New Vehicle Feebates: Theory and Evidence*

Nicholas Rivers† Brandon Schaufele‡

September 3, 2014

Abstract

Revenue neutral new vehicle feebate programs are potentially popular policy options for politically constrained governments; yet analyses of these policies have been limited to *ad hoc* proposals. In this paper, we (1) derive expressions for optimal new vehicle feebate schedules, (2) distinguish between categories of revenue neutral policy design and explore the social welfare implications of a class of self-financing constraints referred to as differentiated characteristic revenue neutrality and (3) exploit an extensive, multi-year dataset to evaluate the welfare implications of a long-standing vehicle feebate program in the Canadian province of Ontario. We find that Ontario's feebate program was welfare-enhancing relative to a no feebate scenario but that an optimally designed policy would have yielded additional vehicle sales while reducing fleet-wide emissions.

Keywords: Feebate; revenue neutrality; vehicle choice; environmental policy; externality; optimal tax.

JEL codes: D49, D62, H23, Q58, R48.

*Dylan Morgan and Soodeh Saberian provided excellent research assistance. Funding from the Social Science and Humanities Research Council is gratefully acknowledged. We appreciate the useful comments and suggestions from Mark Brown and Martino Pelli. Versions of this paper have been presented at the Association of Environmental and Resource Economics Annual Meetings, Canadian Economics Association Annual Meetings, Canadian Resource and Environmental Economics Workshop, the Heartland Environmental and Resource Economics Workshop and the Ontario Network for Sustainable Energy Policy.

†Graduate School of Public and International Affairs and Institute of the Environment, University of Ottawa, nrivers@uottawa.ca.

‡Business, Economics and Public Policy, Ivey Business School, University of Western Ontario, bschaufele@ivey.uwo.ca.

1 Introduction

Many governments are unwilling or unable to impose first-best gasoline or mileage taxes to control the external costs associated with driving.[1] This has motivated a search for alternatives to reduce vehicle-related externalities with new vehicle feebate programs receiving increased attention.[2] Feebates are comprised of "rebates" and "fees" levied on new vehicle purchases – purchases of new fuel efficient vehicles are subsidized, while fuel inefficient vehicles are taxed with the ultimate goal of improving environmental quality. Feebates are often presented as a revenue neutral policy (i.e., revenue collected from fees offsets all rebates). Proponents hope that this policy will appeal to governments facing spending and political constraints. This paper offers a derivation of optimal new vehicle feebate schedules, and also exploits an extensive, multi-year dataset of new vehicle registrations to examine the behavioural responses to an actual feebate program implemented in the Canadian province of Ontario. We investigate the welfare consequences of this policy vis-à-vis the optimal revenue neutral feebate and conclude that feebate programs improve upon the no policy status quo but that an optimally designed revenue neutral program would have yielded greater emissions reductions.

Light duty vehicle fuel efficiency and emissions regulations are frequently enacted to achieve fleet-wide environmental targets. The United States has imposed corporate average fuel economy (CAFE) standards since the 1970s and Canada has used similar voluntary fuel consumption guidelines until recently when it adopted fleet-wide greenhouse gas regulations. Many jurisdictions also offer incentive programs to encourage consumers to adopt fuel efficient, electric and hybrid vehicles. Fuel economy standards, however, can be inefficient (e.g., Greene (1991), Thorpe (1997), Kleit (2004), Austin and Dinan (2005), Fischer et al. (2007)) or are potentially non-binding (Small and Van Dender, 2007), while incentive programs require significant funding to be effective (e.g., Chandra et al. (2010), Gallagher and Muehlegger (2011)). Feebate systems are often discussed in conjunction with these policies; indeed, even though the consumer welfare consequences of these programs are less well understood, feebates are viewed as substitutes for CAFE standards (Klier and Linn, 2012; Roth, 2012; Gillingham, 2013).[3]

[1]Parry and Small (2005), for instance, demonstrate that for the United States the optimal gasoline in tax is more than twice its current level and that the "prospects are remote" (pg. 1287) that it will move towards its optimal level, a sentiment echoed by the lukewarm reception to other recent calls for increased gasoline taxes (e.g., Frank (2006) and Karplus et al. (2013)).

[2]US states which have considered feebates include Arizona, California, Connecticut, Maine, Maryland, Massachusetts, New York, Oregon and Wisconsin (Train et al., 1997). Starting in 2007, the Canadian federal government also imposes a "Green Levy" on fuel inefficient vehicles and an "ecoAUTO" rebate on fuel efficient vehicles (Banerjee, 2009). Among other countries, Austria has a feebate program and, in 2008, France introduced their bonus/malus policy.

[3]Feebate programs are usually connected to vehicle-related environmental issues and this is the focus of this study. Yet, these policies can be applied to a more general suite of scenarios including electricity generation (Johnson, 2006) and highway safety. For example, Anderson and Auffhammer (2011) demonstrate that crash incompatibility – accidents which involve vehicles of disproportionate weights – generates safety externalities which have costs equivalent to approximately $0.25 per litre of gasoline. Pivot points and feebate rates could be selected to create a disincentive

Initial research into feebates found that consumers are largely unresponsive to proposed schedules of taxes and subsidies. Train et al. (1997) find a very small demand response to a hypothetical feebate, yielding minor changes in consumer surplus and Greene et al. (2005) illustrate that welfare calculations from *ex ante* studies of feebate programs vitally depend on assumptions about consumer discount rates. Sallee and Slemrod (2010) examine policies aimed at encouraging the use of fuel-efficient vehicles through the lens of "notches", step-wise approximations to smooth Pigouvian subsidies. They find manufacturers strategically respond to feebates by altering their fuel economy ratings via "local" or small design modifications.[4]

d'Haultfoeuille et al. (2014) and Adamou et al. (2014) are the studies closest to the present one. d'Haultfoeuille et al. (2014) evaluate France's bonus/malus program and find that consumers had a dramatic response to French feebate policy. Counterintuitively, they demonstrate that the feebate generated a large rebound effect whereby, even though the economy was in recession, vehicles sales *increased* by 13 percent following the program's introduction, a result which led to increased vehicle emissions in the short run. French feebate rates however were set in consultation with industry and, as such, auto manufacturer's influence over the schedule is unclear[5]: d'Haultfoeuille et al.'s (2014) analysis of the bonus/malus program only covers a single year and does not have access to a control group unaffected by the program. Adamou et al. (2014) examine a prospective German feebate program. They determine that the German program would be welfare-decreasing and any reductions in carbon dioxide emissions would be insufficient to compensate for the distortionary effect of the program. Yet, they do establish that it is possible for well-designed feebates to be welfare-enhancing.

This study makes two advances over the existing research. First, we estimate behavioural parameters using a rich multi-year dataset that includes all vehicle registrations from each Canadian province from 2000 to 2010. We also map vehicle registrations to the forward sortation areas (FSA).[6] This extensive dataset contains over 16 million observations and allows us move beyond the extant literature as we (1) evaluate an actual long-running new vehicle feebate program (rather than conduct an *ex ante* study of a proposed program) and (2) are able to accurately identify the reduced form behavioural response to the Ontario feebate program without worrying about "transaction price" or unobserved heterogeneity issues that can influence structural vehicle choice models. Second, unlike other analyses of feebate policies, we derive an optimal feebate schedule which is a function of the externality costs of driving and explore the consumer welfare implications

for heavier vehicles purchases and thus improving social welfare (e.g., Greene (2009)).

[4]Sallee and Slemrod (2010) suggest that automobile manufacturers may substitute vehicle parts to reduce weight, use low-friction lubricants or make small body changes (e.g., install spoilers or side skirts) in an effort to surpass some fuel efficiency threshold.

[5]d'Haultfoeuille et al. (2014) note that the feebate rates were established at a multi-stakeholder roundtable "in Autumn 2007 by the newly elected president [Sarkozy]" (pg. 6) and then "implemented with unusual speed" (pg. 5).

[6]FSAs are the first three digits of a Canadian postal code. There are roughly 1,600 FSAs in Canada, with an average population of slightly more than 20,000 individuals.

of this policy. This optimal schedule is used to evaluate the efficacy of the program implemented in Ontario and provides the appropriate benchmark with which to evaluate other proposed and existing programs.

We highlight three main results. First, our empirical results suggest that a $1,000 fee (rebate) consistently causes a 30 percent reduction (increase) in unique vehicle market share. These estimates remain robust even after controlling for unobserved heterogeneity and vehicle-specific preferences, varying the resolution of vehicle-region-time specific fixed effects, changing the level of data aggregation and allowing for different vehicle substitution possibilities. Second, we demonstrate that the optimal revenue neutral feebate can be written as a function of the optimal new vehicle Pigouvian tax. The optimal feebate is always less than the optimal Pigouvian tax and, in the special case where consumers do not adjust along their extensive margin, the two policies generate identical social welfare. Third, we find that Ontario's feebate program actually increased emissions relative to a no feebate scenario, yet as it was revenue positive the program generated an increase in welfare equal to $16.5M. An optimally designed revenue neutral policy, in contrast, would generate a notable increase in vehicle sales while still reducing annual emissions by 3M tonnes of CO_2, ultimately yielding a welfare gain of $81.0M.

Revenue neutrality is at the core of many feebate proposals and is central to our analysis. We start by delineating categories of revenue neutral policy emphasizing how the concept is applied in distinct contexts. Next we address an unanswered question: do modest feebate schedules such as those recommended by many proposed policies actually influence behavior? We answer this in the affirmative – consumers do respond to moderate fees and rebates. We then discuss and present expressions for optimal revenue neutral feebate programs and calculate welfare implications under several counterfactual scenarios. Finally, we consider threats to optimal policy design, which policymakers may need to address in the process of implementing optimal revenue neutral feebates.

2 Revenue Neutral Environmental Policy

A central feature of many new vehicle feebate pitches is the potential for "revenue neutral" designs. Indeed, revenue neutrality is a tenet in many environmental policy recommendations. However, the term revenue neutral is often applied in a vague manner. Distinct types of revenue neutrality exist, each with distinct implications for policy design and social welfare. In particular, we distinguish between cross-base revenue neutrality, as found in most tax shifting studies (for example Bovenberg and de Mooij, 1994; Goulder, 1995), and two classes of intra- or within-base revenue neutrality: deposit-refund systems and differentiated characteristic revenue neutrality. Feebate schemes are a form of differentiated characteristic revenue neutrality. After highlighting the differences between these classifications, the remainder of the paper develops the consumer welfare implications of differentiated characteristic revenue neutrality by working through its application to new vehicle feebates.

2.1 Cross-base Revenue Neutrality

A common environmental policy narrative claims that governments should impose Pigouvian taxes on environmental externalities such as greenhouse gas emissions and use revenues to offset pre-existing distortionary taxes on labor and capital. We refer to this type of scheme as *cross-base* revenue neutrality since a new base is created by taxing environmentally damaging outputs. Revenues from this new tax base are then used to reduce rates on an existing tax base. This revenue neutral tax substitution yields improved environmental quality and in some cases may produce a so-called "double dividend", where the potential for a double dividend depends on the relative magnitudes of the revenue recycling and tax interaction effects (Bovenberg and de Mooij, 1994; Parry, 1995).

The defining characteristic of cross-base revenue neutrality is the ability and willingness of governments to engage in tax substitution and a prerequisite is the existence of both a series of distortionary taxes and the legislative or political capacity to simultaneously adjust these rates. Revenue neutrality in the double dividend context, for instance, involves levying revenue positive environmental charges and simultaneously reducing, say, distortionary personal and corporate income tax rates so that expected total government receipts remain unchanged.[7,8]

Governments however may be explicitly or implicitly constrained in the their ability to engage in the tax substitutions required for cross-base revenue neutrality. Constitutional delineation of taxation authority, for example, may limit distinct levels of government from pursuing particular environmental policies – e.g., many regional and municipal governments are prevented from taxing income or profits. Similarly, political or equity constraints may constrain government's ability to adjust existing tax schedules – decision-makers may feel it is politically inexpedient to impose revenue positive environmental taxes even if they are welfare improving. It is necessary therefore to develop the concept of within-base revenue neutrality. Environmental policy that is within-base or *self-financing* may have a larger probability of implementation as it circumvents key challenges in real-world policy-making.

[7]There are two ways to define this class of cross-base revenue neutrality. It may refer (i) to a situation where the revenue from the environmental tax equals the revenue recycled or (ii) to the case where the environmental charge is set to equate government deficits or expenditure. Under definition (i), the government's budget may change even though the policy satisfies one defintion of revenue neutrality.

[8]The Canadian province of British Columbia provides an example of cross-base revenue neutral environmental policy. British Columbia levied a tax on carbon dioxide emissions and simultaneously reduced personal and corporate income taxes such that expected government revenues remained unchanged (Harrison, 2012; Rivers and Schaufele, 2013).

2.2 Within-base Revenue Neutrality

While cross-base revenue neutrality involves imposition of a new environmental tax and simultaneous reduction of an existing tax, a within-base tax shift achieves revenue neutrality through changes in tax rates within a single tax base. Within-base revenue neutrality implies that a program is *self-financed* with program revenues and expenditures disconnected from the other features of the government's budget. Characteristics of two types of within-base revenue neutral program are discussed: deposit-refund systems and differentiated characteristic programs.

2.2.1 Deposit-Refund Systems

Deposit-refund programs exploit a two-step process: a tax is levied on a product at its point of purchase and then, at a later date, the return of the product or its packaging triggers a refund of all or part of the initial fee (Walls, 2011). Deposits and refunds are directly linked to specific products and some inter-temporal structure of consumption is exploited – i.e., the consumer must undertake multiple transactions to obtain, first, the product where she pays the price plus deposit and then, once she has finished consuming the product, the refund. Deposit-refund schemes are useful if the environmental objective involves central collection of a product at the end of its life, for reuse, recycling or to reduce dumping. Under specific assumptions, deposit-refund systems can replicate first-best outcomes while maintaining revenue neutrality as each deposit received is refunded upon the return of the product or packaging. Importantly, these systems are generally established as standalone programs and are independent of governments' other revenue and expenditure decisions avoiding the complications that may arise with cross-base revenue neutral schemes.

Deposit-refund systems have been applied to beverage containers, batteries, tires, motor oil and electronics (Palmer and Walls, 1997; Fullerton and Wolverton, 2000). In certain circumstances, these policies have advantages over conventional Pigouvian taxation. If pollution monitoring, tax collection or regulatory enforcement is imperfect (e.g., tax evasion or "midnight dumping" problems exist), then deposit-refund systems provide firms and individuals incentives to undertake costly abatement activities. Mrozek (2000) demonstrates that revenue neutral deposit-refund systems approach first-best Pigouvian outcomes if compliance costs are low and demand is inelastic. Fullerton and Wolverton (2000) extend the basic results to show that even if taxes cannot be levied on dirty good production, a two-part instrument, consisting of a clean production subsidy and output tax, can mimic the optimal deposit-refund system.

Several challenges confront deposit-refund systems in practice. First, there will be non-compliance for any subset of depositors whose transactions costs are greater than the refund (Mrozek, 2000). This is a consequence of the multi-transactional design. Second, in general, all products in a class are treated identically even though they make generate heterogeneous social costs. Aligning charges and external costs could lead to welfare improvements (Starrett, 1972). Third, sufficient physical and administrative infrastructure is required to support deposit-refund systems. Consumers, for instance, need a place to take their empty bottles and cans and retailers must collect and hold/transfer deposits. For consumer goods such as cans and bottles, physical infrastructure may be a small barrier; however, not-in-my-backyard concerns may influence the location of other, potentially toxic products. Irrespective of these potential challenges, deposit-refund systems provide a viable structure within which to design revenue neutral environmental policy aimed at recovery of goods at end-of-life.

2.2.2 Differentiated Characteristic Revenue Neutrality

While deposit-refund systems use a time dimension as a central policy parameter, markets that contain many differentiated products, such as the automobile market, are able to exploit heterogeneity of product characteristics to ensure revenue neutrality. Vehicle markets contain a range of non-identical products, each which can be described as a bundle of distinct attributes. These attributes include body color, trim and upholstery in addition to fuel economy, engine size and number of cylinders. Specific attributes within the characteristic space of these differentiated products are associated with different external costs. The link between fuel economy and greenhouse gas emissions is the focus of this paper, but vehicle weights and footprints also yield diverse externalities. If vehicles are differentiated according to fuel economy for instance, we obtain uni-dimensional ranking of the differentiated product space according to each product's externality cost. A differentiated products environmental policy then levies a fee on the purchase of fuel inefficient vehicles that occupy one end of the distribution. The revenue from this fee is used to subsidize the purchase of relatively fuel efficient vehicles situated at the other end of the fuel economy distribution. Similarly, a differentiated characteristic policy based on vehicle weight would tax heavy models while subsidizing low weight vehicles.

The defining property of differentiated products revenue neutrality is that the heterogeneity in the environmental damage caused by characteristics of a class of goods determines tax treatment. By redistributing purchases across the distribution of goods, differentiated characteristic feebate policies lower the sales-weighted average fuel economy and provide continual incentives to reduce external costs. As with deposit-refund systems, feebate systems can be designed to be self-financing or revenue neutral if tax revenue on products with high external costs is completely exhausted in providing rebates on products with low or negative social costs. Indeed, in our treatment of new vehicle feebates, we impose the revenue neutrality constraint.

