Kill Two Birds with One Stone? Environmental Policy Design with Multiple Targets in the Swedish Car Market*

Cristian Huse† Claudio Lucinda‡

July 2026

Abstract

Governments increasingly subsidize dual-fuel vehicles – flex-fuel cars in the 2000s, plug-in hybrids today – on the premise that buyers will use the cleaner fuel. We evaluate this class of policy through an *ex-ante* structural analysis of the 2007-2009 Swedish Green Car Rebate, a subsidy charged with two EU-mandated targets at once: lower new-vehicle CO2 and higher renewable-fuel use. Estimating demand on pre-policy data and simulating a grid of alternative designs, we show the twin failure was predictable. A single rebate cannot serve two targets (a Tinbergen conflict); strong within-segment substitution confines its reach so that even large rebates barely move fleet composition; and asymmetric compliance costs steer manufacturers toward high-emission flex-fuel models rather than genuinely efficient ones. The welfare consequences are counterintuitive. Because the rebate expands total sales, every design *raises* total emissions – and monetized environmental damage – even as it lowers CO2 per vehicle. The program's measured welfare gain is almost entirely an inframarginal transfer: some 90 percent of the outlay subsidizes buyers who would have purchased an eligible vehicle regardless. The lesson transfers directly to today's plug-in hybrid subsidies, which rest on the same optimistic bet about how the vehicle is used.

JEL Classification: H22, H23, L62, Q42, Q48, Q54, Q58.
Keywords: Green car subsidies; Plug-in hybrid vehicles; Tinbergen rule; Multiple policy targets; Ex-ante policy evaluation; Unintended consequences; Alternative-fuel vehicles; Discrete choice demand; CO2 emissions.

*We thank Tobias Olsson at Vroom Stockholm for making data available. We are also grateful to Shaun McRae and the audiences at the International Industrial Organization Conference (IIOC) and the University of Freiburg for helpful suggestions. The usual disclaimer applies.

†University of Oldenburg. Email address: cristian.huse@uol.de. Web: http://cristianhuse.github.io. Correspondence address: Department of Economics, University of Oldenburg, Postfach 5634, 26046 Oldenburg, Germany. Phone: +49 441 798 4670.

‡University of Sao Paulo (FEA). Email address: claudiolucinda@usp.br.Web: http://www.claudiolucinda.net.

1 Introduction

Governments routinely design a single policy instrument to pursue several objectives at once. Nowhere is this more consequential, or more expensive, than in the transport sector: in both the EU and the US, fuel-economy or CO2 standards coexist with renewable-fuel mandates, and the vehicle purchase subsidies that have become a central pillar of climate policy are routinely justified on industrial-policy and stimulus grounds in addition to environmental ones.[1] This practice sits in tension with a long-standing principle of economic policy, due to Tinbergen (1952): achieving multiple independent targets generally requires at least as many independent instruments. The *Tinbergen Principle* is widely invoked but rarely tested structurally in a real market, in part because the data required to separate the targets, the instruments, and the underlying behavioural frictions are seldom available.

We provide such a test by analyzing a uniquely informative historical precedent: the 2007-2009 Swedish Green Car Rebate (GCR).[2] The GCR was explicitly tasked with two distinct, EU-mandated objectives: lowering the average CO2 emissions of newly registered vehicles, and raising the share of renewable fuels toward the Renewable Fuel Standard (RFS). It is an ideal laboratory for the Tinbergen question. The new vehicle market offers unusually rich data – registrations at the brand-model-fuel level, detailed vehicle characteristics, and individual-level transition patterns – and the GCR's asymmetric design, which set sharply different eligibility thresholds for conventional and alternative-fuel vehicles, generates exactly the variation needed to ask whether one instrument can serve two masters.[3] Crucially, the episode is historical: we can use its pre-period to place ourselves in the shoes of a 2006 policymaker and ask, using only pre-policy data, what could have been known *before* the programme was launched.

The exercise is far from a historical curiosity. The GCR subsidised flexible-fuel vehicles (FFVs) on the basis of an optimistic assumption about how often drivers would actually use ethanol rather than gasoline. This is structurally the same bet that underlies today's subsidies for plug-in hybrid electric vehicles (PHEVs), which are credited with low official emissions on the assumption of a high "utility factor" – the share of kilometres driven on electricity rather than fuel. A growing body of evidence documents a substantial PHEV performance gap, with on-road CO2 emissions several times higher than official ratings, precisely because real-world utility factors fall well short of the regulatory assumption (e.g., Plötz et al., 2022). The GCR is thus a completed natural experiment for a policy design that is being repeated, at far greater fiscal scale, across the global transition to electric mobility. The question we ask of the GCR – can a demand-side subsidy green the new vehicle fleet when consumers face strong frictions and the cleaner fuel is optional? – is the question now facing PHEV policy.

[1] Government spending on electric vehicle (EV) incentives is a cornerstone of global climate policy. The International Energy Agency (IEA) estimated that public spending on EV subsidies and incentives reached almost USD 30 billion in 2021 (IEA, 2022), and that annual government spending on electric cars – direct purchase incentives plus foregone purchase and registration tax revenue – has since remained around USD 38 billion per year (International Energy Agency, 2025). Plug-in hybrids continue to absorb a substantial share of this support: in China, the largest market, nearly 45 percent of vehicle subsidies accrue to them.

[2] This paper is distinct from our earlier work on the same policy. Huse and Lucinda (2014) provide an ex-post analysis of the GCR's market impact, and Huse (2026) analyzes how its asymmetric design was exploited by manufacturers. The present paper is ex-ante and structural: it uses only pre-policy data and evaluates a full grid of counterfactual designs, which lets us separate predictable failure – rooted in pre-existing frictions – from outcomes that could only have been learned after the fact.

[3] For the purposes of the GCR, alternative fuels are fuels other than gasoline and diesel, such as ethanol, electricity, and CNG.

Empirical Strategy. We develop and estimate a structural model of demand for the Swedish new-vehicle market. We employ a random coefficients nested logit (RCNL) model estimated on pre-policy data (2004-2006), which captures heterogeneous consumer preferences and the strong within-segment substitution that characterises the market. We then use the estimated model to simulate market outcomes under a comprehensive grid of counterfactual policies. Each policy is a tuple (τ_r, τ_a, R), where R is the rebate amount and τ_r, τ_a are the CO2 eligibility thresholds for regular- and alternative-fuel vehicles; this parametrisation nests the actual GCR and spans both technology-neutral (symmetric) and technology-biased (asymmetric) designs.

The empirical analysis is disciplined by a theoretical framework that is built on the demand framework we take to data and delivers three testable propositions. First, a uniform rebate cannot simultaneously satisfy an arbitrary CO2 target and an arbitrary renewable-fuel target: since the rebate leaves the internal composition of the eligible "green" nest unchanged, both targets collapse onto a single aggregate margin and become perfectly collinear – a structural statement of the Tinbergen conflict. Second, the rebate's power to shift the fleet toward green vehicles is governed by their *baseline* market share, not by the within-segment correlation parameter; when green vehicles are a niche, the subsidy operates largely as an inframarginal transfer. Third, when compliance thresholds are asymmetric, manufacturers direct their effort toward the cheapest pathway – converting existing models into high-emission FFVs rather than engineering genuinely low-emission models. The propositions operate on the very objects the structural model identifies – inclusive values, nest shares, within- versus cross-nest substitution. Therefore, each of them yields a prediction on an object the structural model identifies – the disjointness of CO2- and RFS-optimal designs, the inframarginal share governed by the baseline green share, and the corner solution in product entry – and each is borne out: the theory is confirmed on the estimates, not merely motivated by them.

Main Findings. Our central result is that no policy in our grid meets the CO2 and RFS targets jointly, and the failure is not one of insufficient generosity: rebates several times the actual GCR leave both targets far out of reach. The designs closest to the CO2 target (moderate thresholds near 160 gCO2/km) are systematically different from those that maximize alternative-fuel adoption (the laxest thresholds), exactly as the Tinbergen conflict predicts. We trace this to two reinforcing frictions. On the supply side, the segments Swedish consumers actually bought – large station wagons and family cars – had no low-emission versions on offer, and none were introduced during the program. On the demand side, individual-level transition matrices reveal pronounced state-dependence: buyers overwhelmingly repurchase within the same size class, brand, and model family, so the rebate reshuffles demand within segments rather than pulling consumers toward smaller, cleaner cars. With a pre-policy green share near six percent, even a large rebate is bounded by the niche status of green vehicles, as our framework anticipates.

Two consequences of these frictions are, to us, the paper's most interesting findings. First, the program is environmentally counterproductive on net. By lowering the price of eligible vehicles the rebate expands total sales, and the emissions from the additional vehicles outweigh the modest reduction in CO2 *per* vehicle – so every design we examine *raises* total CO2, and raises total monetized environmental damage, even as it improves the regulated average CO2 emissions per vehicle. Worse, the designs most effective at cutting per-vehicle CO2 do so partly by tilting demand toward diesels, raising local pollutants (NOx and PM) by three to seven percent: a shift of the environmental burden from a global to a local pollutant, and a textbook second-best outcome. The same volume-versus-composition tension shadows the plug-in hybrid credits of the recent U.S. subsidy expansion, where per-vehicle ratings improve even as fleet-wide effects are muted or perverse.

Second, the program's welfare case rests entirely on the surplus it transfers, not on any environmental return. From the government's budget perspective, which counts the rebate outlay as a cost, every design reduces welfare. But a rebate is a transfer: we show it returns to the private sector almost in full – as producer surplus on additional sales and incremental VAT – so that under the standard total-surplus criterion, in which transfers cancel, welfare rises. The gap between the two verdicts is precisely the outlay. What the rebate buys with that outlay, however, is almost entirely inframarginal: only some two to three percent of subsidized purchases are genuinely induced, the rest going to buyers who needed no inducement. The program is thus best understood not as an inefficient purchase of environmental benefit but as an expensive, poorly targeted transfer that leaves the environment worse off.

Contribution and Related Literature. The paper makes three contributions. First, and most directly, it provides what is to our knowledge the first structural, *ex-ante* evaluation of a subsidy designed to meet both a CO2 and a renewable-fuel target. The distinction is methodological as well as substantive. Structural evaluations of transport and environmental policy are overwhelmingly *ex post*: they estimate demand on data that span the policy and quantify what occurred (e.g., Li, Liu and Wei, 2022; Reynaert, 2021). The ex-ante stance we adopt – estimating preferences only on pre-policy data and asking what a 2006 policymaker could have foreseen – is far less common, precisely because it forecasts behaviour out of sample under designs that were never implemented. The closest precedent in this respect is Adamou, Clerides and Zachariadis (2014), who simulate hypothetical feebate schemes for the German market; relative to their work we study an enacted policy with two explicit and potentially conflicting targets, and we ground the exercise in a theoretical framework that delivers the conflict as a structural prediction. In doing so, we turn the Tinbergen rule from an abstract counting principle into a concrete, model-based mechanism. Our setting also connects to the structural analysis of taxation in differentiated-product markets (Anderson, de Palma and Kreider, 2001; Fershtman, Gandal and Markovich, 1999) and to the theory of instrument choice under multiple externalities (Fullerton and West, 2002; Fullerton and Heutel, 2010). The closest empirical antecedent on multiple objectives is Li, Liu and Wei (2022), who study whether a US scrappage programme can jointly serve environmental and stimulus goals; we complement their ex-post analysis with a different policy, a different market, an ex-ante design, and an explicit account of the demand- and supply-side frictions that make the single instrument fail.

