to February 2006) and excluding households who dropped out from the experiment, households which received an adjusted final compensation, households which were not part of the original list of participants, households which changed one of their vehicles during the experiment and households with a baseline period length lower than three months (April-June 2005), for a total of 66 households excluded.[1]

Around 5.97% of the trips in the dataset have zero duration and zero miles. That occurs because a trip is recoded whenever the engine is turned on, regardless of the miles driven. I dropped these trips from my sample and I recalculated

[1]Only 6 households dropped out from the trial before its end and only one changed vehicle during the experiment. The vast majority of dropped households were households not included in the initial recruitment phase and households who received a compensation not corresponding to the difference between endowment and toll costs. In this last case, this was mainly due to serious technical problems with data collection.

information like number of trips per tour and dwelling time.

After taking into account all these aspects, the dataset contains 174,025 tours and 407,828 trips, which is the sample I am considering from now on.

The Traffic Choices database have various ways to flag a trip which can have problems: in the sample, a few trips have incorrect information about trip duration (less than 0.1%) and miles driven (0.046%). About 6.9% of the trips start or end outside the area of the experiment. About 3.4% of trips belong to a tour with dwelling time higher than 48 hours. Less than 0.1% of the trips happened when the household was on vacation. Also, about 0.95% of trips do not have corresponding tour information. This occurs when a trip belongs to a tour which ends or starts off-grid. In 55 instances throughout the trial, a vehicle had a malfunctioning problem with the OBU. The vast majority of these issues were solved within one day.

At the trip level, the dataset contains various variables that indicate the number of miles driven: total miles driven, miles driven in tolled roads between 6am-10pm (toll rate was positive), miles driven in tolled roads between 10pm-6am (toll rate was zero), miles driven in non-tolled roads. The latter measurement is very inaccurate, as includes locations (like Seattle downtown) where the installed OBUs have poor precision. For that reason, I worked almost exclusively with miles driven only in tolled roads.

Because the authors of the Traffic Choice trial worked only at the tour level in their analysis, the available dataset does not contain information of highway miles driven at the trip level, but only at the tour level. However, the trip level data do contain information on toll amount during the trial, and the hypothetical toll amount during the baseline period. With these two variables it is possible to reconstruct the overwhelming majority of highway miles driven at the trip level.

If the trip was the only trip in its tour, the highway miles corresponded to the tour highway miles. If the trip was not the only trip in its tour but it was done between 6am-10pm and within a single toll time band, the number of highway miles driven is given by:

$$M_h = \frac{T - M * \tau_r}{\tau_h - \tau_r}$$

where T is the toll amount, M are the total miles driven in tolled roads, τ_h is the toll rate per mile in highways and τ_r is the toll rate per mile in normal roads. After updating the database, if within one tour there was only one trip left without highway miles information, the highway miles were calculated as the remaining tour highway miles not assigned to the other trips.

After performing these steps, only 1.8% of the sample still did not have information on highway miles driven. These trips were generally trips made during the night, where total toll was always zero, and trips crossing two toll bands.[2] If those trips were crossing 6am or 10pm (i.e. switching from no toll to toll and vice-versa) it was possible to calculate the highway miles for the band subject to the toll based on the total toll amount.

The last step was calculating the total number of miles driven for each band toll and each road type at the week level. For trips crossing two toll bands and driving both in normal roads and highways, the problem was estimating the number of miles

[2] A very small amount of trips crossed more than two bands and were dropped from the sample.

driven in each toll band for each road type. I used a simple algorithm based on the assumption that the share in miles driven in the first time band should be as close as possible to the share of travel time spent in the first time band, subject to a series of constraints. Formally:

$$\min_{M_{r1},M_{h1}} \left(\frac{t_1}{t} - \frac{M_{r1}+M_{h1}}{M}\right)^2 \tag{B.1}$$

$$\text{s.t. } T = M_{r1} * \tau_{r1} + (M_r - M_{r1}) * \tau_{r2} + M_{h1} * \tau_{h1} + (M_h - M_{h1}) * \tau_{h2},$$

$$M_{r1} >= 0,\ M_{r1} <= M_r,$$

$$M_{h1} >= 0,\ M_{h1} <= M_h$$

where M_{r1} and M_{h1} are the miles driven in tolled normal roads and highways in the first toll band, M_r and M_h are the miles driven in normal tolled roads and highways, t_1 is the travel time spent in the first toll band, t is total travel time, T is the total toll, and τ_{r1}, τ_{h1}, τ_{r2}, τ_{h2} represent the toll rate for each type of road and each toll band.