As with any policy, there are challenges to implementing feebates. First, feebates operate by creating winners and losers. Consumers who purchase products with higher than average social costs, such as heavy vehicles with poor fuel economy, bear the cost of a feebate. Consumers who purchase products with low social costs, in contrast, may be subsidized with the policy. Programs structured with transparent winners and losers often face resistance to implementation. Second, although feebates can be designed to be revenue-neutral *ex ante*, they may not be revenue neutral *ex post* if consumer behaviour deviates from expectations (this is true of other revenue neutral policies as well). This is a challenge confronting any attempt to implement optimal policy, one we partially address in section 6. If revenue neutrality is an important program characteristic, continual adjustment of the program may be necessary over time. Third, feebates are generally applied on the sale of the vehicle not on vehicle use. Since it is use of the vehicle generates most external costs, the disconnect between the policy and the social cost reduces efficiency.

Important knowledge gaps also exist on feebate policies. Although several feebate programs have been implemented worldwide, few have received attention from economists and to our knowledge there has been no effort to conceptualize an optimal feebate design. A larger outstanding issue is the dearth of evidence on the degree to which feebates actually affect consumer behavior. This lack of empirical evidence on effectiveness may be an additional barrier to their adoption.

3 Feebates in Practice: Do modest feebates influence behavior?

Feebate policies have received increasing attention in the literature (e.g., Adamou et al., 2011; Gillingham, 2013). However, very few studies evaluate extant feebate programs. Rather, recommendations rely on simulations of hypothetical vehicle feebates using models that are parametrized with data from markets without feebates. To the extent that actual behavior does not match modeled behavior, important questions remain unanswered. Prominent among these is whether modest fees and rebates, such as those in existing and proposed programs, actually influence household vehicle purchase behavior. Automobiles are durable goods that comprise a large share of household budgets. Suggested levies and subsidies are typically a small fraction of a new vehicle's total cost. Little evidence supports the conclusion that limited feebates are sufficient to alter household decisions, particularly if consumers face cognitive costs or anchor on specific model attributes.

Before introducing our approach to optimal feebates, we provide evidence that modest feebates do indeed affect vehicle purchase behavior. We exploit one of the handful of existing vehicle feebate programs, a longstanding policy that existed in the Canadian province of Ontario, using data on the population of Canadian new vehicle registrations over the 2000-2011 period. Despite being one of only a handful of existing programs, Ontario's feebate has received scant attention.[9] Our evaluation of Ontario's feebate enables us to conclude that modest feebate schedules do produce economically and statistically meaningful changes in consumer vehicle purchase decisions – for example, a fee of $1000 reduces the market share of a vehicle by 30 percent.

In the next four subsections, we first review the main characteristics of the Ontario Tax and Credit for Fuel Conservation program (Ontario's feebate program). Second, we use a case study to

[9] Bregha and Moffet (1995) provide a qualitative discussion of the program, focusing on implementation details and political economy; however, they do not provide any quantitative analysis.

show how parameters are identified in our empirical model. Our example focuses on the Ford Mustang, which was redesigned in 2004, leading to a change in its feebate treatment. Next, our dataset is introduced – as mentioned, we exploit the Canadian population of new vehicle registrations categorized by fuel economy over an eleven year period. Results on the changes in vehicle-specific market shares attributable to Ontario's feebate program are then presented. Ultimately we demonstrate that modest fees and rebates do produce notable behavioral changes.

3.1 Overview of Ontario's Policy

Ontario's Tax and Credit for Fuel Conservation or feebate program underwent three iterations. The initial policy, developed in 1989, was a gas guzzler-type tax on fuel inefficient cars (Government of Ontario, 1989).[10] A tax was levied on all new car purchases that had a highway fuel consumption rating exceeding 9.5L/100 km. As shown in Table 1, the initial fee was $600 for passenger cars consuming between 9.5-12.0L/100 km increasing to $3500 for cars with ratings greater than 18.0L/100 km. One year later, in 1990, tax rates were doubled and extended to cover cars with a wider range of fuel economy ratings. Fees on the most fuel inefficient vehicles now equalled $7000, while vehicles with ratings in the 8.0-8.9 and 9.0-9.4L/100 km were subject to taxes of $200 and $700, respectively.[11]

The announcement of the 1990 tax schedule proved extremely controversial (Bregha and Moffet, 1995). Concerns were raised by both the Canadian Auto Workers union and by manufacturers. Lobbying and political pressure persuaded the government to revisit the policy. The version of the program that we analyze was formulated in 1991, following consultation with stakeholders, and remained unchanged until 2010.[12] This final iteration expanded the policy to incorporate a larger number of vehicles, notably sport utility vehicles (SUVs) were included (passenger vans and pick-up trucks remained exempt). The program also started offering rebates to fuel efficient cars and lowered the threshold at which the tax applied to 6.0L/100 km. The tax rate on the bottom two brackets was also reduced relative to the 1990 levels – the fee on cars with fuel efficiency ratings of 8.0-8.9L/100 km dropped from $200 to $75, while for the 9.0-9.4L/100 km bracket it decreased from $700 to $250.[13] Table 1 displays the full schedule of taxes and rebates for each version of the

[10]Canada's fuel efficiency rating are determined by Natural Resources Canada (NRCan) in a near identical manner to the US's Enivonrmental Protection Agency (EPA) methodology.

[11]Note that the Ontario program is based on rated highway fuel consumption, rather than weighted city-highway fuel consumption. Since emissions from vehicles are optimally determined by weighted city-highway fuel consumption rather than highway consumption alone, this choice adds a distortion to the Ontario policy.

[12]Our analysis covers the period 2000-2010, a decade following the last update to the program and sufficiently distant from the program revision to allow virtually complete change of the new vehicle market. Therefore, although the policy was designed with feedback from the vehicle industry in 1990, it is reasonable to treat the policy as exogenous for the period covered by our analysis.

[13]The tax and credit applies to all vehicles which are purchased, leased or rented in the Province of Ontario. Vehicles purchased outside of Ontario, but which are registered in the province, are still subject to the tax and credit (fees are paid or rebates received at the time of registration). Motorcycles and vehicles sold to non-residents are

program. On July 1, 2010, Ontario eliminated its feebate program as the province underwent large scale tax reform.

Over the life of the program, the tax on fuel inefficient vehicles generated an average of approximately $30 million ($2002) per year. However, tax revenue decreased throughout the period of analysis, from a high of over $40 million to a low of less than $20 million in 2009. In contrast, tax credits associated with the program increased over the period of analysis from $1.9 million in 2001 to $6.6 million in 2010. Although net revenue fell throughout the period of analysis, the combination of fees and rebates was revenue positive in every year that the program existed. On a per-vehicle basis, the feebate resulted in an average fee of about $148 per vehicle across all new vehicles purchases. Figure 1 displays these trends.

Figure 2 decomposes total vehicle sales according to tax treatment under the Tax and Credit for Fuel Conservation program. At the start of the program, most of the vehicles sold in Ontario were passenger cars, subject to the $75 tax. A transition occurred, over the next decade, from passenger cars towards SUVs. Along with this car-SUV substitution, mean fuel efficiency improved resulting in observable increases in the proportion of passenger cars eligible for rebates. Evident in Figure 2 is that the majority of Ontario vehicle purchases were subject to relatively modest fees and rebates, motivating our analysis of the behavioral response to small feebate rates.

3.2 Ontario's Feebate and the Ford Mustang

Before providing detailed empirical results on the effect of the feebate, we provide some initial graphical evidence using the Ford Mustang as a case study. The Ford Mustang is a vehicle with a reputation for performance and power; it is not a car known for fuel efficiency. Mustangs were subject to a fee under Ontario's feebate program and, during our period of analysis, underwent a complete redesign. Introduced in 2004, the fifth generation Mustang had an increased engine size and more horsepower than its predecessor. This redesign meant the car also received different treatment under Ontario's feebate program. As such, the Ford Mustang crosses two feebate brackets, provides an example of how a feebate works both in theory and practice and illustrates how we identify the effect of the vehicle feebate using our data.

Figure 3 illustrates how the Mustang was treated under the Ontario feebate program and how its tax treatment is related to its market share. The top panel shows the NRCan rated highway fuel consumption for a new automatic transmission 4.6 litre 8 cylinder Ford Mustang in each year from 2000 to 2010. This vehicle had a fuel consumption of 9.1 L/100km in 2000, a fuel consumption of less than 8.9 L/100 km from 2001 to 2003, and a fuel consumption greater than 9.0 L/100 km in each year from 2004 to 2009 before falling back below 8.9 L/100 km in 2010.[14] A threshold in Ontario's

exempt.

[14]In 2010, the body of the car had minor redesigns which improved its drag coefficient and consequently its fuel efficiency.

feebate scheme exists at 8.9 L/100 km, such that the Ford Mustang was taxed at $75 in 2001 through 2003 and at $250 between 2004 and 2009. The bottom panel of Figure 3 displays the difference of between the market shares of the vehicle in the rest of Canada and the market share of Ontario. Early in the decade, when the Mustang was taxed at $75, the difference between market shares in Ontario and other provinces was small. However, following the redesign, a higher feebate rate applied. The Mustang's market share in Ontario fell relative to the rest of Canada (the difference between the market share in the rest of Canada and Ontario grew). Even a small change in the fee of $175 per vehicle led Ontario consumers to substitute away from the new Mustang. We interpret this as suggestive evidence that consumers responded to the fuel economy disincentive associated with the feebate program. Still the graphical evidence is not over-whelming and many unobserved factors may motivate these trends in market shares. Whether similar substitution patterns hold for other models, across feebate rates and after controlling for confounding factors is not obvious from the figure. In the following analysis, we use a range of high dimensional fixed effects to control for a wide suite of unobserved characteristics in order to identify the true underlying behavioral response to the feebate scheme and, in section 5, determine whether the feebate program is welfare improving.

3.3 Data

Proprietary data from Desrosiers Automotive Consultants, provider of RL Polk data in Canada, is used throughout this study. The dataset covers the population of private vehicle registrations in Canada over the period 2000 to 2010. The focus of this paper is on new vehicle sales and associated policy, rather than registrations. Thus, we drop all observations for which the model year of the vehicle is not equal to the registration year.[15] We treat a registration of a model year t vehicle in year t as a sale.

We observe all vehicle registrations by make, model, series, and model year, in addition to each vehicle's engine characteristics. We define a "unique vehicle" as a combination of make, model, series, engine size, and number of cylinders. We then merge these data on vehicle sales with information on rated city and highway vehicle fuel consumption, from Natural Resources Canada.[16] Because the different data sources have distinct naming conventions, successfully merging the datasets required significant manual processing. Further, for a small number of low-volume vehicles, rated fuel consumption data was not available from NRCan. For these vehicles, we obtained fuel economy ratings from the US Environmental Protection Agency (EPA).[17] In total, we successfully rated

[15] For example, for vehicles registered in 2003, we treat model year 2003 vehicles as vehicle sales. In reality, some 2004 model year vehicles will also be sold in 2003, and we instead treat these as sales in 2004. This should not present a problem in our analysis, since we are examining the long-run impact of the feebate program, which is unchanged during the period we study.

[16] See `http://oee.nrcan.gc.ca/transportation/tools/fuelratings/ratings-search.cfm`.

[17] See `http://www.fueleconomy.gov/`. As EPA measures fuel economy in gallons per mile and NRCan measures fuel

the fuel economy for virtually all passenger vehicles in the Polk data set, yielding a population of vehicles classified according to their fuel efficiency rating. In addition to fuel efficiency and vehicle characteristics, we also incorporate data on retail gasoline prices, demographics, provincial gross domestic product and vehicle sales retrieved from Statistics Canada.

Data are used at two levels of aggregation. First, we estimate models exploiting cross-provincial variation. Feebates are implemented at the provincial-level and we are interested in the implications of the policy at the level of the intervention. Our policy evaluation thus focuses on data at this coarser level of aggregation. After matching, the complete dataset includes approximately 600 unique vehicle models across 10 provinces and 11 years for a total of nearly 60,000 observations. Second, we also map registrations to forward sortation areas (FSAs).[18] There are roughly 1,600 FSAs in Canada. Moving to FSAs enables a more detailed picture of vehicle ownership trends as we are able to investigate intra-provincial variation. There are over 16 million observations in the FSA dataset. Appendix A presents summary statistics for the data.

3.4 Results

We demonstrate that Ontario's feebate had an economically meaningful and statistically significant effect on the mix of vehicles. A $1000 fee reduces per vehicle market share by approximately 30 percent, a result that is robust across a wide range of specifications. Modest fees and rebates do influence behavior.

We present our empirics in several steps. First, the econometric specification and estimation methodology are discussed. Provincial-level results are then presented followed by estimates from disaggregated FSA-level data. Finally, nested logit models are reviewed allowing for correlations of vehicle market shares within pre-defined classes. Coefficients from these nested logit models will be used when we evaluate Ontario's program vis-á-vis an optimal feebate design. Throughout, we employ a conventional differentiated products demand framework (e.g., Berry, 1994; Berry et al., 1995; Nevo, 2000) where markets are defined as province-years or FSA-years. Effort is concentrated on identifying the magnitude of the feebate effect. To this end, we fully exploit the panel structure of our large dataset. A notable innovation within this framework is the inclusion of many high-dimensional fixed effects. Conventional statistical software is unable to invert matrices with so many parameters, so we successively apply a Frisch-Waugh-Lovell demeaning procedure to remove unobservables and obtain unconfounded estimates of our parameter of interest.

---

consumption in litres per 100km, we converted EPA ratings to fuel consumption ratings using appropriate conversion factors.

[18]FSAs are the first three digits of Canadian postal codes.

Econometric Specification

Differentiated products demand models are well-known, so we only briefly present our approach. Appendix B fully describes our application of this methodology.

The effect of the feebate is found by estimating the following equation:

$$\log\left(\frac{s_{kpt}}{s_{0pt}}\right) = \theta_{vt} + \alpha F_{kpt} + \phi \log G_{kpt} + \mu_{kpt} \tag{1}$$

where s is the market share of a unique vehicle, F is the dollar value of the feebate, G is the gasoline cost in dollars per kilometer traveled and μ is the error term. k indexes vehicles (and 0 is the outside option of not purchasing a vehicle), p indexes provinces and t indexes time. The index k is at the resolution of a 'unique vehicle'.

Different specifications of the model include distinct arrays of vehicle-time fixed effects, denoted by θ. These parameters capture average preferences for a vehicle in each year. The index v is used to control the number of fixed effects that enter the estimating equation. Models are estimated where v is (i) the set of all vehicle make-model combinations, (ii) the set of all make-model-series combinations, and (iii) the set of all make-model-series-engine size-cylinder combinations. In (iii), the sets v and k are identical, such that we include a separate fixed effect for each unique vehicle-by-year in the data. The fixed effects, θ, control for unobserved consumer preferences that are based on detailed vehicle characteristics. We employ three levels of fixed effect coarseness for two reasons. First, estimating such a large number of fixed effects reduces the statistical power of the estimation, possibly soaking up too much variation. Consumers may have preferences over more general vehicle characteristics, adequately described by make and model; including these extra parameters needlessly reduces the precision of the estimates. Second, by examining models with different combinations of parameters, we are able to test the sensitivity of the estimates to different assumptions about consumer preferences. That our coefficient of interest is stable across a variety of models credibly supports our assertion that unobservables are not driving the elasticities (Oster, 2013).

α is the coefficient of interest. If the feebate increases by one unit, α measures the responsiveness of the log of the odds ratio from choosing a particular vehicle relative to the outside option. We present our results with feebates measured in thousands of dollars where fees are positive and rebates enter as negatives. We expect and find that α is negative – an increase in the tax reduces vehicle demand, all else constant. The final parameter in (1) is ϕ, which captures the sensitivity of consumers to changes in gasoline cost. Several recent studies have examined the effect of gasoline prices on vehicle purchase decisions (e.g., Busse et al. (2013); Klier and Linn (2010); Li et al. (2009)).

Estimation Procedure

Several of our models contain hundreds of thousands of fixed effects across several dimensions. Standard estimation techniques are not feasible with so many coefficients. Our objective is to accurately identify a single feebate parameter however. As such we invoke the Frisch-Waugh-Lovell Theorem to annihilate the fixed effects via application of a repeated demeaning process (Davidson and MacKinnon, 1993).[19] This Frisch-Waugh-Lovell procedure enables us to shrink the parameter set and use transformed data to correctly estimate effect of feebates on vehicle market shares. There are two advantages of this approach: (i) we are able to flexibly control for a broad array unobserved confounders, while (ii) dramatically relaxing computational costs involved with estimating models with millions of observations and hundreds of thousands of fixed effects. A similar approach was applied in Carneiro et al. (2012).