Second, we add a mechanism to the literature on the limits of demand-side environmental policy. A large body of work attributes the underperformance of such policies to the energy-efficiency gap and consumer myopia (e.g., Busse, Knittel and Zettelmeyer, 2013; Allcott and Wozny, 2014), or to the limited salience of policy instruments (Huse and Koptyug, 2022). We show that policy can fail even absent these channels: powerful state-dependence in brand, model, and size-class choice acts as a structural barrier to substitution, grounding micro-level concepts such as status-quo bias and brand loyalty (Dubé, Hitsch and Rossi, 2010) in aggregate, policy-relevant outcomes. We further provide an empirical counterpart to the theory of directed technical change (Acemoglu et al., 2012): a pure demand-side "pull" fails when the supply-side "push" is absent, because manufacturers – facing asymmetric compliance costs – expand the cheap, high-emission compliance pathway rather than innovating. This is the same loophole that Anderson and Sallee (2011) document for flexible-fuel vehicles under US CAFE, and that Huse (2026) documents directly for the GCR. Whereas the latter paper analyzes the GCR ex post and in reduced form, our contribution is the ex-ante, structural lens: using only pre-policy data, we ask whether the policy's failure was foreseeable from the market's underlying frictions, and we evaluate the full space of alternative designs the policymaker could have chosen.

Third, we quantify a trade-off between global and local pollutants induced by single-metric vehicle policy. While the engineering tension between CO2 and NOx in diesel engines is well known (Reynaert, 2021), our structural model shows how a specific policy design actively drove a fleet-wide shift toward higher-NOx vehicles, providing a sharp empirical illustration of second-best reasoning (Fullerton and Heutel, 2010). Taken together with the growing literature questioning the cost-effectiveness and equity of modern EV and PHEV subsidies (see Sheldon, 2022, for a survey), our findings offer a historical and structural foundation for why dual-fuel subsidies tend to underperform: the lesson of the GCR is being relearned, at far greater expense, in the PHEV era. We further show that this trade-off is not confined to the CO2-NOx margin: once the volume response is accounted for, the program raises *total* emissions of the very pollutant it targets. This distinction between the regulated per-vehicle average and the welfare-relevant total is, we argue, a general hazard of quantity-based green subsidies, and one that ex-post evaluations focusing on average ratings can miss.

The remainder of the paper is organised as follows. Section 2 describes the institutional background. Section 3 presents the theoretical framework. Section 4 describes the data. Section 5 documents the supply- and demand-side frictions. Section 6 presents the demand estimation. Section 7 reports the policy simulations. The final section concludes.

2 Institutional Background

The 2007 Swedish Green Car Rebate (GCR) provides a compelling case study in policy design, as it was implicitly tasked with achieving two distinct, and potentially conflicting, national objectives. The first was to reduce average CO2 emissions from new vehicles. The second was to attain the Renewable Fuel Standard (RFS) of 5.75 percent set by the EU Biofuels Directive.[4]

EU Regulatory Framework. During most of the 2000s, European countries were subject to directives from the European Union (EU) in the same vein as US states are subject to regulations put forth by the US Environmental Protection Agency (EPA). That is, just like US states, EU countries had the freedom to choose the policy instruments in order to comply with those directives.

Since emissions from passenger cars account for roughly half of all CO2 emissions in the transport sector, the EU introduced a directive to tackle CO2 emissions from newly registered vehicles in the 2000, see Regulation 1753/2000/EC. According to this directive, newly-registered passenger vehicles in EU countries would have to attain average emissions of 120 gCO2/km by 2010 in order to combat climate change;[5] this target corresponds to 47 (51.7) mpg running on gasoline (diesel), thus way more ambitious than the post-2011 US CAFE standards.[6]

[4]The EU Biofuels Directive, formally the Directive on the Promotion of the use of biofuels and other renewable fuels for transport, 2003/30/EC, is the EU counterpart of the US Renewable Fuel Standard. According to Swedish Consumer Agency (2011), life-cycle CO2 emissions of the ethanol typically used in Sweden are 55 percent lower than those of gasoline. Moreover, anecdotal evidence suggests that this was a way the newly-formed government managed to get the support of the Green Party.

[5]Such target was later relaxed to 130 gCO2/km to be attained by 2015, see Regulation 443/2009/EC.

[6]Alternatively, such target corresponds to fuel consumption of approximately 5 (4.5) liters of gasoline (diesel) per 100 km. Moreover, diesel cars must also have a particle filter, so that their particle emissions are less than 5 mg/km. For perspective, according to the 2011 CAFE standards, a small vehicle (up to 41 sq ft footprint), e.g. Honda Fit, has to attain fuel economy of 47 mpg by 2019 while a larger vehicle (from 55 sq ft footprint), e.g.

Figure 1: **Average CO2 Emissions of New Vehicles in Selected EU Countries**

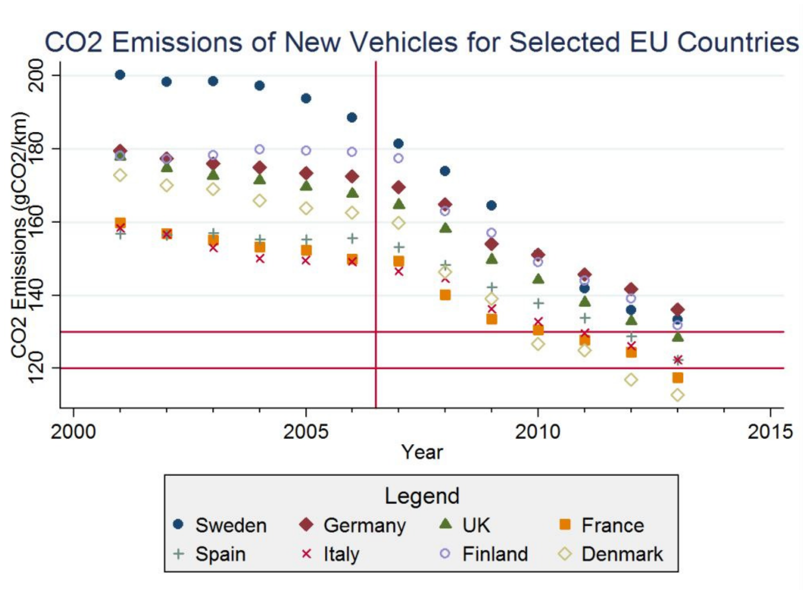

Note: Data based on all (by consumers and businesses) newly registered vehicles from European Environmental Agency (2014). Countries reported in addition to Sweden have the full series of average CO2 emissions available since 2001. Among these, we selected the ones with the five biggest economies within the EU in terms of nominal GDP in 2013, plus Nordic countries Denmark and Finland. The vertical line is drawn between the years 2006 and 2007, when the Swedish Green Car Rebate was introduced (see below). The horizontal lines represent the CO2 emission targets of 120 and 130 gCO2/km.

As in most EU countries average CO2 emission levels of newly registered vehicles in the early 2000s were well above the targets set by the EU (see Figure 1), policymakers were under pressure to meet such stringent targets in a relatively short period. The Swedish case is particularly worrisome for two reasons. First, because CO2 emissions from the passenger car fleet are estimated to account for 19 percent of total CO2 emissions in the country, as compared to the EU average of 12 percent. Second, because the average CO2 emissions of newly registered vehicles in Sweden during the period 2001-2006 have consistently been the highest when compared to EU countries for which such data is available (European Environmental Agency 2014, Table A3).

Given such bleak picture facing EU countries, it was no surprise that a number of policies aimed at the transport sector – new vehicles in particular – were introduced starting from the mid-2000s. Taking advantage of their freedom to choose policy instruments and the strong recession in the wake of the global crisis, most EU countries have launched different versions of stimulus/scrappage programs in 2008 and 2009, very much in the spirit of the US "Cash for Clunkers" program.[7]

Program Background. The Swedish GCR was an environmental program aimed at the new vehicle market. It provided a one-time, lump-sum rebate of SEK 10,000 (approximately USD 1,500, or 6 percent off of the price of a 2009 VW Golf) to private buyers of new "green cars". A car was deemed "green" if it met one of two criteria: (1) it could be powered by an alternative fuel, such as ethanol (E85), or gas (CNG/LPG) and had CO2 emissions below 220 gCO2/km;[8] or (2) it was a conventional gasoline or diesel vehicle with CO2 emissions below 120 gCO2/km. This dual-track eligibility is central to our analysis, as it created distinct incentives for different types of "green" technologies. Vehicles satisfying the former were called (green) Alternative Fuel Vehicles (AFVs) whereas those satisfying the former were called (green) Regular Fuel Vehicles (RFVs). The GCR's asymmetric design created a regulatory loophole that was ultimately exploited by manufacturers, leading to a perverse increase in the emissions of subsidized vehicles (Huse, 2026).[9]

The GCR was passed in Parliament and announced to the public in late March 2007, and effectively started in April 2007, thus before the global crisis. Given its design (and timing, see below), it was closer in spirit to the US hybrid subsidy (Berestenau and Li, 2011) than the "Cash for Clunkers" program.[10]

To better appreciate its design, note that the GCR was a policy within the Swedish National

Mercedes Benz S-Class, has to attain fuel economy of 46 mpg by 2025.

[7]These countries are Austria, Cyprus, France, Germany, Greece, Italy, Ireland, Luxembourg, The Netherlands, Portugal, Romania, Slovakia, Spain and the UK, see AEA Technology (2012).

[8]Formally, energy consumption has to be no greater than the equivalent of 9.2 liters/100 km using gasoline or 9.7 m3/100 km using gas (typically CNG, compressed natural gas); electric cars are considered green if their consumption is no greater than 37 kWh/100 km. Emission thresholds in Sweden apply to individual cars rather than to the brand level as in the US market. This emission threshold is equivalent to about 193 gCO2/mile, thus being already more stringent than the 250 gCO2/mile CAFE standard to take place from 2016 in the US. For perspective, the 2012 Porsche 911 Carrera, emits 205 gCO2/km.

[9]This strategic repositioning was made possible by the low cost of the FFV technology; the necessary hardware modifications, primarily consisting of ethanol-compatible fuel system components and recalibrated engine software resulted in a marginal manufacturing cost of well under USD 200 per vehicle (U.S. Department of Energy, 2008).

[10]The program was originally scheduled to run until December 2009. The Swedish Government initially allocated 250mn SEK (about USD 37.5mn, with 50mn SEK to be spent in 2007 and 100mn SEK to be spent on each of 2008 and 2009) to it, but the program was expanded twice: in Spring 2008, an additional 240mn SEK was made available whereas in fall 2008 both an additional 325mn SEK and the shortening of the program (to June 2009) were announced. Thus, the GCR amounted to a cost of 815mn SEK (about USD 122mn), slightly more than the amount spent in the first three years of California's Clean Vehicle Rebate Project, see https://energycenter.org/clean-vehicle-rebate-project/rebate-statistics.

Climate Policy Bill passed in 2005, which not only ratified the Kyoto Protocol, but also aimed to meet the 2003 EU Biofuels Directive regarding targets for the use of biofuels.[11] According to such fuel standard, 5.75 percent of all transport fuels in the EU should be met by biofuels by 2010.

A few features emerged as a consequence of the GCR's design. First, it had a broad impact on the new car market, in that vehicles benefiting from the GCR (*green cars*) commanded an unprecedented market share of 26.5 percent in 2008, the first calendar year after its inception, as compared to a pre-GCR market share of 6 percent in 2006.

Second, the share of AFVs in the Swedish car fleet turned out to be among the highest in the EU, together with Italy, see Figure 2. Among AFVs, the leading segment was that of flexible-fuel vehicles (FFVs).[12] FFVs were the main gainers following the GCR commanding about 15 percent of registrations in 2008, while CNG and electric vehicles have never commanded more than 1 percent of the market, see Huse and Lucinda (2014).[13]

In sum, the GCR is thus a policy instrument designed to address two targets. It can moreover be seen as a particular case of more general policies defined by a tuple (τ_r, τ_a, R), where R denotes the amount of the rebate paid out to consumers (in thousands of SEK) and τ_r (τ_a) is the emission threshold below which a regular (alternative) fueled vehicle is eligible for the rebate.