B.2 Residual household endowment

To understand how the endowment is calculated, I regress the actual endowment received by each household with mileage in the baseline period, disaggregated by hour of the day, day of the week and type of road in order to match the groups in the toll scheme.

Results show that mileage predict the vast majority of the household endowment (R-squared equal to 0.9517), but there is still a sizeable component of the endowment not explained by driving behavior. To show that, in figure B.1 I plot the distribution of the residuals in absolute terms (panel A) and in terms of share of the predicted endowment.

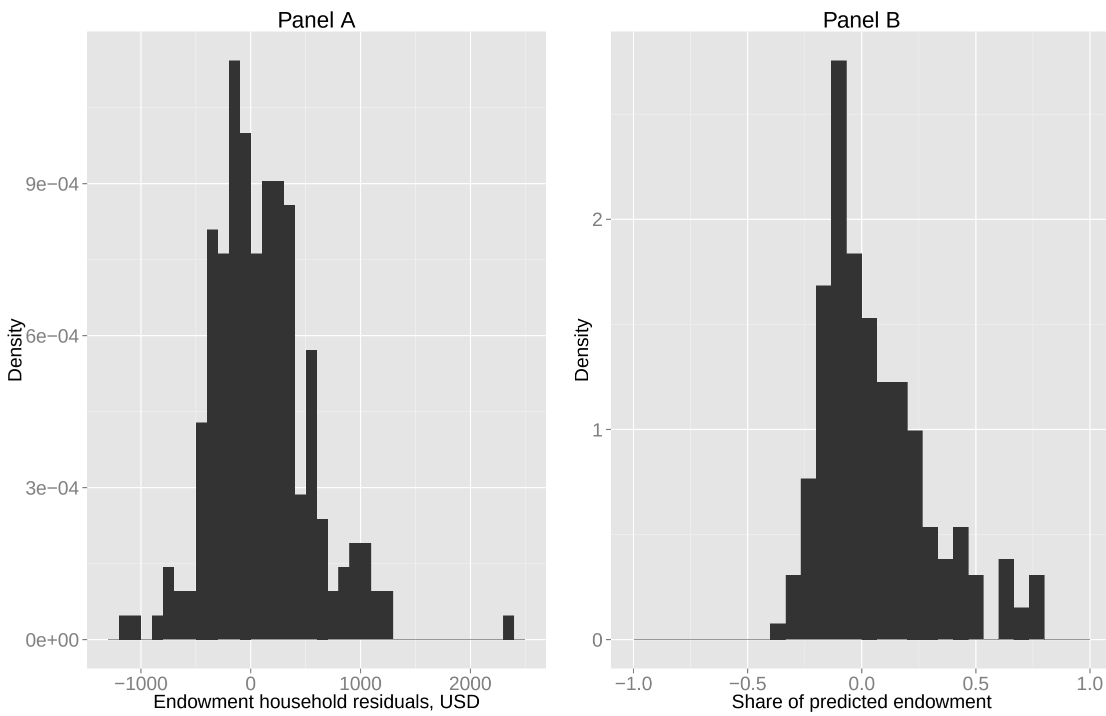

Figure B.1: Distribution of residual endowment **Panel A**: Distribution of residual endowment in US dollars **Panel B**: Distribution of residual endowment in share of the predicted endowment. Observations above 1 are not shown.

Most importantly, the endowment residuals appear to be just weakly correlated with other household characteristics. The correlation between endowment residuals and number of vehicles, number of kids, household income and age of household head are respectively 0.134, -0.114, 0.056, and 0.154.

Bibliography

Abdel-Aziz, A. and Frey, H. C. (2003), 'Quantification of hourly variability in NO_x emissions for baseload coal-fired power plants', *Journal of the Air & Waste Management Association* **53**(11), 1401–1411.

Abou-Zeid, M., Ben-Akiva, M., Tierney, K., Buckeye, K. and Buxbaum, J. (2008), 'Minnesota Pay-as-You-Drive Pricing Experiment', *Transportation Research Record* **2079**, 8–14.