A specific example clarifies this procedure. Column 3 of Table 3 presents results using 16,801,960 observations at the FSA-level. This specification includes two sets of fixed effects. First is a vehicle-by-year fixed effect defined at the resolution of make-model-series-engine size-number of cylinders. These coefficients capture a range of vehicle-specific attributes that are common across markets. With 11 years of data and approximately 600 unique vehicles, 6600 coefficients would need to be estimated, a number which is reasonable for on most personal computers. Next, vehicle-by-FSA effects are added. Unobserved commute times in suburban Toronto, for instance, are presumably longer than those in St. John's, Newfoundland. Thus, residents of Toronto may have different unobserved preferences for fuel economy when compared to Newfoundlanders. This suite of parameters captures these FSA-specific preferences. In this model however, up to 1600*600 or 960,000 additional parameters would need to be estimated. Inverting a 950,000 by 950,00 matrix requires approximately 1Tb of memory, well beyond what is available on most computers. The Frisch-Waugh-Lovell procedure overcomes these computational obstacles.

Step one of the Frisch-Waugh-Lovell procedure calculates the mean for all 6600 vehicle-years. Each of these $k * t$ group means are then subtracted from all remaining variables in the second step. This demeaning procedure is successively repeated for each of the fixed effects (e.g., vehicle-by-FSA). Once all group means have been removed through demeaning, a regression of the transformed dependent variable on the transformed independent variables yields identical estimates of the remaining parameters as a regression that includes the full suite of fixed effects. In our example,

[19] The `areg` package in Stata applies a related procedure, but only works when a single high-dimensional fixed effect. As we employ multiple high-dimensional fixed effects, a recursive method was needed. Recently, user provided Stata code automates a similar procedure in the `reg2hdfe` package (see Guimaraes and Portugal (2010)). Likewise, all of our code will be freely provided on our websites. As a final note, to the best of our knowledge, we are the first to apply this class of estimation approach using instrumental variables.

we estimate:

$$\log \widetilde{\widetilde{\left(\frac{s_{kft}}{s_{0ft}}\right)}} = \alpha \tilde{\tilde{F}}_{kft} + \nu_{kft} \tag{2}$$

using least squares where the $\sim$ over a variable represents the transformed data[20] and ν_{kft} is the error term.

Throughout our analysis, we report standard errors clustered by province-years or FSA-years. Standard errors calculated from the the transformed estimating equations, e.g., (2), are correct, even with clustering, once adjusted for degrees of freedom (Guimaraes and Portugal, 2010). However, not all of the fixed effects are identified in the data. As the demeaning algorithm is automated, we do not precisely count how many parameters are removed. We therefore adopt a conservative rule of thumb and assume that all fixed effects are identified in our dataset. This implies that the reported standard errors are over-estimated and we are less likely to reject a null hypothesis of no effect for our feebate parameter.

Provincial-level Results

Table 2 displays province-level results. In all cases, the dependent variable is the log difference in market share of a unique vehicle type and market share of the outside option. Reported coefficients demonstrate the effect of the feebate on the market share of a vehicle. Three panels of estimates correspond to different resolutions of fixed effects. In Panel A, we include vehicle make-model fixed effects. Panel B has vehicle make-model-series and Panel C includes vehicle make-model-series-engine displacement-number of cylinders fixed effects. In each case, we supplement these with other geographic, vehicle and time fixed effects.

In total, Table 2 contains six specifications for each vehicle definition. The parameter of interest, Feebate, indicates the percent change in vehicle-specific market share that results from a $1000 fee (rebates are treated as negative fees). Upon initial inspection, a key observation about Table 2 is the consistency of sign, significance and magnitude of the coefficient of interest across fixed effect structures. Even modest feebates have meaningful effects on behaviour.

Moving across the six specifications, identification is conditional on capturing different unobserved sources of variation that may be correlated with the feebate. All models control for gasoline cost per kilometer (dropping this variable does not significantly affect our main results, however). Starting on the left-hand side of the table, (1) includes time invariant vehicle and province-class fixed effects. With roughly 700 estimated parameters in Panel A and around 3250 in Panel C, this is our most parsimonious specification with vehicle fixed effects capturing long-standing preferences for specific vehicle types that are constant across the country and province-class fixed effects cap-

[20] This study is interested in the feebate parameter, α, however it is possible to recover the fixed effects.

ture geographically-distinct preferences over vehicle class that could arise as a result of differences in climate, geography, urbanization, culture or other slow-moving factors. Identification might be compromised if there are time-variant unobserved factors at a provincial level that are correlated with the feebate or if province-specific preferences over vehicles are specific to individual vehicles, rather than just vehicle class.

Results in (1) suggests that a $1,000 fee is associated with a reduction in the market share of a vehicle of 40 to 50 percent.[21] The level of analysis is a unique vehicle, so the market share reduction is in comparison to the counterfactual market share of this unique vehicle. The effect is precisely estimated, with a standard error equal to 11 percent.

(2) adds class-year fixed effects to (1). The regression therefore controls for time-varying vehicle class preferences at a national level. Figure 2, for example, shows the large increase in sport utility market share over the decade. Enlarging the set of controls has only a minor influence on the feebate coefficient: a $1,000 fee is now associated with a 36 to 46 percent decline in a vehicle's market share. Once again, the effect is precisely estimated.

Region-class-year fixed effects are added in (3).[22] We now control for heterogeneously changing preferences over time between provinces. Intuitive regional groupings exist in Canada based on common geography, history and economic characteristics. Yet, as with the other specifications, these fixed effects leave our coefficient of interest largely unchanged. The effect is again precisely estimated.

(4) includes vehicle-province and class-year fixed effects, accounting for any province-specific but time invariant preferences over vehicles. This strategy controls for average preferences for each unique vehicle by province, identifying the effect of the feebate from within province-vehicle variation in market shares. This specification also includes class-year fixed effects, which capture the shift in preferences towards larger sport utility vehicles. The results in this column suggest that a $1,000 feebate is associated with a reduction in market share of 27 to 57 percent percent for a vehicle model. The point estimates are nearly identical to previous models, but are not estimated precisely – inclusion so many fixed effects at this level of analysis soaks up much of the variation.

The final two columns, (5) and (6), include vehicle-year as well as province-class and vehicle-province fixed effects, respectively. Vehicle-year fixed effects remove common shocks to particular vehicles across provinces over time. For example, vehicle advertising is typically national in scope and changes in vehicle design that sway preferences over unique vehicles are likely to affect all consumers in a similar way at a given time. Since changes in vehicle design may also result in a difference in feebate treatment over time, omitting this variable could potentially lead to bias in

[21]The effect is linear, so a -$1,000 feebate (i.e., a $1,000 subsidy) would be associated with an increase of similar magnitude.

[22]We separate provinces into four regions (BC, prairies, Ontario, and East) and produce fixed effects from these regional groups. The prairies group includes Alberta, Saskatchewan and Manitoba, and the East group includes Quebec and the four Atlantic provinces.

our results. The results in (5) suggest that a $1,000 feebate is associated with a 32 to 47 percent reduction in market share, depending on the resolution of fixed effects. (6) is the most flexible specification, controlling for a full set of vehicle-year and vehicle-province factors. The vehicle-province fixed effects account for province-specific preferences over particular vehicle models that are influenced by geography, road characteristics or climate. In total, this specification controls for over 13,000 fixed effects (about 5,500 vehicle-province fixed effects and about 7,600 vehicle-year fixed effects) in Panel A and about 38,000 fixed effects for the estimates in Panel C, which provides substantial confidence that our feebate parameter is uncounfounded. Ultimately, the inclusion of this rich set of fixed effects does not have a large effect on our estimate: a $1,000 feebate applied on a particular vehicle approximately causes a 30 percent decline in the market share of the vehicle. While our point estimates are nearly identical to the other specifications, the coefficients are not statistically significant. Increasing the sample size in the next subsections, we do obtain precisely estimated effects with FSA-level data and nested logit specifications.

Overall, our results are robust to a wide-range of specifications. Imposing a fee of $1,000 reduces the market share of a vehicle by approximately 30 percent relative to a counterfactual where no fee is applied. These estimates are in-line with the literature on vehicle taxes and subsidies. Chandra et al. (2010) analyses hybrid vehicle rebate programs in Canada and finds a nearly identical result: a $1,000 hybrid vehicle rebate has the effect of increasing market share by around 35 percent. Diamond (2009) conducts a similar impact in the US, and find that a $830 hybrid vehicle rebate results in an increase in hybrid market share of about 18 percent (for the Ford Escape model). Adamou et al. (2014) find that a €1,000 reduction in vehicle price is associated with a 20 percent increase in vehicle market share. The model of Berry et al. (1995) allows different elasticities based on observed and unobserved price elasticities of different vehicles. They find that a $1,000 increase in vehicle price is associated with a 9 to 125 percent reduction in market share, depending on the vehicle type. d'Haultfoeuille et al. (2014) find that a €1,000 rebate results in a 38 percent increase in vehicle market share.

3.4.1 FSA-level Results

Table 3 replicates column 6 in Table 2 with more disaggregated data. This table is based on vehicle market shares at the forward sortation level. At this level of disaggregation, two changes are required. First, we do not include gasoline prices as a covariate since we do not observe gasoline prices for individual FSAs. Next, the dependent variable is log market share instead of the difference in logged market shares of a vehicle and the outside option, because we do not observe population size in each year at the FSA-level.[23] The results are basically unchanged when we make similar modeling choices using provincial-level data.

[23]Section 6.1 shows that the feebate has minimal effect on total vehicle sales, such that this dependent variable should be an appropriate substitute to the one used in Table 2.

To a reasonable approximation, the results in Table 3 corroborate those in in Table 2. Even though gasoline prices are excluded from the model, feebates have a statistically significant and economically meaningful effect on unique vehicle market shares. In (1), a $1,000 fee leads to a 38 percent reduction in market share. (2) and (3) apply progressively finer fixed effect structures and show that a $1,000 levy reduces market share, respectively, by 39 and 38 percent. These FSA-level specifications closely match those in Table 2; the consistency of the feebate effect across distinct levels of data aggregation adds substantial reassurance to the conclusion that modest rates do influence behavior.

3.4.2 Nested Logit Results

As a final test of the robustness of the feebate elasticity, we return to the provincial data and relax the independence of irrelevant alternatives assumption by exploiting the data's 15 pre-defined vehicle segments (minivans, two-seaters, midsize cars, etc.). Table 4 presents estimates from a nested logit specification where each vehicle class is treated as a nest. The three columns correspond to the three fixed effects resolutions (similar to (6) in Table 2). Results from this table closely match the earlier coefficients and all coefficients are statistically significant at a 1 percent level. In column (1), a $1,000 feebate causes a 25 percent decline in market share. (2) and (3), using the different vehicle definition, have estimated decreases of 24 percent and 30 percent, respectively.

Over a range of specifications, the feebate parameter consistently shows that a $1,000 increase in the fees lead to a reduction in unique vehicle market share equal to approximately 30 percent. Importantly, this result is robust to a wide assortment of fixed effects, distinct substitution structures and alternative levels of aggregation. Ontario's modest feebates and those proposed in other jurisdictions do appear to alter the mix of the new vehicle fleet. We next establish an optimal feebate benchmark and evaluate Ontario's program vis-á-vis this standard and a no policy option.

4 Theoretical Foundation for Welfare Analysis

Consumers respond to modest feebates but it is unclear whether modest feebates are optimal. The next step involves embedding a feebate program in a consumer vehicle choice model and deriving expressions for optimal feebates that account for the demand responses associated with the policy.[24]

[24] In the current analysis, we assume manufacturers do not respond to the feebate program. While this may be restricitve for a Californian or national US feebate program, it is reasonable for many interesting cases for four reasons. First, most feebate programs are directly targeted at consumers. As such, we focus on the same decision as policy-makers. Second, many feebate policies that have been proposed or enacted are in small jurisdictions such as Ontario, Maryland or Oregon. Any potential gain or loss in market share in these jurisdictions likley provides an insufficient incentive to undertake the costs of vehicle redesign. Third, conventional feebates rates are modest and manufacturers should already be optimizing vehicle design to maximize fuel economy conditional on model attributes, price and competitors' options. All else constant, consumers will purchase a more fuel efficient vehicle as it will have lower lifetime operating cost. This implies that any potential gains from redesign are second- or third-order. Finally, assuming negligible manufacturer response is also highly plausible for our empirical application and the model can

This section demonstrates that an optimally designed feebate program can approach the outcome of new vehicle (indirect) Pigouvian taxation depending on the magnitude of the extensive, purchase margin. For expositional reasons, we maintain several assumptions that better align this analysis with the Ontario policy: we restrict attention to uniformly mixed pollutants such as CO_2 and use familiar logit choice structures in an effort to emphasize the underlying intuition of the welfare analysis. Appendix C extends and generalizes the in-text derivations.

4.1 Optimal Feebates

Policy-makers interested in designing new vehicle feebate programs must be cognizant of two margins. First, consumers make purchase or participation decisions: they choose whether to purchase a new vehicle and which model to select. Second, drivers choose whether to undertake a marginal trip. Feebates have the largest impact on model selection along the former margin, yet it is the latter decision that generates most vehicle-related externalities such as congestion, accidents and greenhouse gas emissions. Because new vehicle feebates primarily influence purchase decisions, they are second-best policies. A first-best policy adjusts the price of driving the marginal kilometer, forcing drivers to internalize all costs along the consumption margin. Feebates persist as a politically palatable environmental policy due to the ability to select particular program parameters. In particular, revenue neutral feebates programs are viewed as self-contained policies which can potentially improve social welfare without the perceived political challenges associated with first-best approaches. To date however, scrutiny of feebate programs has been based on *ad hoc* proposals. We derive analytically-tractable expressions for new vehicle feebate schedules based consumer demand responses which allow us to provide a benchmark for welfare analysis.

Feebate programs are usually comprised of two components: a pivot point (the fuel consumption rating at which there are no fees or rebates) and rate schedule (i.e., a system of fees and rebates) (Greene et al., 2005). Each of these components has a distinct role in offsetting the external costs of driving. Choice of feebate rates influences marginal utilities at the point when a consumer decides to purchase a new vehicle. In particular, feebates alter prices such that there is a greater incentive to purchase a relatively more fuel efficient vehicle. We restrict attention to revenue neutral feebates as these are the most frequently proposed policies and they represent one of the most important features of these proposals.

We present a model that includes consumers whose utility depends on vehicle choices and the external costs of driving and a social planner who selects the feebate parameters. The consumer's problem is specified using the standard random utility, differentiated products framework (e.g., Berry et al. (1995), Train and Winston (2007)) and builds on De Borger (2000, 2001). This approach directly relates to the empirical models in the previous section. The planner maximizes

readily be adapted to include manufacturer response. In sum, we believe that using the optimal feebate schedule based solely on consumer response as a benchmark for our welfare analysis is realistic and appropriate.

welfare subject to a revenue constraint. Throughout, we focus on greenhouse gas emissions as the primary externality. Curtailing vehicular greenhouse gas emissions is an objective of many feebate programs (including Ontario) and this simplifies the derivation. All expressions can be generalized to include other externalities. Optimal point-of-purchase feebate rates and pivot points then are a function of vehicle k's market share, the cost of the emissions externality and vehicle k's fuel efficiency relative to the average fuel efficiency of the new vehicle fleet.[25] The optimal revenue neutral feebate expression can be written as a function of the new vehicle Pigouvian tax, but is always of smaller magnitude.[26]

We provide graphical intuition for the derivation underlying optimal feebates next and then discuss the consumer problem, introduce expressions for the greenhouse gas externality and then present the planner's problem.

4.2 Graphical Intuition for Optimal Feebates

A simple two vehicle depiction conveys the intuition underlying optimal new vehicle feebate schemes. Appendix C presents the continuous formal model which motivates this graph. The next sub-section works through an optimal feebate derivation with a discrete choice structure, the most frequently used empirical methodology for modeling vehicle purchase behavior.

Figure 4 depicts a vehicle market with no outside option – i.e., there is no extensive margin decision and consumers must choose one of two models. The two panels in Figure 4 are identical except for the shaded areas that reflect the gain in social welfare due to the feebate program. The horizontal axis reflects the size of the vehicle market which is the sum of sales from a relatively fuel inefficient vehicle – vehicle d (dirty) – whose share is read from left to right, and a relatively fuel efficient vehicle – vehicle c (clean) – whose share is read from right to left. Consumers' marginal private benefits (MPB) for each vehicle are shown by the downward sloping dotted lines. (Marginal costs are normalized to zero.) Without any intervention, the market shares of the vehicles are determined by the intersection of these demand curves – shown as the private equilibrium.

Vehicles d and c generate external costs that are disproportional for a given level of usage: for the same number of kilometers traveled, vehicle d produces more pollution than vehicle c. The solid downward sloping schedules in Figure 4 show these marginal social benefits (MSB) from the purchase of either vehicle. Due to the external costs, the marginal social benefit curves sit below the marginal private benefit lines and, as d produces more pollution than c, its downward shift is larger. The socially optimal allocation is shown at the social optimum. This equilibrium demonstrates the mechanism through which feebates operate: feebate programs alter the mix of vehicles in the market y changing relative prices. The socially efficient equilibrium involves a reduction of vehicle d's share

[25] The role of market share depends on the choice structure applied. For instance, it drops out of the standard logit formulation due to the independence of irrelevant alternatives condition.