3 Theoretical Framework

To ground the empirical analysis theoretically, we develop a stylized model that captures the key features of the Swedish vehicle market: strong within-segment substitution and asymmetric compliance costs across vehicle types. The model is built on the same class of nested logit demand as our empirical specification and operates on the same objects – nest shares, inclusive values, and within- versus cross-nest substitution – so that its predictions can be confronted directly with the estimates. It deliberately simplifies the nesting, however: the theory uses a single green-versus-conventional split to isolate the mechanisms cleanly, whereas the estimated model (Section 6) nests by market segment and adds random coefficients. We argue below that this simplification is conservative – the frictions the theory identifies are, if anything, stronger in the richer empirical structure. We associate a testable implication with each result and defer proofs to Appendix A.

3.1 Setup

Consumers choose among J differentiated vehicles grouped into two nests: a *conventional* nest G_c containing vehicles ineligible for the rebate, and a *green* nest G_g containing eligible vehicles, plus an outside option $j = 0$ (purchasing a used car or not purchasing). The green nest contains two sub-types of vehicles: low-emission regular fuel vehicles (RFVs, with emissions $e_{rfv} \leq \bar{e}_{rfv}$) and alternative fuel vehicles (AFVs, with emissions $e_{afv} \leq \bar{e}_{afv}$), where $\bar{e}_{rfv} < \bar{e}_{afv}$ are the emission eligibility thresholds to the rebate, reflecting the policy's asymmetric design.

[11] In addition to the GCR, other measures were put in place in order to incentivize the purchase of more fuel-efficient cars. For instance, a number of Swedish municipalities have introduced rebates on parking fees for green cars, but the definition of what a green car consists of differs across municipalities.

[12] The main alternative fuel in Sweden is ethanol (E85), a 85:15 mixture of ethanol and gasoline in which the latter works as a lubricant and helps start the engine. In Winter, the ratio is adjusted to 75:25 to mitigate the risk of freezing.

[13] The price of an FFV is slightly higher than that of a comparable gasoline model. FFV engines piggy-back on the standard (Otto cycle) gasoline technology and allow to seamlessly switch between gasoline and ethanol, which may explain their swift adoption by Swedish consumers.

Figure 2: **Share of AFVs in the Car Fleet for Selected EU Countries**

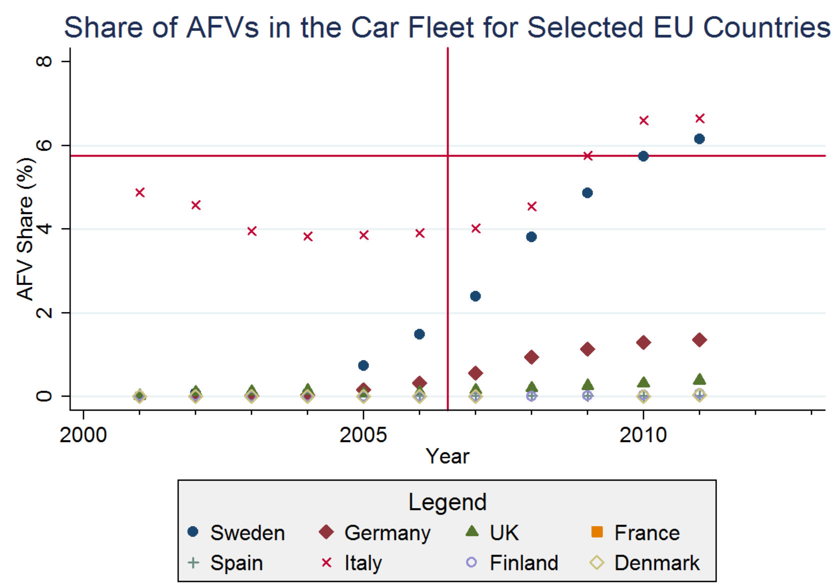

Note: Data based on all (by consumers and businesses) newly-registered vehicles from European Environmental Agency (2014). Countries reported in addition to Sweden have the full series of average CO2 emissions available since 2001. Among these, we selected the ones with the five biggest economies within the EU in terms of nominal GDP in 2013, plus Nordic countries Denmark and Finland. The vertical line is drawn between the years 2006 and 2007, when the Swedish Green Car Rebate was introduced (see below). The horizontal lines represent the CO2 emission targets of 120 and 130 gCO2/km.

Consumer i's utility for vehicle j in nest $h \in \{c, g\}$ is:

$$u_{ij} = \delta_j + \zeta_{ih} + (1 - \sigma)\epsilon_{ij}$$

where $\delta_j = x_j\beta - \alpha\tilde{p_j} + \xi_j$ is the mean utility of product j, with $\tilde{p_j} = (p_j - R \cdot \mathbf{1}[j \in G_g])$; $R \geq 0$ is a uniform rebate applied to all eligible green vehicles, lowering their effective price. Thus the rebate enters mean utility as αR for each $j \in G_g$. Here $\sigma \in [0, 1)$ is the nesting parameter governing within-nest correlation, ζ_{ih} is a nest-specific shock, and ϵ_{ij} is i.i.d. Type I extreme value. The outside option yields utility $u_{i0} = \epsilon_{i0}$, normalized to $\delta_0 = 0$. This is precisely the nested logit structure estimated in Section 6.

The standard nested logit choice probabilities yield the market share of vehicle j in nest h:

$$s_j = s_{j|h} \cdot s_h$$

where the within-nest share is $s_{j|h} = \exp(\delta_j/(1 - \sigma))/D_h$, $D_h = \sum_{k \in G_h} \exp(\delta_k/(1 - \sigma))$ is the inclusive value of nest h, and the aggregate nest share is:

$$s_h = \frac{D_h^{1-\sigma}}{1 + D_c^{1-\sigma} + D_g^{1-\sigma}}$$

where the 1 in the denominator corresponds to the outside option.

To evaluate the policy targets, we must define the internal composition of the nests. Let $S_{AFV}(R) \equiv \sum_{j \in AFV} s_{j|g}(R)$ denote the share of AFVs strictly within the green nest, and let $\bar{E}_g(R) \equiv \sum_{j \in G_g} e_j s_{j|g}(R)$ denote the average CO2 emission rate of the green nest. Similarly, let $\bar{E}_c$ denote the average emissions of the conventional nest.

3.2 Propositions

Proposition 1 (The Tinbergen Conflict). *A single uniform rebate R cannot simultaneously satisfy an arbitrary CO2 emissions target and an arbitrary renewable fuel standard. Both targets reduce to linear functions of the same scalar – the aggregate green nest share – making them perfectly collinear under a uniform instrument.*

The intuition of the result is as follows. Since the rebate R enters the utility of every green vehicle identically, it does not alter their relative attractiveness. Thus, the rebate leaves the internal composition of the green nest unchanged. As shown in the proof, the within-nest AFV share (S_{AFV}) and the average green emission rate ($\bar{E}_g$) are invariant to R.

Since the within-nest composition is invariant to R, the other possibility is that the rebate R shifts the aggregate market share of the green nest. Since both the RFS target and the fleet-wide CO2 target are driven by this single aggregate expansion margin, they become perfectly collinear. Satisfying the RFS target at a given level uniquely dictates the CO2 outcome, and vice versa. We thus conclude that, a single instrument R cannot simultaneously attain two independently specified targets, formalizing the Tinbergen rule.

Testable Implication 1: Across the set of simulated policies, the policy parameters that minimize average CO2 emissions will systematically differ from those that maximize AFV market share. We test this tension in Section 7.

Proposition 2 (Demand Inertia via Niche Positioning). *The power of the rebate to shift consumers from conventional to green vehicles is governed by the baseline market share of the green* *nest, not by the within-nest correlation parameter σ. When green vehicles occupy a niche market segment, the rebate functions largely as an inframarginal transfer with weak aggregate effect.*

Because the rebate applies uniformly across all green vehicles, it shifts the inclusive value of the entire green nest one-for-one. As shown in the proof, the nesting parameter σ cancels out entirely from the cross-nest substitution derivative. As a result, the green nest share behaves exactly as a standard logit, meaning the marginal effect of the rebate on the green nest share is strictly bounded by $s_g(1 - s_g)$.

This effect approaches zero as $s_g \to 0$. In the Swedish market, the pre-policy green share was approximately 6 percent, giving a theoretical maximum marginal effect of only $0.06 \times 0.94 \approx 0.056$ percentage points per unit of rebate. The sizeable demand inertia documented empirically in Section 5 therefore reflects the niche status of green vehicles – driven by the limited portfolio of attractive models – rather than a structural property of the nesting correlation.

Testable Implication 2: Even large rebates will produce modest shifts in the total green vehicle share. Conditional on purchasing a green vehicle, consumers will substitute primarily within model families (documented by the transition matrices), but the aggregate pull from conventional to green vehicles will be weak. We test both predictions in Sections 5 and 6.

Proposition 3 (Directed Technical Change and Supply Frictions). *Under asymmetric emission thresholds, if the fixed cost of meeting the strict RFV threshold substantially exceeds the cost of AFV conversion, profit-maximizing manufacturers predominantly expand AFV product lines, limiting the supply of genuinely low-emission vehicles.*

Let manufacturer m's profit from introducing a vehicle via compliance pathway $j \in \{afv, rfv\}$ be $\Pi_j(R) = (p_j - c_j)D_j(R) - K_j$, where K_j is the fixed cost of R&D and retooling. The policy applies the identical rebate R to both pathways. Converting an existing gasoline engine to accept ethanol requires minimal engineering (K_{afv} is low), whereas meeting the strict 120 gCO2/km RFV threshold requires significant powertrain innovation ($K_{rfv} \gg K_{afv}$).[14] As documented in Section 5, available FFV and RFV variants commanded similar prices, implying comparable gross margins. Under these conditions, $\Pi_{afv}(R) > \Pi_{rfv}(R)$, and manufacturers strictly prefer the AFV pathway.

If K_{rfv} is sufficiently large that $(p_{rfv} - c_{rfv})D_{rfv}(R) < K_{rfv}$ for all R within the program's budget, then $\Pi_{rfv}(R) < 0$ and the RFV pathway is never chosen. This yields a corner solution in which all eligible product introductions are AFVs – vehicles that satisfy the lenient 220 gCO2/km threshold but offer minimal CO2 improvement over the conventional fleet.

Testable Implication 3: Manufacturers will expand their FFV product lines rather than invest in genuinely low-emission conventional models. We document in Section 5 that no new low-emission RFV model below 120 gCO2/km was introduced during the GCR period, while FFV variants of existing high-emission models proliferated.

Relationship to the empirical specification. As noted at the outset, the theory uses a two-nest (green versus conventional) structure to isolate the mechanisms, whereas the empirical RCNL model of Section 6 nests vehicles by market segment and adds random coefficients, so that green

[14] For perspective, TNO, IEEP, and LAT (2006) estimate that the manufacturer costs to reduce emissions from the 2008-2009 baseline average of 140 gCO2/km down to the 2012 target of 120 gCO2/km at roughly EUR 1,700 per vehicle (Figure 3.10). The report also estimates a series of cubic polynomial specifications based on engineering data (Table 3.11). The estimate to reduce emissions from 220 ($\bar{e}_{afv}$) to 120 ($\bar{e}_{rfv}$) gCO2/km is between EUR 3,630 and 22,200, with a central estimate of EUR 8,800. This is one to two orders of magnitude of converting a gasoline vehicle into an FFV, estimated to be approximately EUR 150 (U.S. Department of Energy, 2008).

and conventional vehicles coexist within each nest. The framework's nesting parameter σ corresponds to the empirically estimated ρ. This difference has two implications. First, Proposition 1's cancellation result – that within-nest composition is invariant to R – does not hold literally in the segment-based nesting, since the rebate shifts within-segment demand toward green variants. However, the Tinbergen conflict manifests even more strongly in the richer structure: within each segment, the rebate directs demand toward high-emission FFVs (potentially beneficial for the RFS but not for CO2, depending on fuel choice), while the strong within-segment substitution we estimate ($\hat{\rho} = 0.744$, Section 6) severely limits cross-segment substitution toward the small, genuinely low-emission RFVs that would improve CO2 outcomes. The two targets thus pull along different margins that a uniform R cannot independently control. Second, the segment-based nesting reinforces Proposition 2: the high $\hat{\rho}$ estimate implies that consumers overwhelmingly substitute within their preferred size class, confining the rebate's reach regardless of its generosity. The simplified framework therefore provides a conservative lower bound on the frictions the policy faces.