Adamou, A., Clerides, S. and Zachariadis, T. (2012), 'Trade-offs in CO2-oriented vehicle tax reforms: A case study of Greece', *Transportation Research Part D: Transport and Environment* **17**(6), 451–456.

Adamou, A., Clerides, S. and Zachariadis, T. (2014), 'Welfare Implications of Car Feebates: A Simulation Analysis', *The Economic Journal* **124**(578), 1–22.

Agerholm, N., Waagepetersen, R., Tradisauskas, N., Harms, L. and Lahrmann, H. (2008), 'Preliminary results from the Danish Intelligent Speed Adaptation Project Pay As You Speed', *IET Intelligent Transport Systems* **2**(2), 143.

Alberini, A. and Bareit, M. (2016), 'The Effect of Registration Taxes on New Car Sales and Emissions: Evidence from Switzerland', *CER-ETH Economics Working Paper* **16/245**.

Allcott, H. and Rogers, T. (2014), 'The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation ', *American Economic Review* **104**(10), 3003–3037.

AMMA (2008), Monitoraggio indicatori Ecopass: Prime valutazioni, Technical report.

ARPA (2007), Rapporto sulla qualità dell'aria di Milano e provincia, Technical report.

Auffhammer, M. and Kellogg, R. (2011), 'Clearing the air? the effects of gasoline content regulation on air quality', *American Economic Review* **101**(6), 2687–2722.

Bento, A., Kaffine, D., Roth, K. and Zaragoza-Watkins, M. (2014), 'The effects of regulation in the presence of multiple unpriced externalities: Evidence from the transportation sector', *American Economic Journal: Economic Policy* **6**(3), 1–29.

Berry, S., Levinsohn, J. and Pakes, A. (1995), 'Automobile Prices in Market Equilibrium', *Econometrica* **63**(4), 841–890.

Berry, S. T. (1994), 'Estimating Discrete-Choice Models of Product Differentiation', *The RAND Journal of Economics* **25**(2), 242–262.

Bolderdijk, J., Knockaert, J., Steg, E. and Verhoef, E. (2011), 'Effects of Pay-As-You-Drive vehicle insurance on young drivers speed choice: Results of a Dutch field experiment', *Accident Analysis & Prevention* **43**(3), 1181–1186.

Bun, M. J. and Kiviet, J. F. (2003), 'On the diminishing returns of higher-order terms in asymptotic expansions of bias', *Economics Letters* **79**(2), 145–152.

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), 220–256.

Carrillo, P. E., Arun, S. M. and Yiseon, Y. (2013), 'Driving restrictions that work? Quito's Pico y Placa program', *Working Paper* .

Carslaw, D. C., Beevers, S. D., Tate, J. E., Westmoreland, E. J. and Williams, M. L. (2011), 'Recent evidence concerning higher NOx emissions from passenger cars and light duty vehicles', *Atmospheric Environment* **45**(39), 7053–7063.

Chen, Y. and Borken-Kleefeld, J. (2014), 'Real-driving emissions from cars and light commercial vehicles results from 13 years remote sensing at Zurich/CH', *Atmospheric Environment* **88**, 157–164.

Ciccone, A. (2014), 'Is it all about CO2 emissions? The environmental effects of a tax reform for new vehicles in Norway', *Working Paper, University of Oslo* **19**.

Commissione Ecopass (2010), Relazione conclusiva dei lavori, Technical report.

Congressional Budget Office (2011), Alternative approaches to funding highways, Technical report.

Council of State Governments (2010), Focus on: Vehicle miles traveled fees, Technical report.

Danielis, R., Rotaris, L., Marcucci, E. and Massiani, J. (2011), 'An economic, environmental and transport evaluation of the Ecopass scheme in Milan: Three years later', *SIET Working Papers* .

Davis, L. W. (2008), 'The effect of driving restrictions on air quality in Mexico City', *Journal of Political Economy* **116**(1), 38–81.

D'Haultfuille, 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), 444–480.

Eskeland, G. and Feyziolu, T. (1997), 'Rationing can backfire: The "Day without a Car" in Mexico City', *The World Bank Economic Review* **11**(3), 383–408.

Ferreira, J. and Minikel, E. (2012), 'Measuring per mile risk for Pay-As-You-Drive automobile Insurance', *Transportation Research Record* **2297**, 97–103.

Finkelstein, A. (2009), 'E-ZTax: Tax salience and tax rates', *The Quarterly Journal of Economics* **124**(3), 969–1010.