[26] Feebates can also be designed to replicate any binding CAFE standard (Klier and Linn, 2012; Gillingham, 2013); although the reverse is not necessarily true.

and an increase in vehicle c's share as illustrated by comparing the vehicle purchases underneath the horizontal axes of Panels A (private equilibrium) and B (socially optimal equilibrium), respectively.

The shaded area in Panel A depicts the social welfare at the private equilibrium. Panel B then shows two shaded areas. First, the dark triangle is the *gain* in total surplus from a move from the private equilibrium to the socially optimal outcome. This improvement in social welfare is achieved by varying the relative prices of vehicles d and c and, thereby, inducing marginal consumers to purchase the relatively fuel efficient vehicle rather than the relatively fuel inefficient vehicle. Changing the mix of vehicles – reducing the share of the fuel inefficient vehicle and increasing the share of the fuel efficient vehicles purchased – leads to less pollution for a given level of usage which in turn improves welfare. Next, the lightly shaded areas $ABCD$ and $DEFG$ are of equal size and demonstrate the transfer of surplus within the market that results from the self-financing constraint.

Panel B and the expressions in Appendix C also enable us to compare new vehicle Pigouvian taxation (i.e., a tax that applies to vehicle purchase rather than use in proportion to vehicles' lifetime emissions) to feebate programs. The socially optimal equilibrium is defined by relative prices and is identical for both the Pigouvian tax and feebate scenarios as long as extensive margin adjustments are restricted: feebate policies generate the same outcome as Pigouvian taxes. Similar to the optimal feebate, new vehicle Pigouvian taxes are heterogeneous across vehicles (Knittel and Sandler, 2013; Starrett, 1972). We focus on vehicle d, the relatively fuel inefficient option, but symmetrical arguments apply to the clean vehicle.

In Figure 4, the dollar-valued tax on the relatively fuel inefficient vehicle, vehicle d, equals $t_d^P = MPB_d - MSB_d$. Appendix C shows that the optimal new vehicle tax rate corresponds to $t_d^P = \frac{\delta}{\lambda}\frac{e_d}{p_d}$ where δ is the per unit cost of pollution (e.g., \$tCO_2), λ is the multiplier on the government's revenue constraint, e_d is per vehicle lifetime emissions and p_d is total vehicle cost. Revenue neutral feebates, in contrast, impose a self-financing constraint so that fees and rebates must be redistributed *within the market* based on differentiated fuel consumption ratings. The optimal feebate in this two vehicle market, as shown in Appendix C, can be expressed as a function of the optimal Pigouvian taxes: $F_d = s_c\big(t_d^P - t_c^P\big)$, where $t_d^P > t_c^P > 0$ and $t_d^P > F_d > 0 > F_c$. The optimal feebate schedule is a market share weighted vertical shift of the difference between the indirect Pigouvian tax schedule on the relatively fuel inefficient and efficient vehicles (see Appendix for additional discussion). This revenue neutral rate generates the same equilibrium as the Pigouvian tax.

Beyond the mechanical elements of this correspondence, several additional comments are warranted. There are advantages that support a feebate policy versus a Pigouvian tax for new vehicle: the planner is able to achieve the same equilibrium mix of vehicle by proposing lower per vehicle levies with a feebate than the per vehicle Pigouvian taxes, a result that may smooth potential opposition to program implementation. In real world policy conversations, the absolute size of taxes, fees and subsidies are often important. Second, the distribution of surplus differs between the Pigouvian tax and the feebate. With the tax, the government receives revenue that can be used for other purposes (e.g., cross-base revenue neutrality is possible as is the prospect of a double dividend). Feebates, as illustrated in the figure, redistribute surplus within the market from purchasers of relatively fuel inefficient vehicles to purchasers of relatively fuel efficient vehicles. The politics of this are less clear and may be beneficial or harmful. Finally Figure 4 makes obvious that the welfare gains from feebates emerge from a change in the mix of vehicles that are differentiated across fuel consumption ratings. Neither a new vehicle Pigouvian tax nor a feebate, in this simple context, reduces the number of vehicles purchased – hence both policies are second best. Pigouvian taxes may be expected to yield greater reductions in total pollution via the extensive margin channel, whereas sections 5 and 6.1 shows that feebates have ambiguous effects along the vehicle purchase margin.

We next move from this general model and derive optimal feebate expressions using the conventional discrete good methodology for vehicle purchase decisions, an approach that corresponds to our empirical results and permits welfare analysis.

4.2.1 Consumer's Problem

Consumer i in region j at time t can choose vehicle type k or an outside good, denoted by the subscript 0. The outside good implies that the consumer does not purchase a new vehicle. Utility is comprised of explained, V_{ijkt}, and random, ϵ_{ijkt}, components (i.e., $U_{ijkt} = V_{ijkt} + \epsilon_{ijkt}$). The systematic component of utility (V_{ijkt}) for vehicle model k is defined as:

$$V_k \equiv \{\max_{x_k,z} U(x_k, \theta_k, z, E) \mid p_k + p_g f_k x_k + F_k + z = y\} \tag{3}$$

Utility depends on x_k, the expected kilometers driven in vehicle k, a vector of characteristics describing the vehicle, θ_k, a numeraire good, z, whose price is normalized to one and E, the total level of automobile emissions generated by all drivers. p_k is vehicle k's price which is interpreted as the sum of purchase price and maintenance costs. Operating costs equal $p_g f_k x_k$, where p_g is the price of gasoline and f_k is the fuel efficiency of vehicle k in liters per 100 kilometers. The feebate rate which applies a tax or subsidy on the price of new vehicles according to their fuel efficiency rating is given by F_k. Consumer i's income is y.

Any individual driver has a positive, but negligible, effect on total vehicle-related emissions and as such individuals do not consider their contribution to total emissions when making vehicle purchases. Hence, we decompose V_k as:

$$V_k = \tilde{V}_k - \delta E \tag{4}$$

where δ is the (constant) marginal disutility of emissions which is common to all drivers and the $\sim$ over V indicates the consumer's private decision problem. A first-best policy applies to vehicle use not its purchase and would set marginal cost of driving equal to an expression that depended on δ.

The probability that consumer i selects any given vehicle k over vehicle l is: $P_{ikpt} = \Pr(U_{ijkt} > U_{ijlt} = \Pr(V_{ijkt} - V_{ijlt} > \epsilon_{ijlt} - \epsilon_{ijkt})$. Assuming the ϵ_{ijkt} are independently and identically distributed extreme value, we write the familiar multinomial logit expression which describes the probability that consumer i in region j chooses vehicle k:

$$\pi_{ijk} = \frac{\exp V_{ijk}}{\sum_{k=0}^{K} \exp V_{ijk}} = \frac{\exp \tilde{V}_{ijk}}{\sum_{k=0}^{K} \exp \tilde{V}_{ijk}} \tag{5}$$

Equations (4) and (5) illustrate that, while new vehicle purchases are influenced by the feebate rate, consumer decisions are independent of the aggregate level of the externality.

4.2.2 Externality

Even though feebates apply to new vehicle purchases, the decision to purchase a vehicle does not generate any driving-related externalities. Vehicle-related external costs are from consumers' usage decisions. Conditional on owning vehicle k then, demand for kilometers traveled is derived from V_k via Roy's identity and equals $x_k(p_g, F_k)$. Total emissions then summed across expected vehicle purchases (including the outside option) and all consumers is:[27]

$$E(p_g, \boldsymbol{F}) = \phi \sum_{i=1}^{N} \sum_{k=0}^{K} \pi_{ik}(p_g, \boldsymbol{F}) x_{ik}(p_g, F_k) f_k \tag{6}$$

where ϕ represents emissions per unit of fuel consumed.[28]

Both the probability of purchasing vehicle k and potentially the number of kilometers driven, x_{ik}, depend the the feebate schedule. Partial differentiating (6) with respect to F_l gives:

$$\frac{\partial E}{\partial F_l} = \phi \sum_{i=1}^{N} \left[\sum_{k=0}^{K} \left(\frac{\partial \pi_{ik}}{\partial F_l} x_{ik} f_k \right) + \pi_{il} \frac{\partial x_{il}}{\partial F_l} f_l \right] \tag{7}$$

A change in the feebate on vehicle l affects greenhouse gas emissions through two channels.[29]

[27] Region-specific indices, the js, are suppressed.

[28] Additional flexibility could be added to this formulation by relaxing the uniformly mixed pollutant assumption. This would involve adding a subscript to ϕ and considering hetreogeneous costs.

[29] Our focus is greenhouse gases, so this expression invokes (4). However, if the externality were non-separable from utility (i.e., individual decisions had a meaningful effect on the magnitude of the externality) as would be the case when assessing congestion costs, then (7) would equal: $\frac{dE}{dF_l} = \frac{\phi \sum_{i=1}^{N} \sum_{k=0}^{K} \left(\left(\frac{\partial \pi_{ik}}{\partial F_l} + \frac{\partial \pi_{ik}}{\partial F_k} \frac{\partial F_k}{\partial F_l} \right) x_{ik} f_k + \pi_{ik} \frac{\partial x_{ik}}{\partial F_k} \frac{\partial F_k}{\partial F_l} f_k \right)}{1 - \phi \sum_{i=1}^{N} \sum_{k=0}^{K} \left(\left(\frac{\partial \pi_{ik}}{\partial E} \right) x_{ik} f_k + \pi_{ik} \frac{\partial x_{ik}}{\partial E} f_k \right)}$.

The first expression after the summation operators describes how the composition of new vehicle purchases changes as a function of the feebate rate on vehicle l and how this mix influences emissions. It quantifies the probability that a consumer will purchase a given vehicle as a result of a change in the feebate rate and then how this decision influences the expected fuel economy of all new vehicle purchases. To a large degree, it is through this effect on the composition of new vehicle stock that the feebate has the greatest impact. The second channel, the final term in (7), reflects how kilometers traveled, conditional on owning vehicle k, changes due to the feebate rate. This second term is likely small – in fact, the impact on kilometers traveled may be zero or even positive reflecting a rebound effect.

Overall (7) denotes the change in the emissions as a consequence of changes to the feebate schedule. We refer to this expression as the "marginal feebate effect" on emissions for vehicle l. As written, (7) is not intuitive. By making several assumptions, it is possible to write this marginal feebate effect, so the relationship between feebates and emissions is clearer. Assume that (i) consumer i would drive any new vehicle k an identical number of kilometers, x_i, (ii) there is no change in kilometers traveled conditional on owning vehicle l (i.e., the second term in (7) equals zero (no rebound effect), a plausible assumption (Gillingham et al., 2013)), and (iii) the program is revenue neutral. (7) then reduces to:

$$\frac{\partial E}{\partial F_l} = \phi \sum_{i=1}^{N} \left[\pi_{il} \beta x_i \left(f_l - \bar{f} \right) \right] \tag{8}$$

where $\beta = \frac{\partial V_l}{\partial F_l} < 0$ is the marginal effect on utility from imposing a feebate on vehicle l and $\bar{f}$ is average fuel efficiency of the entire new vehicle fleet. This expression states that *for a fixed number of kilometers traveled* the change in emissions from a feebate on vehicle l equals the probability that consumer i chooses l multiplied by the marginal effect of the feebate and on vehicle l's fuel economy relative to the mean fuel economy. If vehicle l is more fuel inefficient than average, then the term in parentheses is positive; however, the full expression is negative as β is negative. Thus a fee levied on a fuel inefficient vehicle reduces the overall level of emissions.

4.2.3 Social Welfare

A planner whose objective is to maximize total societal welfare must select a schedule of feebate rates subject to a revenue neutrality constraint. Using utilitarian social welfare, the planner's problem takes the form (Small and Rosen, 1981):

$$\max_{F} \sum_{i=1}^{N} \ln \left[\exp V_{i0} + \sum_{k=1}^{K} \exp V_{ik} \right]$$

Likewise, the assumption of no manufacturer response to the feebate policy implies $\frac{\partial f_k}{\partial F_k} \frac{\partial F_k}{\partial F_l} = 0$.

subject to:

$$\sum_{i=1}^{N}\sum_{k=1}^{K}\pi_{ik}F_k \geq R$$

where R, revenue, equals zero for a revenue neutral program.

Assuming homogeneous consumers and normalizing N to 1, we derive an intuitively-appealing expression for the optimal feebate. A social planner looking to maximize social welfare should set the feebate rate on vehicle l according to:

$$F_l = \frac{\delta}{\lambda}\phi x(f_l - \bar{f}) \tag{9}$$

where λ is the Lagrangian multiplier on the revenue constraint and δ is the marginal damage of emissions. This expression demonstrates that the optimal feebate rate and pivot point are jointly determined under self-financing. The pivot point is determined by the attributes of the fleet – fuel economies and relative market shares – while external costs regulate the slope of the schedule. Further, the logit formulation of the problem is a special case that restricts substitution possibilities a feature that is relaxed in (22) in Appendix C. The next section uses (9) to quantify policy outcomes.

5 Counterfactual Simulations

The welfare implications of Ontario's feebate program are calculated vis-à-vis a no feebate baseline and an optimal schedule. As a first step however, we can qualitatively compare the Ontario program to the optimal revenue neutral policy. Figure 5 depicts the Ontario and optimal rate schedules.[30] Four differences between the policies are apparent. First, Ontario's fee and rebate schedule is coarse with 14 distinct levels across two classes of vehicles. This coarseness can be seen via the discrete jumps ("notches") in the Ontario schedules. An optimal schedule imposes a much finer set of rates with a unique fee or rebate for each model-fuel efficiency rating. In Figure 5, the optimal feebate appears continuous across fuel consumption ratings. Second, Ontario's program differentiates between cars and SUVs while omitting trucks and vans. The optimal schedule applies identically to all vehicles at a given fuel consumption rating irrespective of category. Third, Ontario's program was based on rated highway fuel economy, whereas the optimal program uses a weighted measure over all driving (city and highway). Finally, an optimal feebate program taxes very fuel inefficient vehicles at a much higher rate, while providing larger subsidies to fuel efficient vehicles as evident by the steeper slope of its schedule.

Next, the results from Column 3 in Table 4 are used to evaluate the effect of the program on

[30]The optimal schedule is calculated using year 2000 information.

vehicle sales and carbon dioxide emissions. To construct our counterfactual scenarios, we predict vehicle sales (i) in absence of any feebate and (ii) with an optimal schedule imposed and then compare these to the actual observations from Ontario's program.[31] To obtain simulated vehicle purchase behavior, we segment the sample by vehicle-location and simulate new market shares assuming the absolute size of the outside option remains fixed. We repeat this simulation for each market and then sum over all markets. As we have the population of new vehicle registrations, we are able to exhaustively partition the data (Train, 2003). For each vehicle, we assume a lifespan of 15 years and annual vehicle kilometers of 16,000. The product of total vehicle sales, rated fuel consumption, vehicle lifespan and annual kilometers gives predicted lifetime gasoline consumption by unique vehicle type. Gasoline consumption is multiplied by 2.4 kg CO_2/liter to obtain carbon dioxide emissions. We then sum gasoline consumption and CO_2 emissions over all vehicles to obtain the total lifetime consumption for all vehicles sold in a particular cohort.

Table 5 provides results from these policy simulations. Values are based on a single year's vehicle sales, using the 2010 vehicle cohort as our reference point. Actual vehicle sales in 2010 equaled 360,475 as shown on the Ontario feebate row. Under a no feebate scenario, sales remain nearly identical at 360,410, while they actually *increase* with an optimally designed revenue neutral program to 383,610. Across all unique vehicles, Figure 6 illustrates the change in the distribution of vehicle sales for the no feebate and optimal design scenarios relative to the Ontario policy. Only minor differences are visible between the no feebate Ontario scenarios. Conversely, more dramatic fleet-wide changes in the distribution of vehicles is apparent when contrasting the Ontario program to an optimal design. Sales of fuel inefficient vehicles – those with combined fuel consumption ratings that are greater than, say, 12 liters per 100 kilometers – are much greater under the no feebate and Ontario scenarios compared with the optimal feebate; yet, substantially more very fuel efficient, subsidized vehicles are sold with the optimal policy as shown on the left-hand side of the figure.

Table 5 next illustrates the lifetime emissions of this cohort under the three scenarios. Differences in emissions motivate this policy. The fleet of vehicles under the no feebate, Ontario and optimal programs produce 20.1MtCO_2, 20.3 MtCO_2 and 17.4 MtCO_2, respectively, implying that emissions under Ontario's feebate program *increased* relative to a no feebate situation. In 2010 however, the Ontario program collected $21M in revenue. Combining these emissions and revenue values we are able to calculate the gain in welfare of Ontario's and the optimal feebate versus a no feebate baseline. First, if the social cost of carbon equals $10 per tonne CO_2, Ontario's program yielded an increase in social welfare equal to $19.5M compared with a no feebate option. This is

[31] We emphasize that we are interested in calculating the welfare implications the Ontario feebate program, not in using structural parameters *per se*. Based on the formulation of the theoretical model, only a handful of variables are required to calculate welfare – i.e., we employ Chetty's (2008) sufficient statistic approach which enables us to calculate welfare based on "high-level" elasticities that can be precisely estimated and therefore do not need to fully estimate a full set of vehicle attribute elasticities. The main advantage of this approach is that we are able to precisely identify these feebate parameters, the sufficient statistic in our welfare calculations.

the sum of emissions reductions (valued at $10/t$CO_2$) and government revenue. Had Ontario's policy been optimally designed and revenue neutral, the welfare gain would have equaled $27.0M, an additional $7.5M.[32] Assuming larger social costs of carbon increases the appeal of the optimal revenue neutral policy, while decreasing the benefits of Ontario's program. If the social cost of carbon equals $30 per tonne CO_2, the gain in welfare for the Ontario and optimal feebates corresponds to $16.5M and $81.0M. An optimal designed would have generated an additional $64.5M in social welfare over the lifetime of the 2010 vehicle cohort.