Each proposition yields a testable implication that we confront with the data in the sections that follow: Proposition 3's supply-side prediction against the pre-policy product line (Section 5), Proposition 2's inframarginal prediction against the demand estimates and counterfactuals (Sections 6–7), and Proposition 1's Tinbergen conflict against the full policy grid (Section 7). We flag each confirmation in place.

4 Data

Our primary analysis relies on **vehicle registration data** from *Vroom Stockholm*, a consulting firm. The data on privately owned vehicles (i.e., those eligible for the rebate) is recorded at the monthly frequency from January 2004 to December 2009. An observation is a combination of month, brand, model and fuel type.

We enrich this dataset by merging it with several other sources. First, **vehicle characteristics** from the consumer guides *Nybilsguiden* issued yearly by The Swedish Consumer Agency (*Konsumentverket*). For every model available on the Swedish passenger car market the information includes characteristics such as fuel type, engine power and size, number of cylinders, weight, fuel economy (city driving, highway driving and mixed driving, with testing made under EU-determined driving cycle) and list prices. We deflate the vehicle tax, car and fuel prices using the Consumer Price Index from Statistics Sweden. For car prices and vehicle tax we use the yearly average with 2009 as the base year and for fuel prices the monthly average with December 2009 as the base month.

Second, we add **market-level fuel data** recorded at the monthly frequency at the national level. Recommended retail fuel prices for gasoline, diesel and ethanol are obtained from the Swedish Petroleum and Biofuels Institute (SPBI). These prices were deflated using the same CPI used for car list prices.

Third, we use **mileage data** from the Swedish Motor Vehicle Inspection Company (*Bilprövningen*) on yearly average distances covered by Swedish passenger cars. For every year, we observe average odometer readings for cars of 3, 5, 7, 8 and 10 years of age disaggregated by brand, model, fuel type and body type.[15] With this data in place, we adopt the hedonic approach of Huse and Lucinda (2014) to obtain the lifetime utilization of vehicles.[16]

[15]That is, we do not observe micro level data on mileage. As a result, we are unable to estimate a joint model of vehicle choice and utilization as in, e.g., (Goldberg, 1996).

[16]The preferred specification regresses km/year on fixed-effects for fuel, year, fuel-age, fuel-year, and brand-model-

Fourth, we use **pollutants data** from the UK Vehicle Certification Agency on both global (CO2) and local (PM10 and NOx) pollutants.[17] The pollution values are monetized in our cost-benefit analysis using estimates from Parry et al. (2014) and converted to SEK using an exchange rate of 7 SEK per USD.

An important issue arises when combining data sets. As is typical in the literature, one important issue arising when merging characteristics and registration datasets is that the former is observed at a more disaggregated level than the latter. Despite being more aggregated than car characteristics, the level of aggregation in registrations is still more refined than standard market level datasets in that we observe sales for different versions at the fuel level. For each combination of year-brand-model-fuel we use characteristics from the baseline version, i.e., the lowest priced model. Importantly, given the relatively small number of green versions (typically one or two per model), aggregation issues for these models essentially vanish. For merging the resulting dataset with the pollutant (CO2, PM and NOx) data, we employed a two step process. The first step was based on direct matching based on three variables – car brand, model and year. We started by matching all models with same brand, model and year, and then we matched the models with the same brand and model, with one year difference, with the engine capacity in cc to be used in case of a tie; and then with two year difference, with the engine capacity in cc to be used in case of a tie.

5 Key Frictions in the Swedish Vehicle Market

To understand why the GCR and its variants are unlikely to succeed, we first present two foundational, pre-policy stylized facts about the Swedish market: (1) significant supply-side constraints in the availability of green vehicles, and (2) strong demand-side inertia driven by consumer preferences.

Supply-side. Figure 3 focuses on emissions of CO2 and local pollutants from the supply-side. In Panel A, we select the best-selling vehicle models in the Swedish new car market in the period 2004-2006, which amounts to the models commanding at least 2.5 percent market share in a given year. For each such model, we examine the range of CO2 emissions (measured in gCO2/km) disaggregated by fuel, reporting its minimum and maximum levels in Panel A. For instance, the Ford Focus has a gasoline entry-level version emitting 159 gCO2/km, the best-selling Volvo V70 has a diesel entry-level version emitting 152 gCO2/km, whereas the Saab 9-3 has a diesel entry version emitting 147 gCO2/km.

The key finding in Panel A is that most best-selling models have entry-level versions emitting slightly less than 160 gCO2/km, which become eligible for the rebate in case of policies where 160 gCO2/km is the emission threshold for RFVs. That is, by purchasing a model of the desired market segment but opting for a version with slightly lower CO2 emissions (typically done by downsizing to a smaller/less powerful engine, in the absence of technological improvements), a consumer can

segment fixed-effects, with an R-squared of 0.9383.

[17] Available from https://carfueldata.vehicle-certification-agency.gov.uk/ For models whose pollutants are missing, we impute those after merging the data as follows (see description of the merge below). We estimate regression models for each pollutant where the outcome variable was the pollutant emission level and covariates engine capacity (in CC), engine power (in HP), car weight (in kg) and fuel consumption, in addition to brand and year fixed-effects. The regression models are in C.1 in Appendix C. The comparison between the pollutants estimated by this method and the ones merged from the UK data is also in Appendix C, in Table C.2.

take advantage of the rebate. Alternatively, *if substitution occurs mostly within brands if not models,* it is reasonable to expect consumers to mostly downsize with respect to engine (and thus CO2 emissions), conditional on vehicle size, or even model.

We also display the range of CO2 emissions for AFVs, see Panel B. Recalling that CNG and hybrid electric vehicles (HEVs) never (jointly) commanded more than 1 percent in the Swedish market, we focus on FFVs.[18] Not only was there no low-emission FFV marketed pre-GCR (and no such model was introduced during the GCR), but the models with the lowest CO2 emission levels emitted in excess of 160 gCO2/km. As a result, one would expect that technology-biased policies would not improve CO2 emission levels unless there were substantial technological advances in the industry.

Finally, Panel C reports RFVs with emission of up to 125 gCO2/km (recall we define low-emission models as those emitting no more than 120 gCO2/km). The key finding from inspecting the models listed is that – even if some of the models appear among the top-25 best-selling models – these are much smaller vehicles than the best-selling models listed in Panel A. That is, there is little scope for substitution in terms of similarly-sized vehicles from high- to low-emission models, which prevents the policies to have a stronger effect on the new vehicle market.

At the root of our findings lies the fact that while such policies focus on the demand side, there is no instrument directed at the supply side, i.e., the product lines of firms. Thus, the effectiveness of such policies relies on the existence of a portfolio of low-emission and/or AFV models, which may take carmakers some time to introduce following such a policy, and perhaps even longer following normal market developments.

Testable Implication 3, confirmed. No new RFV model below 120 gCO2/km entered during the program, while flex-fuel variants of existing high-emission models proliferated – the corner solution of Proposition 3, in which manufacturers expand the cheap compliance pathway rather than engineer genuinely efficient vehicles.

Demand-side. A key determinant of a rebate's success is its ability to induce consumers to switch from high-emission to low-emission vehicles. However, individual-level registration data from the pre-policy period reveals significant demand-side inertia. Figure 4 displays transition matrices for vehicle size, brand, and model for the period 2003-2006.[19] The strong concentration of observations along the diagonal reveals substantial state-dependence. For instance, consumers show remarkable brand loyalty and, even when switching, tend to stay within a narrow set of models from the same manufacturer or trade up in size. This evidence strongly suggests that consumers are not comparing all cars on the market but are instead operating within highly constrained choice sets. This inertia presents a challenge for any policy seeking to encourage drastic shifts in purchasing behavior, such as switching from a large Volvo station wagon to a small low-emission alternative.

Zooming into Figure 4, Panel A displays transition probabilities across vehicle sizes, as per the definition of the US NHTSA market segments; vehicles may belong to segments mini, light, compact, medium or heavy to which we assign numbers 1-5, respectively. The take-away from Panel A is that consumers tend either remain in the same segment or trade-up when it comes to vehicle size, as documented by the larger circles above the diagonal of the transition matrix.

Panel B displays transition probabilities across brands available in the Swedish market (with

[18]Actually, CNGs and HEVs never commanded more than 0.5 percent in the period 2004-2009.

[19]All panels display transition matrices constructed using individual-level registration data. For each of them, the size of the circle in entry (i, j) is proportional to the transition probability from state i to state j, p_{ij}.

Figure 3: **CO2 Emissions of Selected Models**

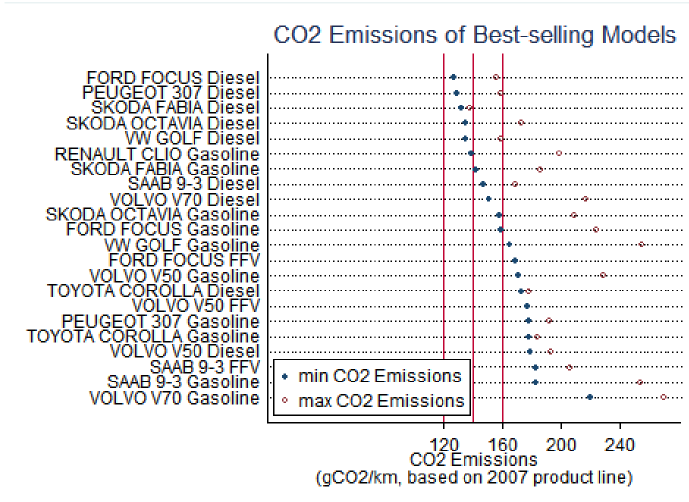

Panel A

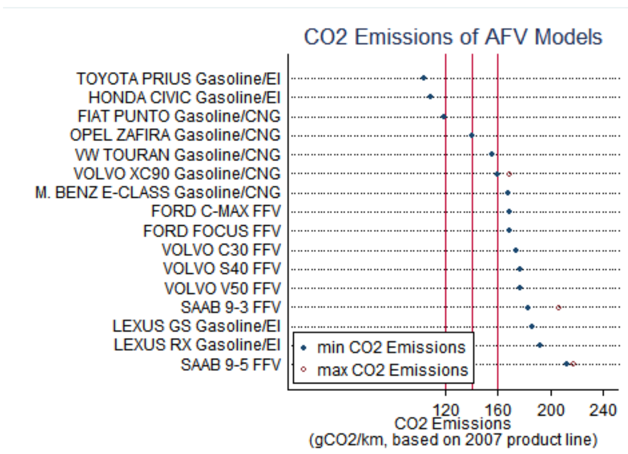

Panel B

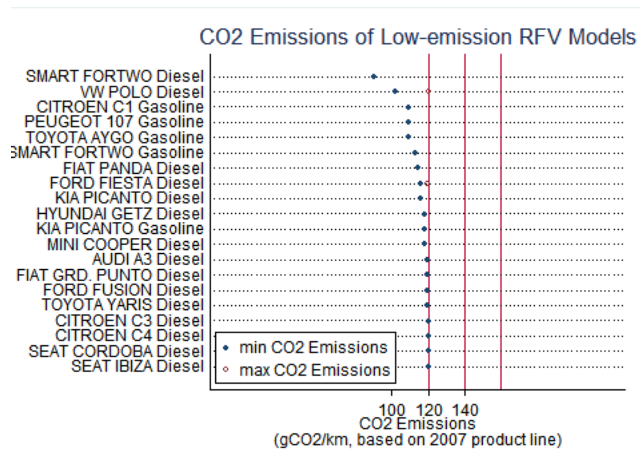

Panel C

Note: The vertical lines denote CO2 emission levels of 120, 140 and 160 gCO2/km.

brands displayed in alphabetical order). Following the dominance of the circles along the diagonal of the transition matrix, the take-away from Panel B is that consumers exhibit substantial state-dependence when it comes to brand choice, often with $p_{ij} \geq 0.5, i = j$. That is, despite choice sets with over 30 brands, consumers are mostly loyal to the current brand upon the purchase of a new vehicle.