Fullerton, D. and West, S. E. (2002), 'Can taxes on cars and on gasoline mimic an unavailable tax on emissions?', *Journal of Environmental Economics and Management* **43**(1), 135–157.

Gallego, F., Montero, J.-P. and Salas, C. (2013), 'The effect of transport policies on car use: Evidence from latin american cities', *Journal of Public Economics* **107**, 47–62.

Gibson, M. and Carnovale, M. (2015), 'The effects of road pricing on driver behavior and air pollution', *Journal of Urban Economics* **89**, 62–73.

Goddard, H. C. (1997), 'Using tradeable permits to achieve sustainability in the world's large cities: policy design issues and efficiency conditions for controlling vehicle emissions, congestion and urban decentralization with an application to Mexico City', *Environmental and Resource Economics* **10**(1), 63–99.

Grigolon, L., Reynaert, M. and Verboven, F. (2015), 'Consumer valuation of fuel costs and the effectiveness of tax policy: Evidence from the European car market', *Working Paper* .

Hanley, P. and Kuhl, J. (2011), 'National evaluation of mileage-based charges for drivers: Initial results', *Transportation Research Record* **2221**, 10–18.

Holland, S. P., Mansur, E. T., Muller, N. Z. and Yates, A. J. (2015), 'Environmental benefits from driving electric vehicles?', *NBER Working Paper* **21291**.

Hultkrantz, L. and Lindberg, G. (2011), 'Pay-as-you-speed: An economic field experiment', *Journal of Transport Economics and Policy* **45**(3), 415–436.

Huse, C. and Lucinda, C. (2014), 'The Market Impact and the Cost of Environmental Policy: Evidence from the Swedish Green Car Rebate', *The Economic Journal* **124**(578), 393–419.

Keeler, T. E. and Small, K. A. (1977), 'Optimal peak-load pricing, investment, and service levels on urban expressways', *The Journal of Political Economy* **85**(1), 1–25.

Khanna, M., Ando, A. W. and Taheripour, F. (2008), 'Welfare effects and unintended consequences of ethanol subsidies', *Review of Agricultural Economics* **30**(3), 411–421.

Kiviet, J. F. (1995), 'On bias, inconsistency, and efficiency of various estimators in dynamic panel data models', *Journal of Econometrics* **68**(1), 53–78.

Klier, T. and Linn, J. (2015), 'Using Taxes to Reduce Carbon Dioxide Emissions Rates of New Passenger Vehicles: Evidence from France, Germany, and Sweden', *American Economic Journal: Economic Policy* **7**(1), 212–242.

Konishi, Y. and Meng, Z. (2014), 'Can Green Car Taxes Restore Efficiency? Evidence from the Japanese New Car Market', *Tokyo Center for Economic Research (TCER) Papers* **E-82**.

Larsen, L., Burris, M., Pearson, D. and Ellis, P. (2012), 'Equity evaluation of fees for vehicle miles traveled in Texas', *Transportation Research Record: Journal of the Transportation Research Board* **2297**, 11–20.

Lin, C.-Y. C., Zhang, W. and Umanskaya, V. I. (2014), 'On the design of driving restrictions: Theory and empirical evidence', *Working Paper* .

Maryland Department of Transportation (2011), Maryland Climate Action Plan - Appendix, Technical report.

McMullen, B. S., Zhang, L. and Nakahara, K. (2010), 'Distributional impacts of changing from a gasoline tax to a vehicle-mile tax for light vehicles: A case study of Oregon', *Transport Policy* **17**(6), 359–366.

Memmott, J. and Young, P. (2008), Seasonal variation in traffic congestion: A study of three U.S. cities, Technical report, U.S. Department of Transportation.

National Surface Transportation Infrastructure Financing Commission (2009), Paying Our Way: A New Framework for Transportation Finance, Technical report.

Newbery, D. (1988), 'Road damage externalities and road user charges', *Econometrica* **56**(2), 295–316.

Newbery, D. (2005), Road user and congestion charges, *in* S. Cnossen, ed., 'Theory and practice of excise taxation: Smoking, drinking, gambling, polluting and driving', Oxford University Press, Oxford.

Nickell, S. (1981), 'Biases in dynamic models with fixed effects', *Econometrica* **49**(6), 1417.

Parry, I. W. (2004), 'Comparing alternative policies to reduce traffic accidents', *Journal of Urban Economics* **56**(2), 346–368.