Both the Ontario feebate program and an optimally designed revenue neutral program generate positive social welfare compared to a no feebate option. On a per vehicle basis, our assumptions imply that Ontario's program yielded an increase in welfare equal to $46 (at $30/t$CO_2$), while an optimal revenue neutral policy produced gains of $211 per vehicle. These values are likely sufficiently large to induce politically constrained policy-makers to seriously consider feebate programs as alternatives to Pigouvian taxation.

6 Potential Challenges to Optimal Policy Design

Heretofore the analysis has focused on the relative merits of an existing feebate program versus an optimally designed revenue neutral program. We demonstrated that optimal revenue neutral policy can produce gains over a sub-optimal revenue positive program. Many real world scenarios however do not match the idealized situation under which we posited the optimal program. Practical and political complications may confront optimally designed policies. This section empirically investigates three challenges to optimal policy implementation finding that two may only present minor difficulties for program design. This analysis has the added benefit of supplementing our earlier econometric results and considers: (i) the effect of Ontario's feebate on the extensive, vehicles purchased margin, (ii) potentially asymmetric responses to fees versus rebates and (iii) potential heterogeneity along an urban-rural dimension.

6.1 Extensive Margin

Feebates operate by altering the mix of vehicles in the fleet. Any program would expect ambiguous effects for the extensive, vehicle purchase decision compared with a new vehicle Pigouvian tax. Still, the implications of extensive margin adjustments may be important (and can be seen in the simulations in section 5). The analysis in section 3 is confined to a period when Ontario's feebate existed. As such, it is less amenable to evaluating whether the existence of a feebate program affected total vehicles sales. To estimate the impact of the feebate on aggregated new vehicle sales, a separate dataset spanning the introduction and removal of the feebate is required. We use

[32]The social cost of carbon that equates the welfares of the revenue positive Ontario program and the revenue neutral optimal policy is $7.37.

monthly data from Statistics Canada, on aggregate provincial passenger car sales from 1980 to 2012 and estimate the following model:[33]

$$y_{it} = \alpha \tilde{F}_{it} + \beta X_{it} + \delta_i + \gamma_t + \epsilon_{it}$$

where y_{it} is log per capita passenger car sales, $\tilde{F}_{it}$ is a dummy variable capturing the presence of a feebate program, X_{it} are time-variant provincial characteristics including log real per capita income, log unemployment rate and log real gasoline price, δ_i is a province fixed effect and γ_t is a time fixed effect. To estimate the model, γ_t is decomposed such that $\gamma_t = \gamma_y + \gamma_m$, where γ_y is a year dummy and γ_m is a month dummy.

Results are presented in Table 6. (1), the most parsimonious model, suggests that the feebate program yielded a statistically insignificant decrease in sales of 7 percent. Including per capita GDP in (2) attenuates this estimate to a 2 percent decrease. Including additional covariates in (3) and (4) actually flips the sign of the feebate coefficient, suggesting the vehicle sales increased as a result of the feebate program. Lack of robustness and statistical insignificance across these four specifications leads us to conclude that Ontario's feebate program had minimal implications for total vehicle sales in the province. Altering the fleet's composition – encouraging the purchase of more fuel efficient vehicles – appears to be the mechanism driving emissions reductions.

6.2 Asymmetric Responses

The feebate response parameters estimated in section 3 assume that consumers respond to fees and rebates symmetrically – a subsidy is simply a negative tax. Recent evidence suggests that the characteristics of the policy intervention may influence household decisions (e.g., Chetty et al. (2008); Gallagher and Muehlegger (2011)). In the feebate context, consumers may react differently when faced with a rebate or a fee. A plausible mechanism actually predicts asymmetric responses to taxes and subsidies: during an initial vehicle test drive, dealers are more likely to emphasize subsidies associated with a vehicle's fuel efficiency. Similarly, they may attempt to lump fees resulting from a vehicle's relative fuel inefficiency with other administrative and freight costs. As such, we may expect that consumers have differential response to fees and rebates.

Table 7 presents a model using the FSA data where each feebate rate is separately binned. (1) replicates model (1) of Table 3 and illustrates that consumers are substantially more sensitive to rebates relative to fees (recall: rebates enter as negatives).[34] Consumers' asymmetric responses to

[33]Data are from Table 0079-0003. Statistics Canada divides the new vehicle market into passenger car and truck segments. The truck segment includes minivans, sport utility vehicles, pickup trucks, buses, heavy trucks and other heavy-duty road vehicles. Because the truck segment mixes passenger and freight vehicles, our regressions use data only for passenger cars. We obtain similar results in regressions using total vehicles sales or trucks as the dependent variable, further supporting the conclusion that the feebate had a minimal effect on total vehicle sales.

[34]Pairwise F-tests of the rebate vis-à-vis each of the fee categories confirm that the difference is statistically significant.

taxes and subsidies pose a challenge for policy-makers: it is more difficult to design policy that approximates the optimal benchmark when decisions are a function of the policy's characteristics. Estimated coefficients show a non-linear and decreasing responsiveness to increasing fees. Consumers are less responsive, on the margin, to a \$1600 fee than a \$250 fee. Of course, the logit framework imposes specific vehicle substitution possibilities, particularly with respect to increasing prices; nonetheless, both the rebate-fee asymmetry and the decreasing responsiveness to increasing fees can pose notable challenges for optimal program design.

6.3 Heterogeneity

Consumers are not homogeneous and distinct groups may attempt to influence the politics of a program design process if they believe that they are disproportionately affected. While rebate-fee asymmetry poses one challenge to the optimal design of feebate schemes, consumer heterogeneity poses another. Rural residents, for instance, may be more likely to be employed in agriculture, construction and resource industries. As such, larger, relatively fuel inefficient vehicles may be necessary for work. The burden of high fees on these vehicles would then fall on a group that has less ability to respond to the policy's incentives. Politicians may be receptive to constituents' concerns over this unequal burden.

Table 8 segments the sample into rural and urban FSAs[35] and demonstrates that rural households are actually more sensitive to feebates compared with rural residents. Table 8 estimates a model for all vehicles in (1) and cars only in (2). Two parameters are shown for each formulation. First, the feebate coefficient is statistically significant and matches the earlier results, equaling -0.35 and -0.36 in columns (1) and (2) respectively. Next, an interaction term between the feebate and a rural dummy variable is included. This parameter equals a statistically significant -0.21 in (1) and -0.20 when only cars are considered in (2). Rather than being less sensitive to the feebate schedule, rural residents are more responsive to fees and rebates. Based on these results, it is not obvious that consumer heterogeneity due to a rural-urban divide presents a major challenge for effective policy-making.

7 Conclusion

New vehicle feebate programs have been proposed by several jurisdictions. The properties of this class of programs are potentially appealing to politically constrained governments. Previous literature has only evaluated proposed feebate programs from an *ex ante* perspective; no prior study has empirically evaluated the implications of an existing program nor considered characteristics of an optimally design feebate policy. This paper addresses this gap by studying Ontario's long-running

[35] Assignment of FSAs into "rural" and "urban" was done according to Canada Post's definition of rural and urban mail delivery routes.

Tax and Credit for Fuel Conservation and by deriving expressions for optimal program parameters using a familiar discrete, vehicle choice framework. We also clarify terminology related to revenue neutral and self-financing policy design.

Using a dataset comprised of the population of Canadian vehicle registrations over an 11 year span, we find that, compared to a no feebate option, Ontario's feebate was welfare-enhancing but less so than an optimally designed revenue neutral program. Our calculations suggest that Ontario's policy actually increased lifetime vehicle CO_2 emissions by about 0.2 million tonnes per vehicle cohort. An optimal revenue neutral formulation, in contrast, would have reduced emissions by 2.7 million tCO_2 relative to the no feebate scenario. We further demonstrate that, although Ontario's program did not appear to affect total vehicle sales, important practical complications may impede policy-makers' efforts to implement an optimal program design – in particular, we highlight that consumers have asymmetric responses to fees and rebates and that rural households are more sensitive to feebate schedules than their urban counterparts.

Our theoretical and empirical results suggest that new vehicle feebates may be a feasible policy option to confront the external costs associated with driving. Effective policy design ensures that these policies provide the greatest benefit at the lowest social cost. Throughout we emphasize that feebates are second-best policies that may be appealing to politically constrained governments. Obviously first-best programs, policies that directly influence vehicle usage decisions, are preferred; yet, in the real world of policy-making and program design, substantial merit exists in exploring the properties and prospective outcomes of a class of *realistic* programs. Revenue neutral feebates seem to belong to this class. Insofar as they remain a viable policy option for decision-makers, it is worthwhile advocating for well-conceived programs – especially as even modest feebates affect behavior and are welfare-enhancing.

References

Adamou, A., S. Clerides, and T. Zachariadis (2011). Environmental and economic effects of co2-based automobile taxes in germany. Technical report, Working paper.

Adamou, A., S. Clerides, and T. Zachariadis (2014). Welfare implications of automobile feebates: A simulation analysis. *The Economic Journal*.

Anderson, M. and M. Auffhammer (2011). Pounds that kill: The external costs of vehicle weight. Technical report, National Bureau of Economic Research.

Austin, D. and T. Dinan (2005). Clearing the air: The costs and consequences of higher cafe standards and increased gasoline taxes. *Journal of Environmental Economics and Management 50*(3), 562–582.

Banerjee, R. (2009). *Deals on Wheels*. CD Howe Institute.

Berry, S. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242–262.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, 841–890.

Bovenberg, A. L. and R. A. de Mooij (1994). Environmental levies and distortionary taxation. *The American Economic Review 84*(4), 1085–1089.

Bregha, F. and J. Moffet (1995). The Tax for Fuel Conservation in Ontario. In R. J. Gale, S. Barg, and A. Gillies (Eds.), *Green Budget Reform: An International Casebook of Leading Practices*, Chapter 4, pp. 68–78. Earthscan and International Institute for Sustainable Development.

Busse, M. R., C. R. Knittel, and F. Zettelmeyer (2013). Are consumers myopic? evidence from new and used car purchases. *The American Economic Review 103*(1), 220–256.

Carneiro, A., P. Guimarães, and P. Portugal (2012). Real wages and the business cycle: Accounting for worker, firm, and job title heterogeneity. *American Economic Journal: Macroeconomics 4*(2), 133–152.

Chandra, A., S. Gulati, and M. Kandlikar (2010). Green drivers or free riders? an analysis of tax rebates for hybrid vehicles. *Journal of Environmental Economics and Management 60*(2), 78–93.

Chetty, R. (2008). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. Technical report, National Bureau of Economic Research.

Chetty, R., A. Looney, and K. Kroft (2008). nsalience and taxation: Theory and evidence, oamerican economic review, forthcoming.

Davidson, R. and J. G. MacKinnon (1993). Estimation and inference in econometrics. *OUP Catalogue*.

De Borger, B. (2000). Optimal two-part tariffs in a model of discrete choice. *Journal of Public Economics 76*(1), 127–150.

De Borger, B. (2001). Discrete choice models and optimal two-part tariffs in the presence of externalities: optimal taxation of cars. *Regional Science and Urban Economics 31*(4), 471–504.

d'Haultfoeuille, X., P. Givord, and X. Boutin (2014). The environmental effect of green taxation: the case of the french bonus/malus. *The Economic Journal*.

Diamond, D. (2009). The impact of government incentives for hybrid-electric vehicles: Evidence from us states. *Energy Policy 37*(3), 972–983.

Fischer, C., W. Harrington, and I. W. Parry (2007). Should automobile fuel economy standards be tightened? *The Energy Journal*, 1–29.

Frank, R. H. (2006). A way to cut fuel consumption that everyone likes, except the politicians. *New York Times*.

Fullerton, D. and A. Wolverton (2000). Two generalizations of a depost-refund system. *American Economic Review 90*(2), 238–242.

Gallagher, K. S. and E. Muehlegger (2011). Giving green to get green? incentives and consumer adoption of hybrid vehicle technology. *Journal of Environmental Economics and Management 61*(1), 1–15.

Gillingham, K. (2013). The economics of fuel economy standards versus feebates.

Gillingham, K., M. J. Kotchen, D. S. Rapson, and G. Wagner (2013). Energy policy: The rebound effect is overplayed. *Nature 493*(7433), 475–476.

Goulder, L. H. (1995). Environmental taxation and the double dividend: a reader's guide. *International Tax and Public Finance 2*(2), 157–183.

Government of Ontario (1989). 1989 Ontario Budget. Technical report, Ministry of Treasury and Economics.

Government of Ontario (1991). 1991 Ontario Budget. Technical report, Ministry of Treasury and Economics.

Government of Ontario (2010). Tax for fuel conservation.

Greene, D., P. Patterson, M. Singh, and J. Li (2005). Feebates, rebates and gas-guzzler taxes: a study of incentives for increased fuel economy. *Energy Policy 33*(6), 757–775.

Greene, D. L. (1991). Short-run pricing strategies to increase corporate average fuel economy. *Economic Inquiry 29*(1), 101–114.

Greene, D. L. (2009). Feebates, footprints and highway safety. *Transportation research part D: transport and environment 14*(6), 375–384.

Guimaraes, P. and P. Portugal (2010). A simple feasible procedure to fit models with high-dimensional fixed effects. *Stata Journal 10*(4), 628.

Harrison, K. (2012). A tale of two taxes: The fate of environmental tax reform in canada. *Review of Policy Research 29*(3), 383–407.

Johnson, K. (2006). Feebates: An effective regulatory instrument for cost-constrained environmental policy. *Energy policy 34*(18), 3965–3976.

Karplus, V. J., S. Paltsev, M. Babiker, and J. M. Reilly (2013). Should a vehicle fuel economy standard be combined with an economy-wide greenhouse gas emissions constraint? implications for energy and climate policy in the united states. *Energy Economics 36*(0), 322 – 333.

Kleit, A. N. (2004). Impacts of long-range increases in the fuel economy (cafe) standard. *Economic Inquiry 42*(2), 279–294.

Klier, T. and J. Linn (2010). The price of gasoline and new vehicle fuel economy: evidence from monthly sales data. *American Economic Journal: Economic Policy 2*(3), 134–153.

Klier, T. and J. Linn (2012). Using vehicle taxes to reduce carbon dioxide emissions rates of new passenger vehicles: evidence from france, germany, and sweden. Technical report, Working Paper, Federal Reserve Bank of Chicago.

Knittel, C. R. and R. Sandler (2013). The welfare impact of indirect pigouvian taxation: Evidence from transportation. Technical report, National Bureau of Economic Research.

Langer, A. and N. H. Miller (2012). Automakers' short-run responses to changing gasoline prices. *Review of Economics and Statistics* (0).

Li, S., C. Timmins, and R. H. von Haefen (2009). How do gasoline prices affect fleet fuel economy? *American Economic Journal: Economic Policy 1*(2), 113–37.

Mrozek, J. R. (2000). Revenue neutral deposit/refund systems. *Environmental and Resource Economics 17*(2), 183–193.

Nevo, A. (2000). A practitioner's guide to estimation of random-coefficients logit models of demand. *Journal of Economics & Management Strategy 9*(4), 513–548.

Oster, E. (2013). Unobservable selection and coefficient stability: Theory and validation. Technical report, National Bureau of Economic Research.

Palmer, K. and M. Walls (1997). Optimal policies for solid waste disposal taxes, subsidies, and standards. *Journal of Public Economics 65*(2), 193–205.

Parry, I. W. (1995). Pollution taxes and revenue recycling. *Journal of Environmental Economics and management 29*(3), S64–S77.

Parry, I. W. and K. A. Small (2005). Does britain or the united states have the right gasoline tax? *The American Economic Review 95*(4), 1276–1289.

Rivers, N. and B. Schaufele (2013). Salience of carbon taxes in the gasoline market. *Available at SSRN*.

Roth, K. (2012). The unintended consequences of uncoordinated regulation: Evidence from the transportation sector. *Unpublished working paper*.

Sallee, J. M. and J. Slemrod (2010). Car notches: strategic automaker responses to fuel economy policy. Technical report, National Bureau of Economic Research.

Small, K. A. and H. S. Rosen (1981). Applied welfare economics with discrete choice models. *Econometrica: Journal of the Econometric Society*, 105–130.