The extent of state dependence can be further appreciated when looking at transition matrices at the model level. Panel C considers the top-25 best-selling models in the Swedish market in the period 2003-2006 (with models displayed according to market share). The take-away from Panel C is that state dependence may be stronger than suggested by the evidence in Panel B: not only there is substantial state dependence it the brand level, but if there is any switching, this will be between at most a handful of models within a brand. For instance, Volvo customers will mostly switch between the station wagons V70 (the best-selling car in Sweden during the sample period) and the slightly smaller V50; VW (the second best-selling brand in Sweden) customers tend to switch among the Golf, Passat and Polo models, the same holding for customers of Skoda, Toyota, Renault and Peugeot models.

Taken together, the above findings suggest that the extent to which consumers are likely to switch is quite limited; consumers are more likely to forego a rebate than to purchase a low-emission variant of the model they currently own.

6 Estimation

6.1 Demand

Setup. Our ultimate goal is to simulate market outcomes for different policy designs which nest the GCR. We take the perspective of a policymaker in that we estimate our model for the period *pre-policy* and use the estimated parameters to quantify the effects of alternative policies.

To do this, we first need to understand consumer preferences from their pre-policy choices. We employ a random coefficients nested logit model (RCNL), which allows us to achieve two objectives. First, it captures the reality that consumers are heterogeneous, thus valuing attributes like brand, size, and fuel economy differently. Second, it accounts for the fact that some products are closer substitutes than others (e.g., two sedans are closer competitors than a sedan and a large SUV). By estimating this model on pre-policy data, we recover the distribution of consumer preferences, which we then use as the foundation for our counterfactual policy simulations.

Model Specification. We estimate the demand for cars using discrete choice models for market level data, following Berry, Levinsohn and Pakes (1995). The starting point is a microeconomic model of rational behavior for individual households which is aggregated to generate market demands. Consumers buy at most one of the products available on the market and, if so, the one yielding the highest utility among the available products. The econometrician does not observe individual choices, only market level data, i.e. prices, quantities and a set of characteristics for each of the J products available on the market for a number of periods (we suppress the index t below to avoid clutter). These "inside" products are indexed by $j = 1, ..., J$, and the outside good, the option to buy a used car or to not buy a car at all is represented by $j = 0$. Define the

Figure 4: **Transition Matrices of Market Segments, Brands and Models**

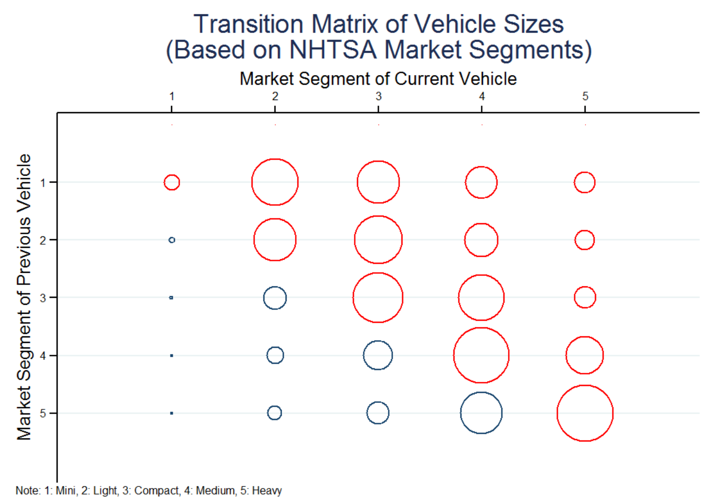

Panel A: Transition Matrix of Market Segments

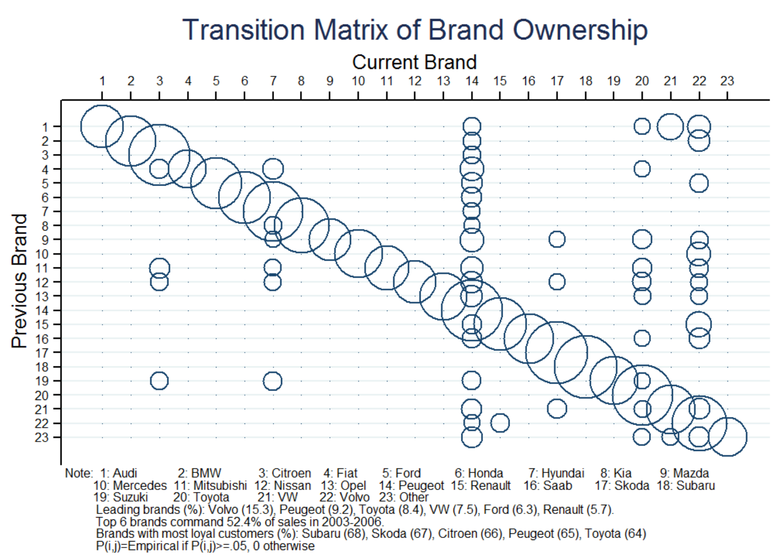

Panel B: Transition Matrix of Brand Ownership

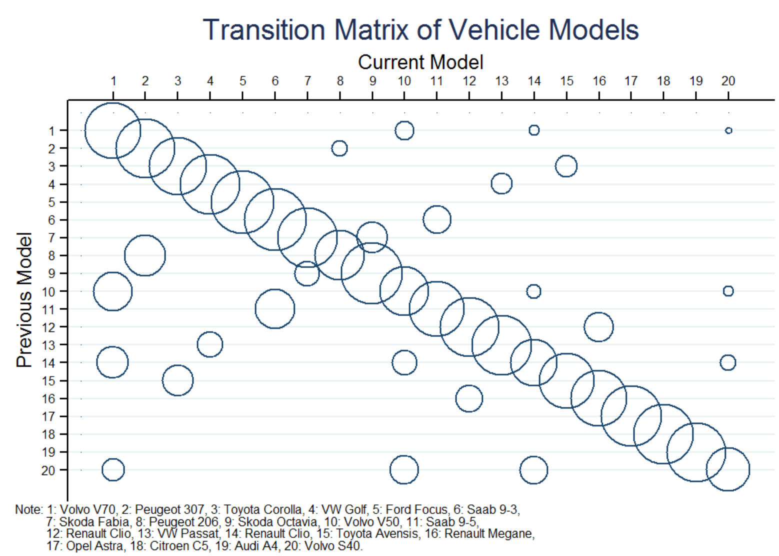

Panel C: Transition Matrix of Model Ownership

Note: Panel C is constructed based on the top-25 best selling vehicle models in the Swedish market in the period 2003-2006.

conditional indirect utility of individual i when consuming product j as

$$u_{ij} = V_{ij} + \varepsilon_{ij} = \alpha_i p_j + \sum_{k=1}^{K} x_{jk}\beta_{ik} + \xi_j + \varepsilon_{ij},\ i = 1, ..., I;\ j = 1, ..., J$$

where p_j are (likely endogenous) vehicle prices, x_{jk} are exogenous observed product characteristics such as engine power and vehicle size while ξ_j are characteristics observed by the market participants but not the econometrician (such as quality, style etc). We decompose the individual coefficients as $\beta_{ik} = \overline{\beta}_k + \sigma_k v_{ki}$, where $\overline{\beta}_k$ is common across individuals, v_{ki} is an individual-specific random determinant of the taste for characteristic k, which we assume to be Normally distributed, $(v_{1i}, ..., v_{Ki})' \sim \mathcal{N}(0, \Sigma)$, and σ_k measures the impact of v on characteristic k. Following (Berry, Levinsohn and Pakes, 1999), the term α_i measures the consumer's distaste for price increases and its distribution varies with income; we assume that α_i has a time-varying distribution following the empirical distribution function of income of Swedish consumers. If y_i is a draw from the empirical distribution function of income, then $\alpha_i = \alpha/y_i$, where α is the parameter to be estimated, with the price sensitivity thus being modeled as inversely proportional to income.[20] Finally, ε_{ij} is an individual and option-specific idiosyncratic component of preferences, assumed to follow the distributional assumptions of the Nested Logit model, allowing product within-group correlation. Formally, ε_{ij} can be written as:

$$\varepsilon_{ij} = \zeta_{ig} + (1 - \rho)\epsilon_{it}$$

where ζ_{ig} has the distribution so as ε_{ij} is extreme value and ϵ_{ij} is also extreme value. The ρ parameter is the so-called "nesting parameter" and approximates the level of correlation of product preferences between products in the same group, and is bounded between zero and one. If ρ tends to zero, the correlation tends to zero and the resulting model converges to the standard Random Coefficients Logit.

Since consumers may decide not to buy a new car, the specification of the demand system is completed with an outside good yielding conditional indirect utility $u_{i0} = \xi_0 + \sigma_0 v_i + \varepsilon_{i0}$, where ε_{i0} is an individual market and time specific idiosyncratic term and v_i is an individual specific component reflecting heterogeneity in tastes.

Our estimation approach represents a pragmatic modeling approximation to actual consumer choice behavior in the industry, assuming away a number of important features in the car market. For instance, cars are durable products, so current ownership of a car is likely to affect the current demand for automobiles, see Hendel and Nevo (2006) and Gowrisankaran and Rysman (2012). Moreover, primary and secondary car markets (new and used cars) do coexist and should ideally be modeled jointly, see e.g., Bento et al. (2009). On the supply-side we also assume a static framework, thus not accountig for product innovation and consistent with a short-run modelling of the new vehicle market.

Identification. Besides the exogenous characteristics, we use the set of "BLP instruments", following Berry, Levinsohn and Pakes (1995). That is, we use a set of polynomial basis functions of characteristics (Engine Power/Weight Ratio and Engine Size (CC)) of other products of the same firm, besides fuel consumption in liters/100km, the standard EU measure of fuel consumption. BLP instruments implicitly assume a form of localized competition among products, and this

[20] As in Berry, Levinsohn and Pakes (1999), this can be obtained as a first-order Taylor series approximation to the "Cobb-Douglas utility function" often used in the literature.

seems consistent with anecdotal evidence for the automobile industry, characterized by a number of market segments and highly differentiated products. We also included the number of products per nest as an additional instrument for the nested logit specifications.

Estimates. We consider demand specifications with the following characteristics: price, engine power (measured in HP, 1 HP = 0.7457 kw) and engine size (measured in cubic centimeters, CC). We also include fixed-effects for automatic transmission, time (month), market segment, fuel segments (as well as their interaction with a *low-emission* dummy). Finally, we interact a measure of the cost of driving (in km/SEK, inspired by BLP's miles per dollar measure) with fuel segment indicators.[21] Consumer heterogeneity is introduced in the RC logit specification via the interaction of price and income draws and onto the transmission coefficient via 500 antithetic pairs of random draws of the standard Normal distribution. (The Appendix discusses a number of alternative specifications also experimented with.) Finally, we used the nesting structure based on market segments (compact, small, medium, large and luxury).

We report alternative demand estimates in Table 6.1. Our estimation period is January 2004-December 2006. Specification 1 ("OLS") reports the estimates obtained when price is assumed to be exogenous, i.e. it is a standard OLS logit regression with market level data. Columns 2 and 3 report IV logit and RC logit estimates, respectively, whereas Columns 4 and 5 report Nested Logit and RC Nested Logit estimates using the instruments described in the previous section. The F-statistic of the excluded instruments of the first-stage regression of price on the exogenous characteristics and the instruments used is 45.92, suggesting that the instruments are not weak.