Parry, I. W. (2005), 'Is Pay-As-You-Drive insurance a better way to reduce gasoline than gasoline taxes?', *American Economic Review (Papers and Proceedings)* **95**(2), 288–293.

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. and Timilsina, G. R. (2010), 'How should passenger travel in Mexico City be priced?', *Journal of Urban Economics* **68**(2), 167–182.

Paz, A., Nordland, A., Veeramisti, N., Khan, A. and Sanchez-Medina, J. (2014), 'Assessment of economic impacts of vehicle miles traveled fee for passenger vehicles in nevada', *Transportation Research Record* **2450**, 26–35.

Percoco, M. (2013), 'Is road pricing effective in abating pollution? evidence from milan', *Transportation Research Part D: Transport and Environment* **25**, 112–118.

Percoco, M. (2014), 'The effect of road pricing on traffic composition: Evidence from a natural experiment in milan, italy', *Transport Policy* **31**, 55–60.

Puget Sound Regional Council (2008), Traffic Choices study final report, Technical report.

Reese, C. A. and Pash-Brimmer, A. (2009), North Central Texas Pay-As-You-Drive insurance pilot program, *in* 'Proceedings of the Transportation, Land Use, Planning and Air Quality Conference, Denver'.

Robitaille, A., Methipara, J. and Zhang, L. (2011), 'Effectiveness and equity of vehicle mileage fee at federal and state levels', *Transportation Research Record* **2221**, 27–38.

Rotaris, L., Danielis, R., Marcucci, E. and Massiani, J. (2010), 'The urban road pricing scheme to curb pollution in Milan, Italy: Description, impacts and preliminary cost–benefit analysis assessment', *Transportation Research Part A: Policy and Practice* **44**(5), 359–375.

Safirova, E., Houde, S. and Harrington, W. (2007), 'Marginal social cost pricing on a transportation network: A comparison of second-best policies', *RFF Discussion Papers* **07-52**, 1–23.

Sana, B., Konduri, K. and Pendyala, R. (2010), 'Quantitative analysis of impacts of moving toward a vehicle mileage-based user fee', *Transportation Research Record* **2187**, 29–35.

Santos, G. (2004), 'Urban congestion charging: A second-best alternative', *Journal of Transport Economics and Policy* **38**(3), 345–369.

Santos, G., Rojey, L. and Newbery, D. (2000), 'The environmental benefits from road pricing', *Cambridge Working Papers in Economics* **20**.

Schmutzler, A. (2011), 'Local transportation policy and the environment', *Environmental and Resource Economics* **48**(3), 511–535.

Shi, J. P. and Harrison, R. M. (1997), 'Regression modelling of hourly NO_x and NO_2 concentrations in urban air in london', *Atmospheric Environment* **31**(24), 4081–4094.

Vance, C. and Mehlin, M. (2009), 'Fuel Costs, Circulation Taxes, and Car Market Shares: Implications for Climate Policy', *Transportation Research Record: Journal of the Transportation Research Board* **2134**, 31–36.

Vickrey, W. (1963), 'Pricing in urban and suburban transport', *The American Economic Review (Papers and Proceedings)* **53**(2), 452–465.

Virginia Department of Transportation (2008), Vehicle miles traveled (VMT) tax: An alternative to the gas tax for generating highway revenue, Technical report.

Walters, A. A. (1961), 'The theory and measurement of private and social cost of highway congestion', *Econometrica* **29**(4), 676–699.

Washington State Transportation Commission (2014), Washington state road usage charge assessment, Technical report.

Weatherford, B. (2011), 'Distributional implications of replacing the federal fuel tax with per mile user charges', *Transportation Research Record* **2221**, 19–26.

West, S. E. (2004), 'Distributional effects of alternative vehicle pollution control policies', *Journal of Public Economics* **88**(3-4), 735–757.

West, S. E. (2005), 'Equity implications of vehicle emissions taxes', *Journal of Transport Economics and Policy* pp. 1–24.

Wolff, H. (2014), 'Keep Your Clunker in the Suburb: Low-emission Zones and Adoption of Green Vehicles', *The Economic Journal* **124**(578), F481–F512.

Zhang, L. and Lu, Y. (2012), 'Marginal-cost vehicle mileage fee', *Transportation Research Record* **2297**, 1–10.