Small, K. A. and K. Van Dender (2007). Fuel efficiency and motor vehicle travel: the declining rebound effect. *The Energy Journal*, 25–51.

Starrett, D. A. (1972). Fundamental nonconvexities in the theory of externalities. *Journal of Economic Theory 4*(2), 180–199.

Thorpe, S. G. (1997). Fuel economy standards, new vehicle sales, and average fuel efficiency. *Journal of Regulatory Economics 11*(3), 311–326.

Train, K. (2003). *Discrete choice methods with simulation.* Cambridge university press.

Train, K., W. Davis, and M. Levine (1997). Fees and rebates on new vehicles: Impacts on fuel efficiency, carbon dioxide emissions, and consumer surplus. *Transportation Research Part E: Logistics and Transportation Review 33*(1), 1–13.

Train, K. E. and C. Winston (2007). Vehicle choice behavior and the declining market share of us automakers*. *International Economic Review 48*(4), 1469–1496.

Walls, M. (2011). Deposit-refund systems in practice and theory. *Resources for the Future Discussion Paper* (11-47).

8 Figures

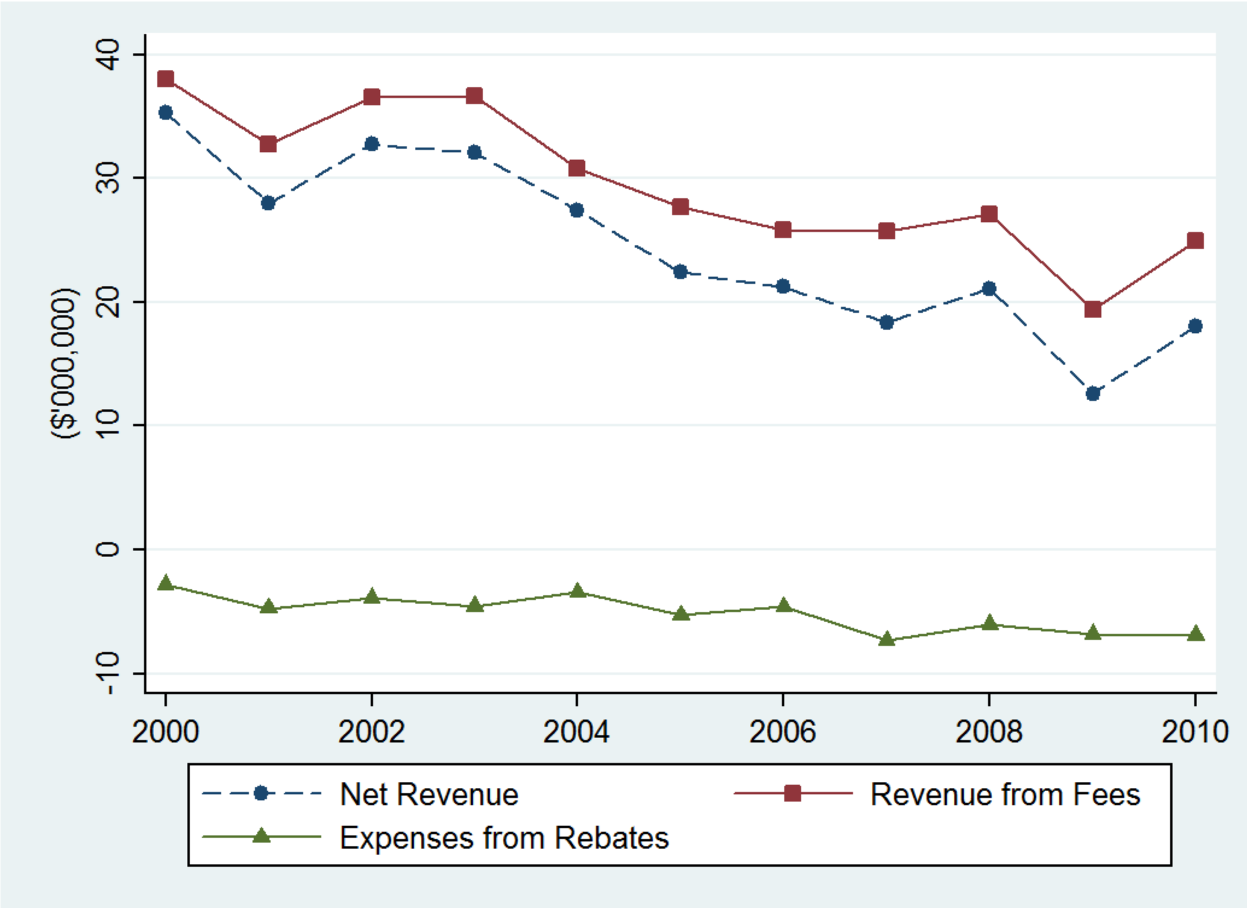

Figure 1: Revenue and expenditures associated with the Ontario Tax for Fuel Conservation and the Tax Credit for Fuel Conservation.

All values are in real terms ($2002) via the all-items national consumer price index (Statistics Canada 326-0021) and calculated by the authors using data from Desrosiers Automotive Consultants.

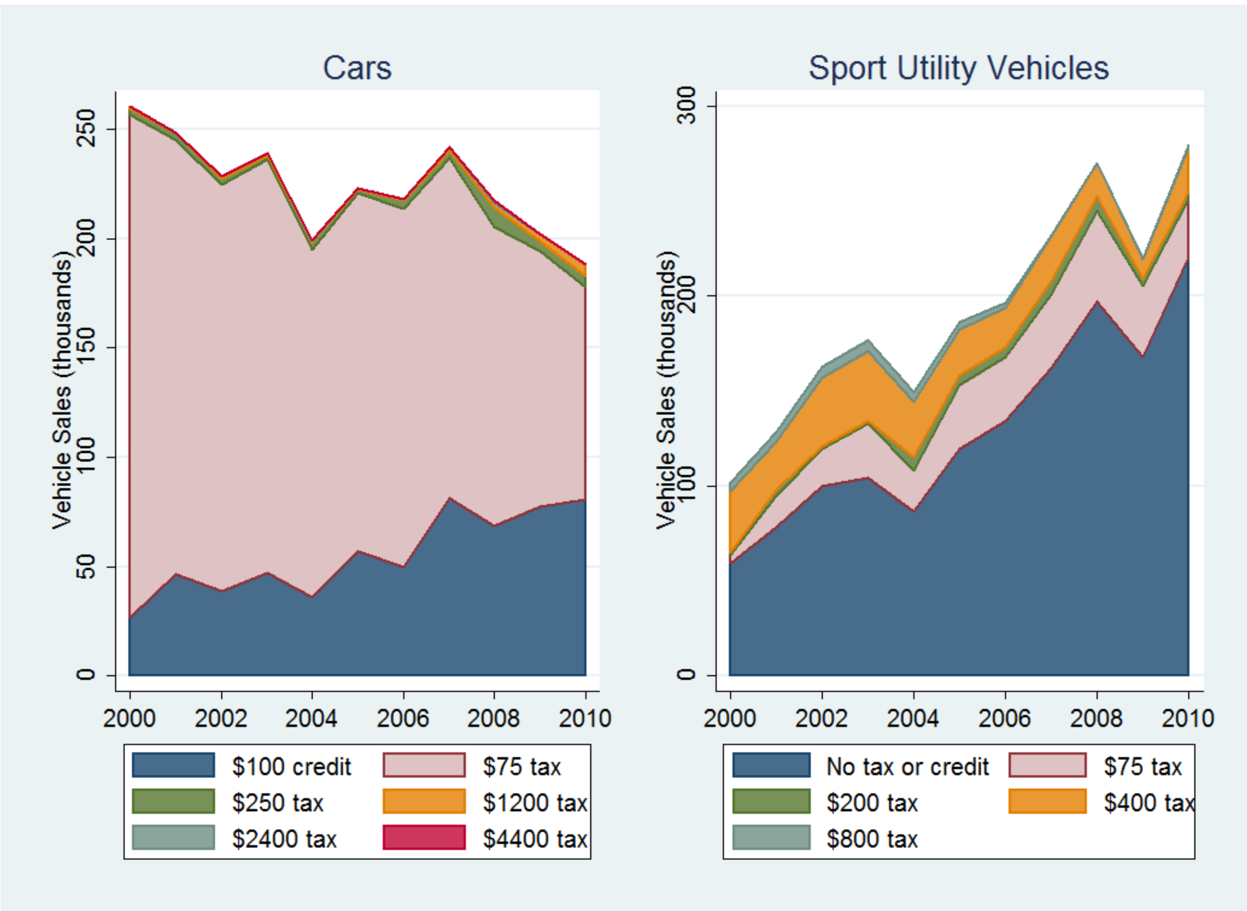

Figure 2: Sales of passenger cars and sport utility vehicles in Ontario by vehicle feebate class.

Feebate classes are expressed in nominal dollars and sales are calculated by the authors using data from Desrosiers Automotive Consultants.

Figure 3: Fuel Consumption Rating (top panel) and Difference Between Rest-of-Canada and Ontarian Market Shares (bottom panel) of 4.6 litre, 8 Cylinder Ford Mustang