Specification 1 features a negative and significant price coefficient of -0.004, positive and significant estimates for both engine size and power, and a negative and significant coefficient for the automatic transmission indicator. Estimates for the market segment indicators suggest that consumers prefer larger vehicles (baseline is the compact segment) whereas estimates for the fuel segment indicators suggest that consumers prefer FFVs and low-emission gasoline vehicles. The cost of driving variables load negatively and are statistically significant for vehicles able to operate on gasoline, except for CNGs, which command a negligible market share and are typically purchased by high utilization consumers. Unfortunately, own-price elasticities are typically less than one in absolute value, which is inconsistent with the assumption of profit-maximizing firms.

Accounting for price endogeneity as in Specification 2 results in a steeper demand curve. In fact, the estimated price coefficient increases a fourfold as compared to its OLS counterpart. An immediate result from controlling for price endogeneity is the improvement on the estimates of own-price elasticities; the 1st and 3rd quartiles are given by 4.7 and 2.4, respectively. Moreover, most coefficients load with the same signs as before.

Introducing consumer heterogeneity in Column 3 results in a price coefficient of -9.104.[22] Although the price coefficient is significant, engine size and the automatic transmission indicator become insignificant. Finally, the random coefficient for automatic transmission is not statistically significant. The heterogeneity induced by the income draws results in own-price elasticities in the range 3.9-5.0 for the 3rd and 1st quartiles; such values imply markups in the range 19-29 percent and are broadly in line with standard estimates for European markets using market level data, e.g. Goldberg and Verboven (2001) report elasticities in the range 3-6 in their Table 6. The remaining estimates of the RC logit specification are also broadly in line with economic theory and the lit-

[21] We have also experimented with product fixed-effects, with unsatisfactory results. This is likely to be due to the relatively short sample period coupled with moderate product entry and exit.

[22] Recall that this is not directly comparable with the logit specifications in (1) and (2).

Table 1: **Demand Estimates**

	OLS	IV	RC Logit	Nested Logit	RC Nested Logit
Price ('1000 SEK)	−0.004***	−0.015***		−0.006***	
	(0.000)	(0.002)		(0.001)	
$\alpha (= \alpha_i / y_i)$			−9.104***		−2.102***
			(2.544)		(0.647)
Engine Power (HP)	0.005***	0.026***	0.012***	0.010***	0.005***
	(0.001)	(0.003)	(0.003)	(0.002)	(0.001)
Engine Size (CC)	0.000**	0.001***	0.000	0.000**	−0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Automatic Transmission (dv)	−0.388***	−0.769***	1.930	−0.276***	−0.148*
	(0.052)	(0.084)	(3.602)	(0.052)	(0.074)
$\sigma_{Transmission}$			4.699		0.147
			(6.254)		(0.798)
Market Segments					
Small (dv)	−0.011	−0.168**	0.677	−0.071**	0.126**
	(0.064)	(0.071)	(0.510)	(0.028)	(0.048)
Medium (dv)	0.200***	−0.044	1.632*	−0.033	0.307***
	(0.062)	(0.075)	(0.892)	(0.030)	(0.089)
Large (dv)	0.531***	0.295***	2.525*	0.078*	0.512***
	(0.071)	(0.085)	(1.112)	(0.036)	(0.131)
Executive/Luxury (dv)	0.428***	0.606***	2.926*	0.201***	0.602***
	(0.085)	(0.096)	(1.276)	(0.049)	(0.161)
Fuel Segments					
Diesel	0.548	1.159*	−0.290	0.341	0.123
	(0.471)	(0.506)	(1.479)	(0.221)	(0.529)
FFV	12.388***	13.075***	10.558***	4.506***	3.079***
	(1.683)	(1.648)	(2.001)	(0.906)	(0.929)
CNG	0.314	2.803*	2.132	0.877*	0.499
	(1.336)	(1.556)	(1.769)	(0.506)	(0.672)
Gasoline Low-emission	9.873***	6.722**	4.631	3.940**	2.703
	(2.535)	(2.356)	(5.327)	(1.687)	(1.994)
Gasoline	4.907***	6.073***	3.539*	2.034***	1.169*
	(0.471)	(0.529)	(1.583)	(0.374)	(0.572)
Cost of Driving x Fuel Segments					
Cost of Driving	−0.006	0.027*	−0.014	0.010*	0.002
	(0.012)	(0.014)	(0.035)	(0.006)	(0.012)
Cost of Driving x Diesel	−0.008	−0.024*	0.017	−0.007	−0.001
	(0.012)	(0.013)	(0.036)	(0.005)	(0.012)
Cost of Driving x FFV	−0.187***	−0.213***	−0.149***	−0.074***	−0.048**
	(0.032)	(0.032)	(0.042)	(0.016)	(0.018)
Cost of Driving x CNG	−0.018	−0.063**	−0.028	−0.021**	−0.010
	(0.018)	(0.021)	(0.028)	(0.008)	(0.013)
Cost of Driving x Gasoline Low-emission	−0.181**	−0.112*	−0.091	−0.080*	−0.060
	(0.062)	(0.057)	(0.144)	(0.041)	(0.048)
Cost of Driving x Gasoline	−0.047***	−0.078***	−0.022	−0.026***	−0.012
	(0.012)	(0.014)	(0.036)	(0.007)	(0.012)
ρ				0.660***	0.744***
				(0.048)	(0.038)
N	9,751	9,751	9,751	9,751	9,751
GMM Crit. Fcn.	0.000	50.170	46.125	16.702	18.718
1st Q.	−1.072	−4.700	−5.030	−4.893	−4.548
Median	−0.817	−3.584	−4.639	−3.642	−3.775
Mean	−0.914	−4.006	−4.297	−4.119	−3.907
3rd Q.	−0.568	−2.489	−3.885	−2.485	−3.157

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

erature, i.e. *ceteris paribus*, consumers value engine power and larger vehicles (as proxied by the market segment dummies).

Specification 4 estimates are roughly comparable to Specification 2 whereas Specification 5 estimates are broadly in line with the RC logit, with a nesting parameter of 0.744 that indicates strong within-segment substitution, consistent with the transition matrices of Figure 4, market segments estimates increasing in size, and slightly less elastic yet more dispersed demand elasticities.

Testable Implication 2, confirmed. The high nesting parameter, together with a pre-policy green share near six percent, is exactly the configuration under which Proposition 2 predicts the rebate operates as a near-inframarginal transfer: consumers substitute overwhelmingly within their size class, so cross-segment pull toward green vehicles is weak regardless of the rebate's size. The counterfactuals in Section 7 bear this out, and the welfare decomposition later quantifies it – only two to three percent of subsidized sales are marginal.

Supply-side assumptions. Our counterfactuals hold both the product line and vehicle prices fixed at their pre-policy levels: we do not model manufacturers' price responses to the rebate, nor endogenous product entry. Three considerations make this the appropriate frame for the question we ask. First, the exercise is deliberately *ex ante*. We place ourselves in the position of a policymaker in 2006 and ask what the pre-policy market structure implied about the rebate's likely effects. That structure – the models on offer, their prices, and the substitution patterns among them – is precisely what a policymaker would have taken as given, and it is what our estimates recover.

Second, the assumption is milder in this market than the label suggests. We observe manufacturer list prices, which adjust at an annual frequency; and in Sweden these track transaction prices more closely than in markets such as the United States, because vehicles are typically ordered in advance and delivered with a lag of weeks or months, leaving little room for the dealer discounting and on-the-lot incentives that open a wedge between list and transaction prices elsewhere (Huse and Lucinda, 2014). That is, there is little high-frequency price variation within the year and list prices are good proxies for transaction prices. Sweden is moreover a small, open market for a global industry: manufacturers set model prices and product lines with an eye to the European market as a whole, so a short-lived national rebate is unlikely to move either quickly.

Third, to the extent that firms *would* respond, the direction of the bias favours the policy, so our fixed-price simulations give it its best case. The most plausible supply response to a consumer rebate is for firms to raise prices and capture part of it; such pass-through would attenuate consumer switching further, weakening the policy relative to our estimates. Our headline conclusion – that the rebate cannot reach its targets – is therefore conservative with respect to this channel. The one supply response we do characterize is qualitative and works in the same pessimistic direction: Proposition 3 predicts that firms expand the cheap, high-emission compliance pathway rather than engineer genuinely low-emission models, a product-line adjustment that our fixed pre-policy choice set omits and that would only compound the failure we document.

This assumption has one important implication for the welfare analysis of Section 7.4: since prices are fixed and marginal costs are never re-solved, the welfare account is a partial-equilibrium accounting of the rebate's incidence rather than a re-optimized equilibrium. As we show there, this is precisely what makes the producer-surplus and VAT components recoverable from the simulation output without additional supply-side assumptions.

7 Assessing Alternative Policies

We evaluate different policies encompassing the actual GCR. We let a policy be defined by a tuple (τ_r, τ_a, R) where R denotes the rebate paid out to consumers upon the purchase of a green vehicle (in thousands of SEK), and τ_r, τ_a denote the CO2 emission thresholds (measured in gCO2/km) below which, respectively, RFVs and AFVs qualify as green cars and are thus eligible for the rebate. Rebate values range from zero to $30,000$ SEK in $5,000$ SEK increments, so $R = 0, 5, ..., 30$ whereas $\tau_r = 120, 160, 200, 240$ and $\tau_a = 120, 160, 200, 220, 240$. The set of policies we examine is detailed in Table 2.[23]

For the sake of clarity, we report results of alternative policies in two graphs per panel, including the actual policy. On the left-hand side, we report policies for which τ_r is fixed at the 120 gCO2/km threshold; this is to highlight how changing the degree of asymmetry in the GCR affects different variables. In particular, one may be interested in asking whether an asymmetric policy would have been at all necessary. On the right-hand side, we maintain the assumption of symmetry – or technology-neutrality – when it comes to CO2 emission thresholds and examine what would have happened as the (common) CO2 emissions threshold changes.

Table 2: **Description of Alternative Policies**

Asymmetric Policies (Left-hand side)	**Symmetric Policies (Right-hand side)**
$(120, 220, R)^*$	$(120, 220, R)^*$
$(120, 120, R)$	$(120, 120, R)$
$(120, 160, R)$	$(160, 160, R)$
$(120, 200, R)$	$(200, 200, R)$
$(120, 240, R)$	$(240, 240, R)$

Note: This table describes the policies examined in this section. CO2 emissions range from 120 gCO2/km, which amount to 46.9 (51.8) mpg running on gasoline (diesel), to 240 gCO2/km, which amount to 23.5 (25.9) mpg running on gasoline (diesel). The benchmark set of policies encompassing the actual GCR $(120, 220, 10)$ are marked with an asterisk $(^*)$ and are reported in both left- and right-hand side graphs to facilitate comparison to other policies.

Model Validation. Before examining the full grid of counterfactual policies, we assess the model's predictive accuracy by comparing its predictions under the actual GCR policy $(120, 220, 10)$ to observed outcomes during April 2007–June 2009. Table 3 reports this comparison for some key variables. The fit is close on every behavioural margin: predicted and actual green-car and AFV market shares agree to within six percent, total new-vehicle sales to within one percent, and the sales-weighted mean CO2 of new registrations almost exactly (166.4 versus 166.5 gCO2/km). This is notable given that the model is estimated entirely on pre-policy data and is asked to predict outcomes two to five years out, under a policy absent from the estimation sample. Reassuringly, the model reproduces the composition and environmental outcomes of the policy.

[23] In earlier versions of the paper we considered rebate values of up to 50,000 SEK, but the current one restricts the analysis of up to 30,000 SEK, with qualitative results unchanged. This decision was taken because the Swedish Government initially allocated a relatively low value of 250mn SEK to the program (which was expanded twice and eventually was allocated 815mn SEK in total). As shown below, even rebate values of up to 30,000 SEK imply a total cost of several times the amount allocated to the program.

Table 3: Model Validation: Predicted vs. Actual Outcomes under the GCR $(120, 220, 10)$

Variable	Predicted	Actual	Ratio
Green-Car (Eligible) Market Share (%)	20.9	19.7	1.06
AFV Market Share (%)	15.0	14.2	1.05
Mean CO2, New Registrations (gCO2/km)	166.4	166.5	1.00
Total New Vehicle Sales (units)	273,736	271,824	1.01
AFV Sales (units)	41,021	38,631	1.06

Note: Predicted values are from the RCNL model estimated on pre-policy data (January 2004–December 2006) and simulated under the actual GCR parameters, pooled over the full GCR period (April 2007–June 2009). Actual values are computed from observed registration data over the same period. Ratio = Predicted / Actual.

7.1 Program Costs

We start by examining the monetary cost of the program. Panel A in Figure 5 displays the costs of alternative policies alongside the budget allocated to the actual program (815mn SEK; horizontal dotted line), which we take as the relevant ex-ante budget constraint against which a policymaker in 2006 would have assessed feasibility.[24] A few lessons emerge from Panel A, beyond the intuitive monotonic patterns in both rebate amounts and emission thresholds.

First, a number of policies – especially symmetric ones – would by far exceed this budget and thus not be politically feasible. To fix ideas, compare a $(120, 160, R)$ policy in the left-hand graph, with an estimated total cost of approximately 800mn SEK at a rebate of 25,000 SEK, with its symmetric counterparts, which typically cost more even at rebates as low as 15,000 SEK. A number of lax policies in Panel B would cost in excess of 2bn SEK and could not feasibly be implemented. We thus conclude that, treating the 815mn SEK allocation as a hard budget constraint, few of the policies we consider would have been feasible.

We also examine both VAT revenues and the net revenues of alternative policy configurations in Figure 5.[25] VAT revenues increase in the rebate amount and are higher under stricter (lower-threshold) policies than under laxer ones. Because the share of consumers choosing the outside good is largely insensitive to the policy, the total volume of VAT-able transactions is roughly constant, so VAT revenue is governed by the price mix of what is sold rather than by market expansion. Laxer thresholds extend eligibility toward near-universal coverage, so a larger fraction of transactions occurs at rebate-reduced net prices, eroding the VAT base; under stricter thresholds most consumers continue to purchase higher-priced, non-eligible vehicles at full price, sustaining VAT revenue.

The comparison of Panels A and B raises the question of the net monetary effect of alternative policies, reported in Panel C of Figure 5. The net fiscal cost (program cost net of VAT revenue) rises with the rebate amount and with the laxity of the threshold.

[24]The 815mn SEK figure is the cumulative appropriation for the program, including estimated tax effects, across successive budget bills. Realized rebate disbursement was lower, at 713mn SEK, reflecting both the program's curtailment in June 2009 and its design feature that rebates were paid six months after purchase, so that payouts on late-window registrations fell into the 2009 and 2010 budget years (Riksrevisionen, 2020). For the feasibility analysis the relevant benchmark is the allocation, since this is the budget envelope known to the policymaker ex ante.

[25]VAT is the main tax levied on vehicles, charged at a 25 percent rate on net vehicle prices.

Figure 5: **Monetary Costs and Benefits**

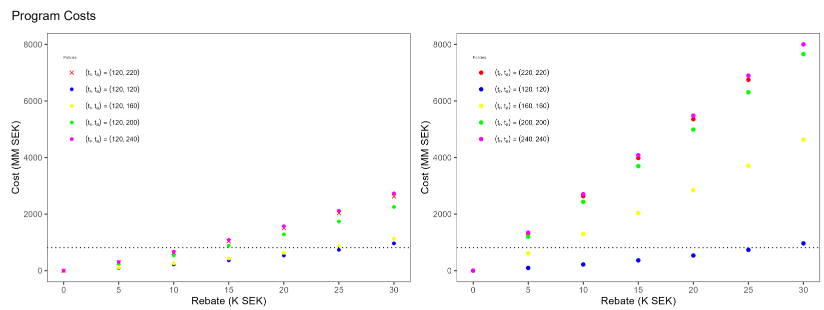

Panel A

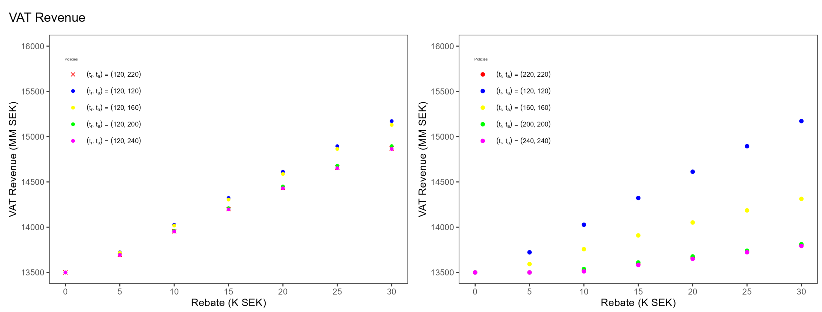

Panel B

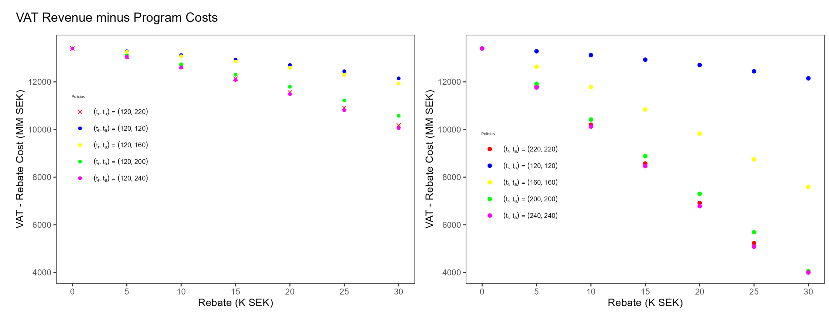

Panel C

Note: Horizontal dotted line in Panel A displays the budget allocated to the actual GCR (815mn SEK).

7.2 Policy Objectives: CO2 Emissions and Renewable Fuel Standard

We examine the impact of alternative policies on the two main targets of relevance to the GCR, namely the new vehicle average CO2 emissions of 120 gCO2/km and the EU renewable fuel standard (RFS) of 5.75 percent.

CO2 Emissions. To compute average CO2 emissions, we estimate the quantity sold of every car model for each simulated policy and combine it with information on its CO2 emissions (see the Appendix for details), calculating then the average CO2 emissions of new registrations – see Figure6. The results for the overall fleet are displayed on Panel A, whereas those of the RFV and AFV fleets are displayed on Panels B and C, respectively.

Panel A highlights how average CO2 emissions are monotonic in the rebate amount, but not in the emission thresholds. Among asymmetric policies, the family $(120, 160, R)$ attains the largest decrease in sales-weighted average CO2 emissions, tracked closely by the strict family $(120, 120, R)$. Among symmetric policies, the minimizer is the moderate family $(160, 160, R)$. However, regardless of policy design, the reductions in average CO2 emissions are modest at best and the average CO2 emissions target of 120 gCO2/km is far from being attained.

Panel B focuses on RFVs and highlights how – again – average CO2 emissions are mostly monotonic in rebate amounts but not monotonic in emission thresholds, even though the stricter policy of the family $(120, 120, R)$ dominates laxer ones. When it comes to symmetric policies, a moderate policy of the family $(160, 160, R)$ would result in the lowest average CO2 emissions for new vehicles.

Panel C focuses on AFVs. Here, average CO2 emissions are minimized for policies of the type $(120, 200, R)$ and $(200, 200, R)$ in the asymmetric and symmetric cases, respectively, but the average CO2 emission levels are consistently above 120 gCO2/km, thus far from attaining the corresponding target.

The relatively low response in demand to changes in policy parameters comes from the existence of few RFV and AFV low-emission models. In particular, Panels B and C suggest that emissions appear largely insensitive to changes in emission thresholds, likely due to the limited substitution towards more fuel-efficient vehicles.

Summing up, Figure 6 shows a few stylized features. First, the steeper gradient of average CO2 emission changes across policies for the symmetric policies. This is a direct consequence of the more pronounced changes of policy parameters in this case.

Second, the uniformity along the axis of the rebate amount, i.e., no crossing of curves when changing the rebate amount, R.

Finally, an intriguing non-linearity emerges from our simulations. Policies with a moderate threshold around 160 gCO2/km often produce lower overall emissions than stricter ones. This "sweet spot" effect likely occurs because the 120 gCO2/km threshold is too restrictive conditional on the product lines, applying only to smaller vehicle segments that buyers of popular, larger models are unwilling to switch to, regardless of the rebate. In contrast, the 160 gCO2/km threshold is attainable for "downsized" versions of these popular models (e.g., the same model with a less powerful engine). This encourages marginal but widespread substitution within preferred model families, leading to a greater overall emissions reduction than a stricter policy that induces almost no substitution from the mass market.

Renewable Fuel Standard. As one of the objectives of the Swedish policies is to attain the RFS, Figure 7 examines the sales of AFVs. In the absence of detailed fuel sales data, we posit different scenarios in terms of fuel use among owners of FFVs under which such target could be attained and compare them to the sales of AFVs.

Under the most favorable (and unrealistic) assumption that all AFVs run exclusively on ethanol, attaining a 5.75 percent renewable share of the Swedish fleet (about 4mn vehicles, i.e. 230,000 AFVs) would require promoting the sale of roughly 170,000 new AFVs over April 2007–June 2009, starting from a pre-policy stock of about 60,000 (a 1.48 percent share in 2006). This 170,000 figure is therefore a lower bound on what any policy must achieve: it is the requirement in the best case for the policy, when every AFV is driven entirely on the renewable fuel.[26][27] Increasing the fleet of AFVs does not necessarily mean that less fossil fuels will be consumed as consumers may choose the cheapest between gasoline and ethanol if driving a FFV (Huse and Lucinda, 2014; Salvo, 2018; Huse, 2026). That is, fuel choice becomes another dimension to consider when evaluating transport policies. More recently, Huse (2018) shows the low substitutability of ethanol and gasoline among Swedish motorists. Swedish policymakers have worked under the assumption that FFV owners use ethanol 75 percent of the time in their analyses. Thus, if the alternative fuel standard is to be attained by newly-registered vehicles, a first approximation is that AFVs have to command a share of at least 7.7 percent ($\cong 5.75/0.75$), or 308,000 units, of the Swedish car fleet. Therefore, any simulated policy meeting the RFS has to promote the sales of at least 238,000 AFVs.

Looking at AFV sales implied by alternative policies, the central point is that even the lower bound of 170,000 AFVs – the requirement under the most favorable fuel-use assumption – is not met by any policy we consider; the realistic requirement of roughly 238,000 AFVs is far further out of reach. That is, the results in Figure 7 show the little traction AFVs gain under all the policies considered. Notice, for instance, that in the best-case scenario, sales of AFVs would not even reach the 100,000 mark. Among all designs, AFV sales are maximized by the laxest asymmetric family $(120, 240, R)$, yet even there sales fall well short of the lower-bound target. Intuitively, strict asymmetric policies have hardly an effect on the sales of AFVs given the lack of low-emission AFVs in the Swedish market.[28]

Attainment of the "Twin Targets". The above results for CO2 emissions and the RFS suggest that meeting the "twin targets" of CO2 emissions and the RFS would be unrealistic without significant technological progress, combined with higher rebate amounts which would likely make the cost of the policy go well above budget.

At the root of these findings lies the fact that there are very few low-emission models being marketed when the GCR is introduced (and even thereafter when looking at actual data), and one would either need an aggressive policy targeting the supply-side or allow carmakers more time to react to the policy accordingly.