Figure 4: Feebate Policies in a Two Vehicle Market

(a) Welfare at the Private Equilibrium

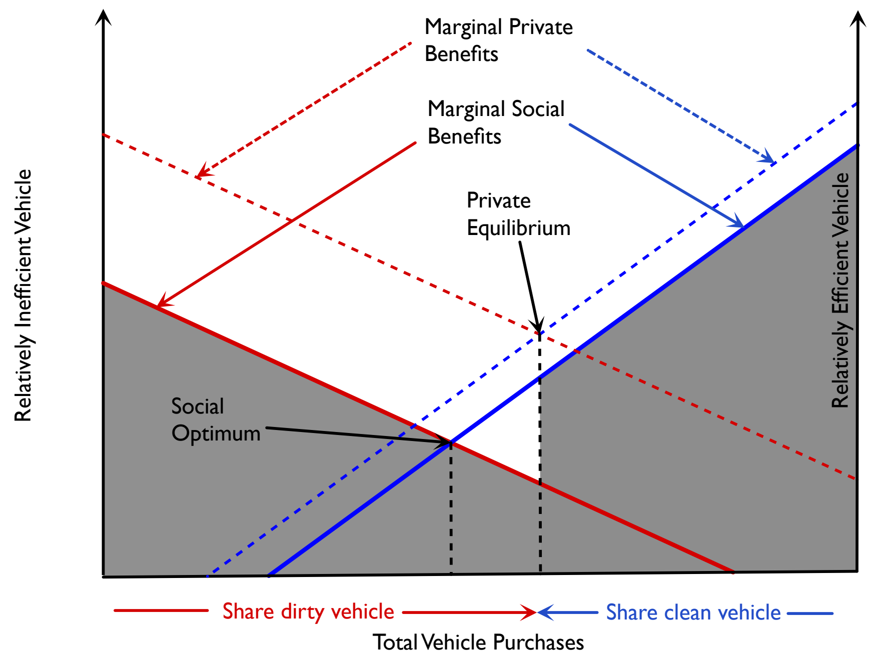

(b) Gain in Social Welfare from Move to the Social Optimum

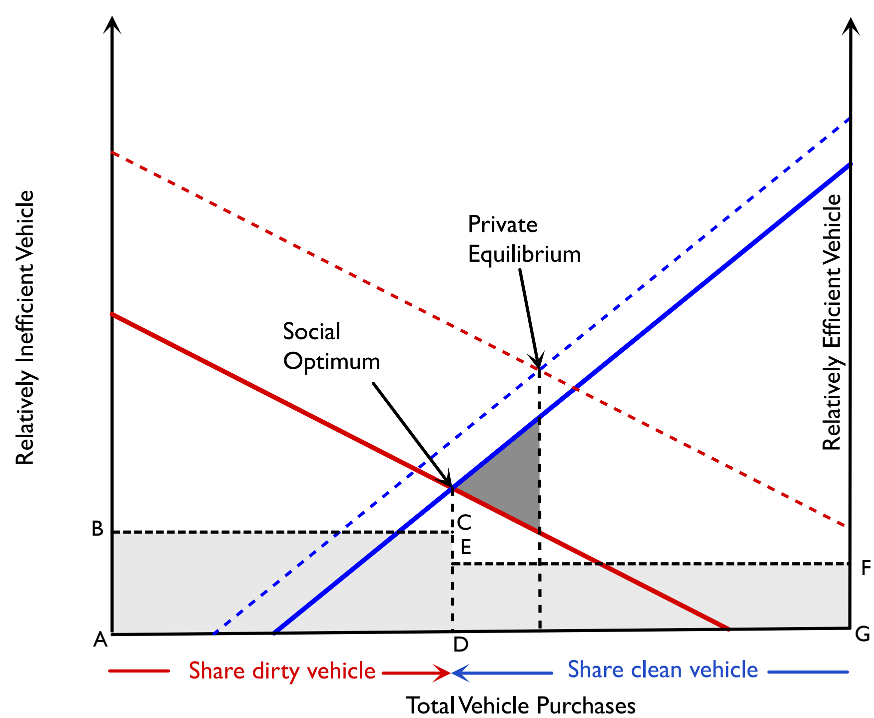

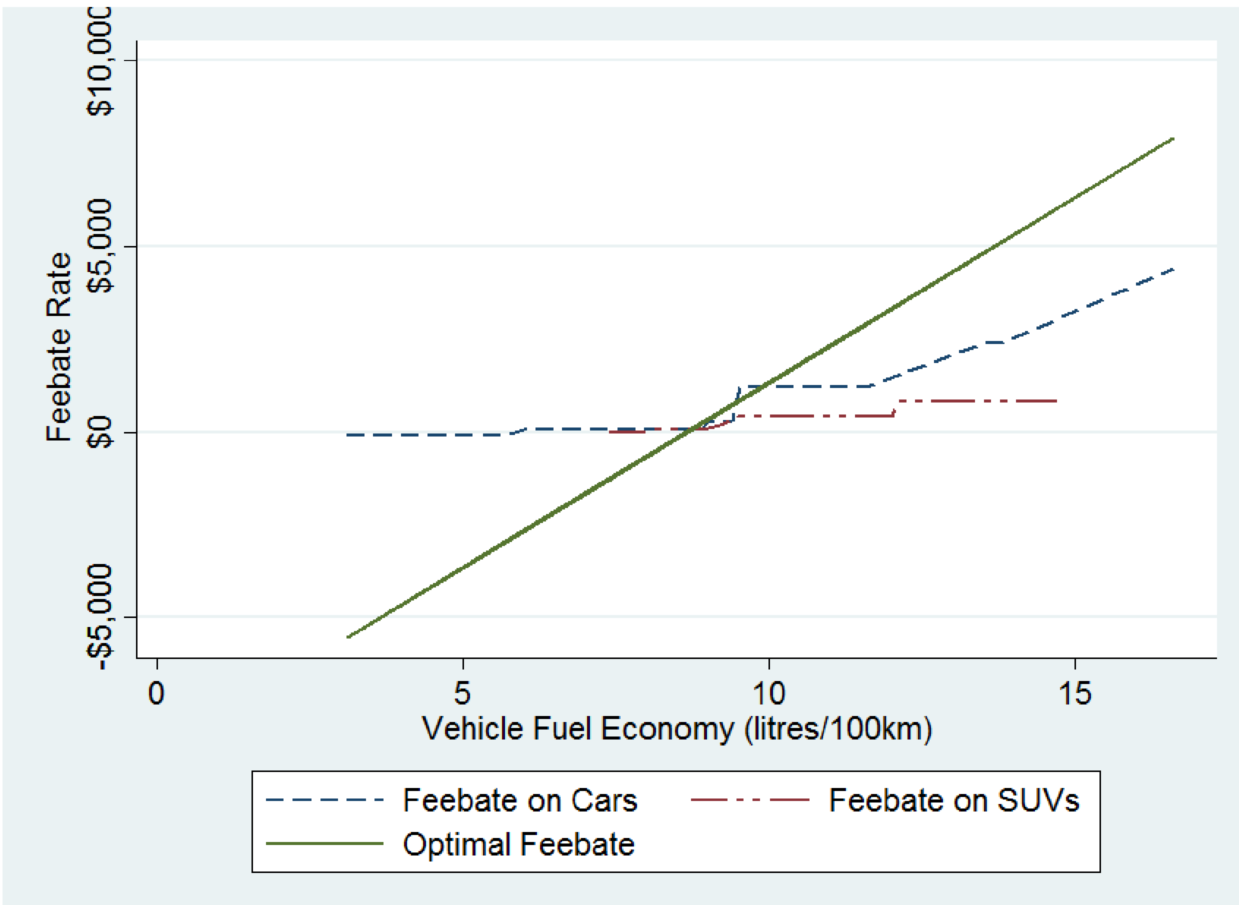

Figure 5: Comparison of the Ontario's Feebates Schedules to the Optimal Schedule

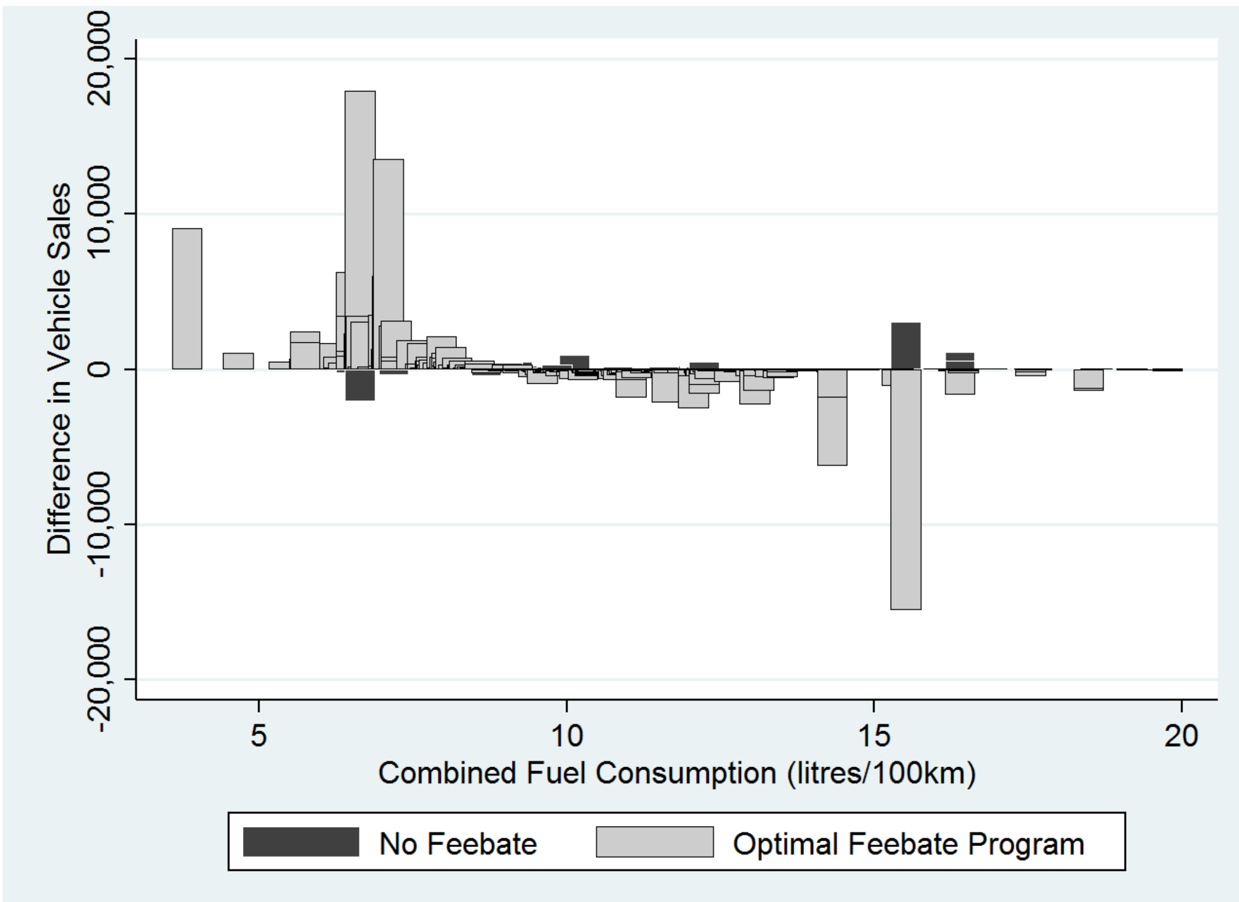

Figure 6: Counterfactual Change in Vehicle Sales Relative to Ontario's Feebate Program

9 Tables

Table 1: Schedule of New Vehicle Fees and Rebates for the Ontario Feebate Program

Highway fuel efficiency rating (L/100km)	1989 Cars	1990 Cars	1991-2010 Cars	1991-2010 SUVs
less than 6.0	-	-	-100	-
6.0-7.9	-	-	75	-
8.0-8.9	-	200	75	75
9.0-9.4	-	700	250	200
9.5-12.0	600	1200	1200	400
12.1-15.0	1200	2400	2400	800
15.1-18.0	2200	4400	4400	1600
over 18.0	3500	7000	7000	3200

All fees and subsidies are in nominal Canadian dollars. Sources: Government of Ontario (1989), Government of Ontario (1991) and Government of Ontario (2010).

Table 2: Effect of Feebates on Vehicle Market Shares Using Provincial-level Data

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Vehicle fixed effects at make-model level						
Feebate ($1,000)	-0.51**	-0.46**	-0.44**	-0.57**	-0.47**	-0.64**
	(0.13)	(0.10)	(0.09)	(0.12)	(0.12)	(0.11)
Panel B: Vehicle fixed effects at make-model-series level						
Feebate ($1,000)	-0.40**	-0.36**	-0.37**	-0.27	-0.32**	-0.27
	(0.11)	(0.08)	(0.07)	(0.25)	(0.09)	(0.27)
Panel C: Vehicle fixed effects at make-model-series-engine size-number of cylinders level						
Feebate ($1,000)	-0.42**	-0.37**	-0.37**	-0.34	-0.33**	-0.30
	(0.11)	(0.08)	(0.08)	(0.32)	(0.09)	(0.29)
Vehicle fixed effects	✓	✓	✓			
Province-class fixed effects	✓	✓			✓	
Class-year fixed effects		✓		✓		
Province-class-year fixed effects			✓			
Vehicle-province fixed effects				✓		✓
Vehicle-year fixed effects					✓	✓
Observations	59,579	59,579	59,579	59,579	59,579	59,579

** - significant at 1%. Values in parentheses are standard errors clustered by province-class.

Coefficients in the table reflect the change in unique vehicle market shares resulting from a $1,000 fee. Gasoline costs are included in all models. Estimates are robust across specifications indicating minor influence of unobservables (Oster, 2013). Panels A, B and C represent distinct vehicle fixed effect resolutions as indicated in the Table. Each of the six columns contains different combinations of fixed effects.

Table 3: Effect of Feebates on Vehicle Market Shares Using FSA-level Data

	(1)	(2)	(3)
Feebate ($1000)	-0.38**	-0.39**	-0.38**
	(0.02)	(0.01)	(0.02)
Make-model-year	✓		
Make-model-series-year		✓	
Make-model-series-engine size-cylinders-year			✓
Make-model-FSA	✓		
Make-model-series-FSA		✓	
Make-model-series-engine size-cylinders-FSA			✓
Observations	16,801,960	16,801,960	16,801,960

** - significant at 1%. Values in parentheses are standard errors clustered by FSA-class.

Coefficients in the table reflect the change in unique vehicle market shares at the FSA-level resulting from a $1,000 fee. Estimates are robust across specifications indicating minor influence of unobservables (Oster, 2013). Each column reflects a distinct resolution of fixed effects corresponding to the Panels in Table 2.

Table 4: Nested Logit Estimates of the Effect of Feebates on Vehicle Market Share

	(1)	(2)	(3)
Feebate (\$1000)	-0.25**	-0.24**	-0.30**
	(0.03)	(0.04)	(0.04)
Share-in-class	0.48**	0.36**	0.17**
	(0.02)	(0.02)	(0.05)
Make-model-year	✓		
Make-model-series-year		✓	
Make-model-series-engine size-cylinders-year			✓
Make-model-province	✓		
Make-model-series-province		✓	
Make-model-series-engine size-cylinders-province			✓
Observations	59,579	59,579	59,579

** - significant at 1%. Values in parentheses are standard errors clustered by province-class.

Coefficients in the table reflect the nested logit effect of a \$1,000 fee on vehicle share accommodating within segment correlation of errors. Each column reflects a distinct resolution of fixed effects corresponding to the Panels in Table 2: (1) uses make-models; (2) uses make-model-series; and (3) uses make-model-series-engine size-number of cylinders. There are 15 classes in the data, segmented according to conventional automobile categorization. All specifications control for gasoline cost. The share-in-class variable is instrumented using the engine size, number of cylinders and fuel economy of all other vehicles in the segment. Estimates are robust across specifications and closely match those in Table 2.

Table 5: Welfare Gains of the Ontario and Optimal Feebate Schedules relative to a No Feebate Scenario

	Annual Vehicles Sold	Lifetime Emissions ($MtCO_2$)	Government Revenue (\$M)	Welfare Gain \$10/tCO_2 (\$M)	Welfare Gain \$30/tCO_2 (\$M)
No Feebate	360,410	20.1	-		
Ontario Feebate Program	360,475	20.3	21.0	19.5	16.5
Optimal Revenue Neutral Feebate	383,610	17.4	-	27.0	81.0

All calculations are based on 2010 data and reflect the welfare gains per vehicle cohort. Emissions calculations assume that each vehicle is driven 16,000km/yr and has a 15 year lifespan. Welfare gains are based on per vehicle emissions and are calculated relative to the no feebate scenario. The social cost of carbon that equates the welfare of the revenue positive Ontario feebate program and the optimal revenue neutral program is \$7.37 per tonne CO_2.

Table 6: Effect of Feebates on Total Passenger Vehicle Sales

	(1)	(2)	(3)	(4)
Feebate program	-0.07	-0.02	0.01	0.02
	(0.04)	(0.04)	(0.04)	(0.04)
GDP per capita		0.58**	0.59**	0.59**
		(0.10)	(0.11)	(0.11)
Unemployment rate			-0.16	-0.15
			(0.08)	(0.08)
Gas price				0.14
				(0.13)
Province fixed effects	✓	✓	✓	✓
Month fixed effects	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓
Observations	3,600	3,600	3,600	3,600

** - significant at 1%. Values in parentheses are standard errors.

Data in this table span 1980-2012 and reflect monthly passenger vehicle sales by province. The dependent variable in each column is logged per capital vehicle sales. *Feebate program* takes a value of one for the province of Ontario for the years 1989 through 2010. GDP per capita, the unemployment rate and gas prices are all logged. Province and month-year fixed effects are included in each model.

Table 7: Asymmetric Consumers Responses to Fees versus Rebates

	(1)
-\$100 fee (rebate)	-5.63**
	(0.42)
\$75 fee	-0.81**
	(0.34)
\$250 fee	-0.86**
	(0.14)
\$400 fee	-0.10
	(0.13)
\$800 fee	-0.13
	(0.08)
\$1200 fee	-0.62**
	(0.03)
\$1600 fee	-0.42**
	(0.09)
≥ \$2400 fee	-0.23**
	(0.02)
Make-model-year	✓
Make-model-FSA	✓
Observations	16,801,960

** - significant at 1%. Values in parentheses are standard errors clustered by FSA-class.

Coefficients in the table reflect the change in unique vehicle market share resulting from the indicated fee. Pairwise F-tests of the null hypothesis of equal coefficient sizes demonstrates that rebates yield a statistically significantly different consumer response than fees.

Table 8: Heterogeneity of Feebate Responsiveness in Urban versus Rural Markets

	(1) *All vehicles*	(2) *Cars only*
Feebate ($1000)	-0.35**	-0.36**
	(0.02)	(0.02)
Feebate*Rural FSA	-0.21**	-0.20**
	(0.05)	(0.05)
Make-model-year	✓	✓
Make-model-FSA	✓	✓
Observations	16,801,960	8,846,661

** - significant at 1%. Values in parentheses are standard errors clustered by FSA-class.

Coefficients in the table reflect the change in unique vehicle market share resulting from a $1000 fee. Feebate*Rural FSA is the interaction term of Feebate and Rural FSA, a dummy variable that takes a value of 1 if Canada Post classifies the FSA as a rural route, which illustrates that rural residents are more sensitive to the feebate schedule than their urban counterparts. All vehicles are included in (1). (2) restricts the sample to passenger cars only.

A Summary Statistics

Table 9 displays the summary statistics for the provincial-level data. Table 10 presents summary statistics for the FSA-level data. Engine displacement is in cubic centimetres; sales are per region-year; fuel economy ratings are in litres/100 km; light duty truck comprise the remainder of the vehicle market not captured by cars and suvs; feebates are in nomial dollars for all vehicles sold in Ontario (including trucks); and, gasoline cost is calculated annually. All observations with zero sales in a region-year are dropped.

Table 9: Summary Statistics for Provincial-level Data

	Mean	**Std. Dev.**	**Min.**	**Max.**
Number of cylinders	5.9	1.55	2	12
Sales (unique vehicle)	182.50	652.88	1	23348
Highway fuel economy	8.70	2.04	3.1	16.6
City fuel economy	12.73	2.99	3.7	27.7
Car share	0.53	0.50	0	1
SUV share	0.25	0.43	0	1
Feebate	178.13	351.70	-100	4400
Gasoline cost	916.19	239.65	237.90	2014.79
Observations	59,579			

Table 10: Summary statistics for FSA-level data

	Mean	**Std. Dev.**	**Min.**	**Max.**
Number of cylinders	6.0	1.63	2	12
Sales (unique vehicle)	1.35	7.33	0	2251
Highway fuel economy	8.76	2.06	3.1	16.6
City fuel economy	12.87	3.05	3.7	27.7
Car share	0.53	0.50	0	1
SUV share	0.25	0.43	0	1
Feebate	147.77	317.11	-100	4400
Observations	16,801,960			

B Econometric Specification

Define the utility that consumer i in province p at time t derives from the purchase of vehicle type k as:

$$U_{ikpt} = \alpha_i(y_{ipt} - q_{kpt}) + \mathbf{X}_{kpt}\beta_i + \zeta_{kpt} + \epsilon_{ikpt} \quad (10)$$

where y is income, q is vehicle price, $\mathbf{X}$ is a vector of observed vehicle characteristics, ζ captures unobserved vehicle characteristics and ϵ is a mean-zero error term that captures the consumer's idiosyncratic preferences over each vehicle type. The parameters α_i and β_i capture individual-level preferences over vehicle price and other characteristics.[36] We decompose utility into a systematic and random component: $U_{ikpt} = V_{ikpt} + \epsilon_{ikpt}$.

Assuming that in the absence of feebates vehicle prices are identical across provinces, we account for the effect of feebates by treating provincial vehicle prices as equal to the Canada-wide vehicle price net of feebates:

$$q_{kpt} = q_{kt} + F_{kpt}$$

where as above F_{kpt} is the dollar value of the feebate applied on vehicle k, with negative values for subsidies (rebates) and positive values for fees (taxes).

Recent evidence using transaction-level data in the US suggests that dealers adjust incentives in response to market conditions (Langer and Miller, 2012). If this occurs in Canadian markets, we expect that a portion of the feebate would be absorbed by vehicle manufacturers or retailers such that the effect of the feebate on vehicle transaction prices would be muted compared to the specification above. To the extent that this occurs, our econometric model should be interpreted as a reduced-form model that captures the equilibrium impact of the policy on vehicle demand, not as a structural specification. Our objective is to evaluate the efficacy of actual feebate policies, thus reduced-form equilibrium impacts are the desired elasticities as they are more interesting and useful for policy analysis.

We treat observable vehicle characteristics $\mathbf{X}_{kpt}$ as consisting of two parts. The first component $\mathbf{X}_k$ is invariant across provinces and time and reflects characteristics such as engine size, vehicle design and safety ratings. Constant for particular vehicles, vehicle-specific fixed effects net out these characteristics in our analysis of the feebate policy. The second component we refer to as G_{kt}, with $G_{kpt} = p^G_{pt} f_k$ being the gasoline cost per kilometer for vehicle k. This is calculated as the price of gasoline multiplied by the vehicle specific fuel consumption rating (measured in liters

[36] This specification imposes a constant marginal utility of income, α_i. This is potentially problematic as vehicles represent a large purchase, and marginal utility may not be constant over the range of vehicle prices. Unfortunately, our data do not include vehicle price, we are unable to adopt a specification where the marginal utility of income is decreasing (Berry et al., 1995).

of gasoline per 100 km of driving).

We now write utility as:

$$U_{ikpt} = \alpha_i y_{ipt} - \alpha_i(q_{kt} + F_{kpt}) + \mathbf{X}_k \beta_i + \phi_i G_{kpt} + \zeta_{kpt} + \epsilon_{ijpt} \tag{11}$$

We next introduce an outside good, indexed by the subscript '0', which involves the option of not purchasing a vehicle. Setting the reference utility for this option to $U_{i0pt} = \alpha_i y_{ipt}$, we assume that consumers purchase one unit of the good (vehicle or outside) that provides the highest utility. Implicitly, this defines values of observed and unobserved vehicle characteristics and preferences that are consistent with the choice of a particular vehicle. The probability that a consumer i will choose any given vehicle k is:

$$P_{ikpt} = \Pr(U_{ikpt} > U_{ijpt}) \quad \forall j \neq k$$
$$P_{ikpt} = \Pr(V_{ikpt} - V_{ijpt} > \epsilon_{ijpt} - \epsilon_{ikpt}) \quad \forall j \neq k$$

To allow a closed-form solution, we eliminate systematic consumer preference heterogeneity, such that $\alpha_i = \alpha$, $\beta_i = \beta$ and $\phi_i = \phi$. The market share of each of the k choices can then be calculated by integration:

$$s_{kpt} = \int_{\epsilon} I(V_{ikpt} - V_{ijpt} > \epsilon_{ijpt} - \epsilon_{ikpt} \quad \forall j \neq k) f(\epsilon) d\epsilon$$

where $I(\cdot)$ is an indicator that equals one if the condition in parentheses is satisfied and zero otherwise. Assuming that ϵ is independent and identically distributed (iid) Type-I extreme value allows evaluation of the integral and yields the familiar logit model:

$$s_{kpt} = \frac{e^{V_{ikpt}}}{\sum_j e^{V_{ijpt}}}$$

Dividing through by s_{0pt} and taking logs gives the log-odds ratio of purchasing each vehicle as follows:

$$\log\left(\frac{s_{kpt}}{s_{0pt}}\right) = -\alpha q_{kt} - \alpha F_{kpt} + \beta X_k - \phi G_{kpt} + \zeta_{kpt} \tag{12}$$

The coefficients of (12) can be estimated using least squares. Unfortunately directly estimating this equation, as is, is problematic: unobserved vehicle characteristics are likely correlated with prices, generating biased estimates. As the goal of this study is to estimate the reduced-form impact of the vehicle feebate policy, we are able to circumvent this problem by introducing vehicle-year fixed effects, θ_{kt}. These fixed effects absorb all vehicle characteristics that are invariant across provinces, including vehicle prices exclusive of the feebate rates. We further decompose the error term into a component that is equal for all provinces and potentially correlated with price, ν_{kt}, and a mean-zero component that varies across markets, μ_{jpt}. The vehicle-time fixed effects absorb the error component ν_{kt}, such that the basic estimating equation is:

$$\log\left(\frac{s_{kpt}}{s_{0pt}}\right) = \theta_{kt} - \alpha F_{kpt} - \phi G_{kpt} + \mu_{kpt}$$

This is equation (1) presented in section 3.

Nested Logit Specification

We also estimate a variant of this equation that relaxes the iid error assumption implicit in the standard logit model. In particular, we estimate a nested logit model in which we group vehicles into 15 nests reflecting the vehicle segment (c). This grouping allows correlation of the unobserved vehicle characteristics according to vehicle segment and permits vehicles within a segment to be treated as closer substitutes for one another than vehicles in different segments. The nested logit specification is estimated as follows:

$$\log\left(\frac{s_{kpt}}{s_{0pt}}\right) = \theta_{kt} - \alpha F_{kpt} - \phi G_{kpt} + \sigma \log(s_{kpt/cpt}) + \mu_{kpt} \tag{13}$$

where $s_{kpt/cpt}$ is the market share of vehicle k within segment c, and σ is the inclusive value parameter, and reflects the substitutability between vehicles within the same segment. This variable is clearly endogenous, so apply the conventional instrumental variable solution. As is typical, our instruments are the sum of all vehicle attributes in segment c, excluding vehicle k (Berry et al., 1995).

C Optimal Feebate Expressions

Three derivations of an optimal feebate schedule are presented. We start with a model where a representative consumer chooses between two vehicles, one which is clean (relatively fuel efficient) and the other which is dirty (relatively fuel inefficient). This is the model discussed in section 4.2. Next, we present the steps which lead to the logit derivation from section 4. Finally, a fully general model is introduced. In each case, consumers select vehicles to maximize utility taking emissions as given, consumer welfare is defined as utility net of emission damages and the government selects tax and rebate policies with the objective of maximizing welfare. Throughout, vehicle supply (vehicle characteristics, prices) is treated as exogenous (see footnote 24 for discussion of this assumption).

C.1 Two Vehicle Model

C.1.1 Consumer

There are two vehicles available to the consumer, clean, indexed with c and dirty, indexed with d. The representative consumer's utility is given by:

$$U(v_c, v_d, x) \tag{14}$$

where x is consumption of other goods. The consumer maximizes utility subject to a wealth, M, constraint:

$$M = (1 + t_c)p_c v_c + (1 + t_d)p_d v_d + x$$

where t_i is the tax (fee or rebate) on vehicle i (a negative value denotes a subsidy), p_i is the tax-exclusive price and x is the numeraire good. First order conditions are given by:

$$\frac{\partial U}{\partial v_c} = \lambda(1 + t_c)p_c \qquad \frac{\partial U}{\partial v_d} = \lambda(1 + t_d)p_d \qquad \frac{\partial U}{\partial x} = \lambda$$

where λ is the marginal utility of wealth. Consumer welfare is utility less damages from emissions:

$$W = U(v_c, v_d, x) - \delta E \tag{15}$$

where δ is the constant marginal disutility from emissions and where emissions, E, are caused by vehicles, such that $E = v_c e_c + v_d e_d$. e_c and e_d are the per vehicle lifetime emissions from v_c and v_d, respectively.[37]

C.1.2 Government

The government uses taxes on vehicles to maximize social welfare, subject to a revenue constraint, $\bar{R}$:

$$t_c p_c v_c + t_d p_d v_d = \bar{R}.$$

The government's optimal tax problem can be set up by differentiating welfare:

[37] See main text for a more detailed emissions function.

$$dW = \frac{\partial U}{\partial v_c} dv_c + \frac{\partial U}{\partial v_d} dv_d + \frac{\partial U}{\partial x} dx - \delta dE$$

Substituting in first order conditions and the differentiated emissions function ($dE = e_c dv_c + e_d dv_d$) gives:

$$dW = \lambda(1 + t_c) p_c dv_c + \lambda(1 + t_d) p_d dv_d + \lambda dx - \delta(e_c dv_c + e_d dv_d) \tag{16}$$

Differentiating the wealth constraint ($dM = 0 = p_c v_c dt_c + (1+t_c) p_c dv_c + p_d v_d dt_d + (1+t_d) p_d dv_d + dx$), substituting and rearranging gives:

$$\frac{dW}{\lambda} = -p_c v_c dt_c - p_d v_d dt_d - \frac{\delta}{\lambda}(e_c dv_c + e_d dv_d)$$

Casting the optimal tax problem in terms of the tax on dirty vehicles yields:

$$\frac{1}{\lambda}\frac{dW}{dt_d} = -p_c v_c \frac{dt_c}{dt_d} - p_d v_d - \frac{\delta}{\lambda}(e_c \frac{dv_c}{dt_d} + e_d \frac{dv_d}{dt_d})$$

Substituting in the differentiated revenue constraint gives:

$$\frac{1}{\lambda}\frac{dW}{dt_d} = p_c t_c \frac{dv_c}{dt_d} + p_d t_d \frac{dv_d}{dt_d} - \frac{\delta}{\lambda}(e_c \frac{dv_c}{dt_d} + e_d \frac{dv_d}{dt_d})$$

Setting $dW = 0$ and solving for the optimal dirty tax gives:

$$t_d = \left(\frac{\delta}{\lambda}\frac{e_c}{p_d} - \frac{p_c}{p_d} t_c \right) \frac{dv_c}{dv_d} + \frac{\delta}{\lambda}\frac{e_d}{p_d} \tag{17}$$

We can further refine this by substituting in the government budget constraint. We consider only a revenue neutral feebate such that $\bar{R} = 0$:

$$t_d = \frac{\delta}{\lambda} \frac{\frac{e_c}{p_d}\frac{dv_c}{dv_d} + \frac{e_d}{p_d}}{1 - \frac{v_d}{v_c}\frac{dv_c}{dv_d}}. \tag{18}$$

Using the consumer's optimality conditions, it is possible to further simplify:

$$t_d = \frac{\delta}{\lambda}\left(\frac{\frac{e_d}{p_d} - \frac{e_c}{p_c}}{1 + \frac{w_d}{w_c}}\right)$$

where w_i is the expenditure share on vehicle i. The model is general, so symmetric expressions can be derived. The clean vehicle's rebate, for instance, is:[38]

$$t_c = -\frac{\delta}{\lambda}\left(\frac{\frac{e_d}{p_d} - \frac{e_c}{p_c}}{1 + \frac{w_c}{w_d}}\right)$$

Next, if we consider a special case where $\frac{dv_c}{dv_d} = -1$, such that there is no extensive margin adjustment and the feebate does not change the total amount of vehicles purchased. This is the scenario examined in section 4.2. Using (18) we obtain:

$$t_d = \frac{\delta}{\lambda}\left(\frac{\frac{e_d - e_c}{p_d}}{1 + \frac{v_d}{v_c}}\right) \tag{19}$$

It is also possible to compare the optimal feebate to the optimal indirect Pigouvian tax, which is rebated through lump-sum transfers to the consumer:[39]

$$t_d^P = \frac{\delta}{\lambda}\frac{e_d}{p_d} \tag{20}$$

Using (20) it is possible to rewrite the optimal feebate expressions as:

$$t_d = s_c\big(t_d^P - t_c^P\big) \tag{21}$$

where s_c is the share of the clean vehicle purchased. Optimal feebates therefore are a function of the optimal indirect Pigouvian taxes.[40] Feebates are also everywhere less than optimal Pigouvian taxes. The following figure illustrates this across a continuum of emissions.

[38]Of course only one fee/rebate is independent; the other tax is determined from the revenue constraint.

[39]This is solved as follows. First, define the new wealth constraint as $M = (1 + t_c)p_c v_c + (1 + t_d)p_d v_d + x - Z$, where $Z = t_c p_c v_c + t_d p_d v_d$ is tax revenue. Differentiate and substitute into (16) to get $\frac{dW}{\lambda} = -p_c t_c dv_c - p_d t_d dv_d - \frac{\delta}{\lambda}(e_c dv_c + e_d dv_d)$. Divide by dt_d, set $\frac{dv_c}{dt_d} = 0$ and get (20).

[40]An equivalent expression for the clean vehicle is: $t_c = s_d\big(t_c^P - t_d^P\big)$.

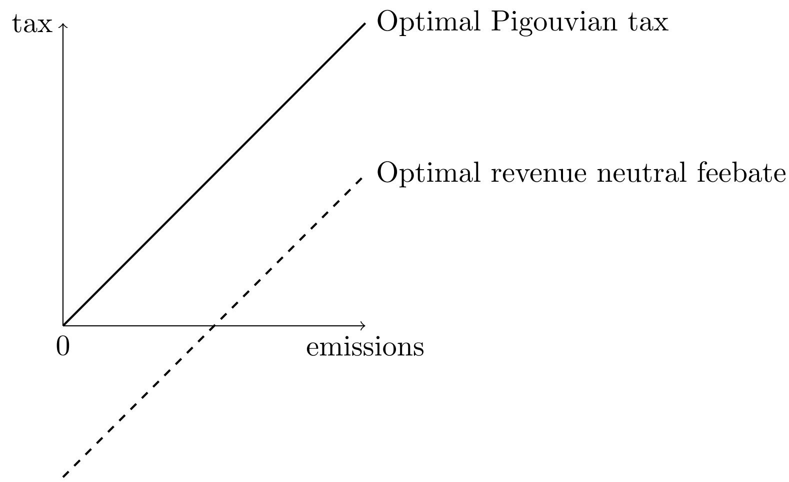

Figure 7: Optimal feebate schedule

C.2 Discrete Choice/Logit Decision Problem

Next, we derived the expression in the text which employs a multinomial logit choice structure. The social planner's problem is:

$$\max_{\boldsymbol{F}} \ln \left[\exp \tilde{V}_0 + \sum_{k=1}^{K} \exp \tilde{V}_k \right] - \delta E$$

subject to:

$$\sum_{k=1}^{K} \pi_k F_k \geq R$$

Under revenue neutrality, we can rewrite the budget and obtain some additional intuition on self-financing constraints:

$$\sum_{k=1}^{K} \pi_k F_k = R = 0 \text{ which can be written as:}$$
$$\sum_{k=1}^{K} \pi_k F_k - \bar{F} = 0$$

where $\bar{F}$ is the average feebate. As K is large and fuel consumption ratings are continuous, any specific F_k will have a negligible effect on $\bar{F}$ (i.e., ${}^{\partial \bar{F}}/_{\partial F_l} = 0$). Thus, it is possible to write the first order conditions as:

$$\pi_l \frac{\partial \tilde{V}_l}{\partial F_l} - \delta \frac{\partial E}{\partial F_l} + \lambda \sum_{k=1}^{K} \frac{\partial \pi_k}{\partial F_l} F_k = 0$$

Rearranging gives:[41,42]

$$\sum_{k=1}^{K} \frac{\partial \pi_k}{\partial F_l} F_k = \frac{1}{\lambda}\big(\delta\phi\pi_l\beta x(f_l - \bar{f}) - \pi_l\beta\big)$$

Decomposing the LHS sum gives:

$$\frac{\partial \pi_l}{\partial F_l} F_l + \sum_{k\neq l}^{K} \frac{\partial \pi_k}{\partial F_l} F_k = \frac{1}{\lambda}\big(\delta\phi\pi_l\beta x(f_l - \bar{f}) - \pi_l\beta\big)$$

Substituting in conditional logit marginal effects gives:

$$\beta\pi_l(1-\pi_l)F_l - \sum_{k\neq l}^{K} \beta\pi_k\pi_l F_k = \frac{1}{\lambda}\big(\delta\phi\pi_l\beta x(f_l - \bar{f}) - \pi_l\beta\big)$$

Removing $\beta\pi_l$ from LHS gives:

$$\beta\pi_l \left(F_l - \pi_l F_l - \sum_{k\neq l}^{K} \pi_k F_k \right) = \frac{1}{\lambda}\big(\delta\phi\pi_l\beta x(f_l - \bar{f}) - \pi_l\beta\big)$$

For a revenue-neutral feebate we have $\pi_l F_l = -\sum_{k\neq l}^{K} \pi_k F_k$, so:

$$F_l = \frac{\delta\phi\pi_l\beta x(f_l - \bar{f}) - \pi_l\beta}{\lambda\pi_l\beta}$$
$$F_l = \frac{\delta\phi x}{\lambda}(f_l - \bar{f}) - \frac{1}{\lambda}$$

[41] This step invokes the envelope theorem which is possible as consumers are optimizing (see (3)), the externality is well-defined at the optimum (see (8)) and the constraint binds due to revenue neutrality.

[42] We use marginal effects from conditional logit functions, $\frac{\partial \pi_k}{\partial F_l} = -\beta\pi_k\pi_l$ and $\frac{\partial \pi_l}{\partial F_l} = \beta\pi_l(1-\pi_l)$ in formulating the expression.

which is the expression in the text.

C.3 A More General Model

In a complete model, consumers choose between $k = 1 \dots K$ vehicle models. Consumer utility is given by $U = U(v_k, x)$, and the wealth constraint is $M = \sum_K (1 + t_k) p_k v_k + x$. Differentiating consumer welfare as above gives:

$$dW = \sum_k \frac{\partial U}{\partial v_k} dv_k + \frac{\partial U}{\partial x} dx - \delta dE$$

Substituting first order conditions and the differentiated wealth constraint gives:

$$\frac{dW}{\lambda} = -\sum_k p_k v_k dt_k - \frac{\delta}{\lambda} \sum_k e_k dv_k$$

Dividing through by dt_j gives K optimal feebate conditions:

$$\frac{1}{\lambda} \frac{dW}{dt_j} = -\sum_{k \neq j} p_k v_k \frac{dt_k}{dt_j} - p_j v_j - \frac{\delta}{\lambda} \sum_k e_k \frac{dv_k}{dt_j}$$

Differentiating the revenue constraint gives $p_j v_j = -\sum_{k \neq j} p_k v_k \frac{dt_k}{dt_j} - \sum_k t_k p_k \frac{dv_k}{dt_j}$. Substituting this into the expression above gives:

$$\frac{1}{\lambda} \frac{dW}{dt_j} = \sum_k t_k p_k \frac{dv_k}{dt_j} - \frac{\delta}{\lambda} \sum_k e_k \frac{dv_k}{dt_j}$$

Solving for t_j at the welfare maximum gives:

$$t_j = \sum_{k \neq j} \left(\frac{\delta}{\lambda} \frac{e_k}{p_j} - \frac{p_k}{p_j} t_k \right) \frac{dv_k}{dv_j} + \frac{\delta}{\lambda} \frac{e_j}{p_j}$$

which directly compares with (17), the expression derived for the two-vehicle case. Applying the consumer's optimality conditions, we can further simplify:

$$t_j = \sum_{k \neq j} t_k + \frac{\delta}{\lambda} \left(\frac{e_j}{p_j} - \sum_{k \neq j} \frac{e_k}{p_k} \right)$$

In the multiple vehicle case, $K-1$ of the optimal tax expressions are independent; the remaining tax is derived from the revenue constraint, which equals:

$$t_K = -\frac{\sum_{j=1}^{K-1} p_j v_j t_j}{p_K v_K}.$$

Solving for t_j involves $K-1$ of the tax expressions and the revenue expression. Therefore, the optimal t_j, written in terms of K, is:[43]

$$t_j = \frac{w_K}{w_K + w_j} \left(\frac{\delta}{\lambda} \right) \left(\frac{e_j}{p_j} - \frac{\bar{e}}{\bar{p}} \right) \tag{22}$$

where $\bar{e}$ and $\bar{p}$ are the average emissions and price over all $k \neq j$ vehicles and w_i is the expenditure on any vehicle i.

[43]This assumes that $\bar{t} \approx 0$; otherwise a small adjustment term would be included in this expression.