Testable Implication 1, confirmed. The policies that minimize sales-weighted average CO2 – the asymmetric (120,160,R) and symmetric (160,160,R) families – are disjoint from those that maximize AFV sales – the laxest asymmetric family (120,240,R). The CO2-minimizing and RFS-maximizing designs do not overlap, so no single rebate reaches both targets. This is the Tinbergen

[26]The RFS is a target on the renewable share of fleet-wide fuel consumption, whereas our policy acts on new-vehicle sales. We translate the fuel-share target into an equivalent number of AFVs using the mileage and utilization assumptions detailed in Appendix C, so that the 170,000 and 238,000 figures should be read as the new-vehicle inflow consistent with the fleet-level target under those assumptions.

[27]Allowing for realistic fuel switching makes the requirement substantially more demanding.

[28]For perspective, recall that however small, such figures are still above to the market shares of about 2 percent experienced by hybrid electric vehicles (HEVs) under a similar policy conducted in the US in the same period (Berestenau and Li, 2011).

conflict of Proposition 1, recovered directly from the simulated policy grid rather than assumed.

What drives these results is the interaction of the demand- and supply-side frictions documented in Section 5, which together confine the rebate to reshuffling demand rather than redirecting it. On the demand side, the transition matrices in Figure 4 show that consumers overwhelmingly repurchase within their current size class, brand, and often model family: the mass of each matrix lies on or near the diagonal, and cross-segment transitions — precisely the switches a green rebate would need to induce — are rare. The demand estimates quantify the force behind this inertia. Using Specification (5) of Table 6.1, the implied willingness to pay for moving up one size class (to a large or executive vehicle, the most conservative estimate) is one order of magnitude larger than the willingness to pay for a marginal unit of engine power. Since the only way to lower a vehicle's CO2 within a segment is to downsize the engine, and the only low-emission vehicles available sit in smaller segments (Section 5), the rebate asks consumers to give up the size they value most in exchange for the attribute they value least. A transfer of 10,000-30,000 SEK is simply too small to overcome a size premium of this magnitude. The result is that even the most generous rebates move fleet composition only at the margin: the CO2 target requires substitution the market will not supply, and the RFS target requires an AFV volume that the same inertia, compounded by the scarcity of low-emission AFVs, keeps out of reach.

7.3 Unintended Consequences: Local Pollutants

While the Swedish GCR aimed primarily at attaining the EU-mandated CO2 emission targets, the importance of local pollutants was already clear and put into regulation at the EU level. In fact, the Euro emission standards had already been introduced in the early 1990s, with the Euro 1 regulaton on local pollutants applying to vehicles registered from 1993.

The local pollutants regulated by the Euro standards include NOx, PM, CO and HC. The regulation in place throughout most of our sample period was Euro 4 (1st January 2006 to 31st December 2010), which had tightened PM emission levels for diesel vehicles and both CO and NOx emission levels for gasoline vehicles.

Given data availability (see Data section and Appendix), in what follows we focus on NOx and PM emissions. Panels A and B of Figure 8 display, respectively, NOx and PM levels for the policies we simulate. The key pattern is that the policy families that perform best on CO2 perform worst on local pollutants. Relative to the actual GCR baseline, the CO2-minimizing asymmetric family $(120, 160, R)$ raises NOx emissions from about 7.0 to 7.2 thousand tonnes at a rebate of 10,000 SEK (roughly +3 percent), widening to about 6.8 to 7.25 thousand tonnes at 30,000 SEK (roughly +7 percent); PM rises correspondingly by about 2 to 5 percent over the same range. The loosest asymmetric policies $(120, 220, R)$ and $(120, 240, R)$, by contrast, leave local pollutants roughly unchanged or slightly lower. The symmetric families display the same tension, with the CO2-minimizing $(160, 160, R)$ family generating the largest local-pollutant increases.

The intuition for these results becomes clearer when comparing CO2 emissions at the vehicle level to both NOx and PM emissions, which we do in Figure 9. While both relationships are nonlinear, the association between the emissions of CO2 and the local pollutants is mostly positive where best-selling vehicles tend to locate, be it pre-GCR (product line for year 2007) of post-GCR (product line for year 2009). Moreover, as we document in Appendix C, diesel vehicles play a crucial role: policies with higher sales of diesels result in higher concentrations of local pollutants.

The mechanism is the switch from gasoline to diesel vehicles. Policies that effectively reduce average CO2 emissions do so partly by shifting demand toward diesel vehicles, which have lower CO2 emissions per kilometer than their gasoline counterparts but substantially higher NOx and

Figure 6: **CO2 Emissions for New Vehicles**

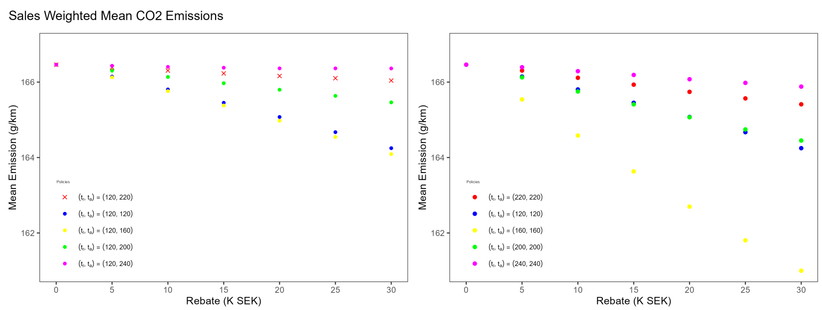

Panel A

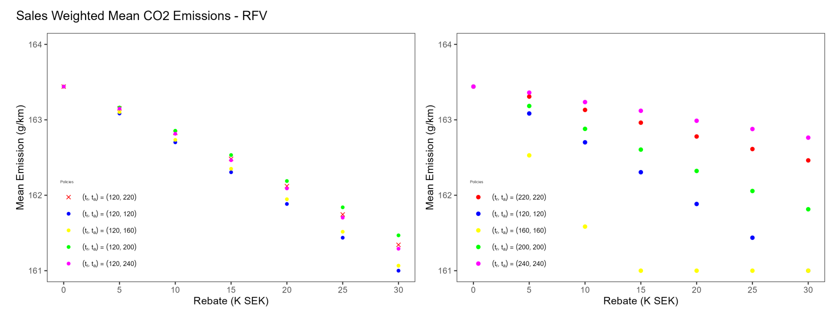

Panel B

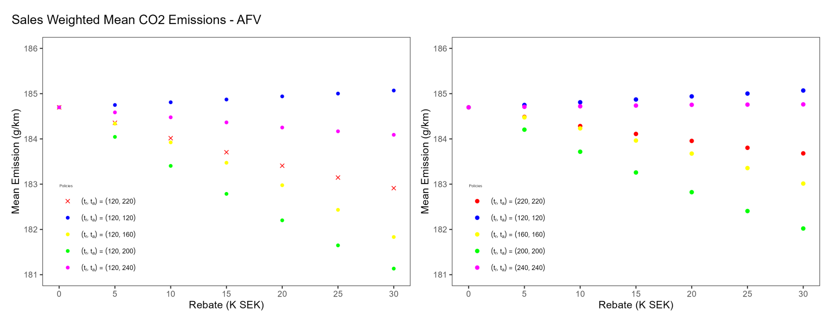

Panel C

Figure 7: **Renewable Fuel Standard**

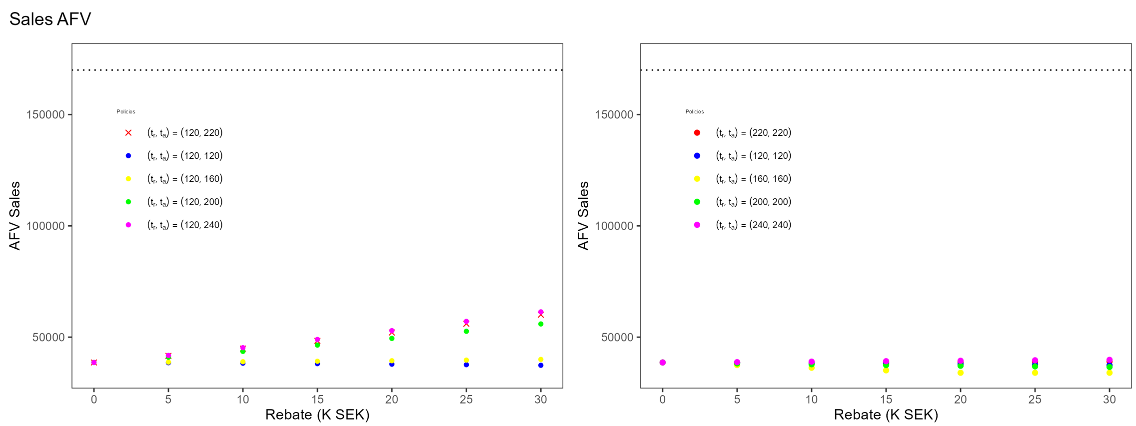

Note: The horizontal dotted line displays the minimum new vehicle sales required to meet the RFS of 5.75 percent.

PM emissions. Consequently, the families of policies that succeed in lowering fleet-average CO2 – namely the $(120, 160, R)$ and $(160, 160, R)$ families – achieve this partly by encouraging diesel adoption, which simultaneously worsens local air quality. This is a classic second-best outcome: in a market where diesel vehicles are undertaxed relative to their local pollution externality, a CO2-focused subsidy exacerbates the pre-existing distortion.

7.4 Cost and Welfare

The preceding subsections showed that no policy in our grid reaches either target, and that the designs which lower average CO2 per vehicle do so partly by shifting demand toward higher-NOx diesels. We now put these strands on a common monetary footing and ask the question a cost-benefit analyst would: what does each policy cost, and what does it deliver in environmental terms? The answer, which we develop below is that across the entire policy space, more spending buys more environmental damage. This reframes how the program's welfare consequences should be read.

Cost versus environmental damage. Figure 10 plots, for every policy in our grid, the program cost against the monetized *reduction* in total environmental damage – the sum of CO2, NOx and PM valued at the unit damages of Parry et al. (2014). Two features stand out. First, every policy lies below the horizontal axis: relative to no rebate, each design *raises* total environmental damage. This is the monetary expression of the volume effect documented above – the rebate expands total sales, and the extra emissions from additional vehicles outweigh the modest fall in CO2 per vehicle, so total damage rises even as the regulated per-vehicle average falls. Second, the relationship is monotone in cost: more generous designs, which spend more, cause more damage. There is no interior design that spends money to *reduce* damage; the policy frontier runs the wrong way. Whatever else the rebate achieves, it does not purchase environmental improvement – it purchases environmental harm, and proportionally more of it the more is spent.

This is why we do not present a conventional cost-effectiveness or efficiency frontier. Such an object ranks policies by benefit purchased per krona and presumes there is a benefit to buy. Here there is none: the good the program was meant to deliver – lower total emissions – moves the wrong way under every design. The relevant summary is therefore not how efficiently the rebate abates, but that it does not abate at all, and that its welfare case must rest entirely on the surplus it transfers rather than on any environmental return.

Welfare decomposition. Table 4 puts this on a full welfare context, decomposing the change from the no-rebate baseline into the surplus accruing to consumers, the gross fiscal cost, and the monetized environmental damage. Because our counterfactuals hold prices and product lines fixed (Section 6), the exercise is a partial-equilibrium accounting of the rebate's incidence. We report two welfare measures that differ only in how they treat the rebate outlay. The government-budget measure counts the outlay as an unrecovered social cost, $W_1 = \Delta\text{CS} - \text{FC} - \Delta\text{Env}$. The total-surplus measure recognizes the outlay as a transfer: as shown in Appendix D, the rebate returns to the private sector almost in full, as producer surplus on additional sales and as incremental VAT, so that $\Delta\text{PS} + \Delta\text{VAT} \approx \text{FC}$ and the fiscal transfers cancel, leaving $W_2 = \Delta\text{CS} - \Delta\text{Env}$.[29]

[29] VAT is levied on the full list price the consumer pays at purchase; the rebate is a separate transfer, paid directly to the buyer and outside the VAT base. Producer surplus is recovered from simulation aggregates without estimating marginal costs; the derivation and the accounting identity are in Appendix D.

Figure 8: **Local Pollutants**

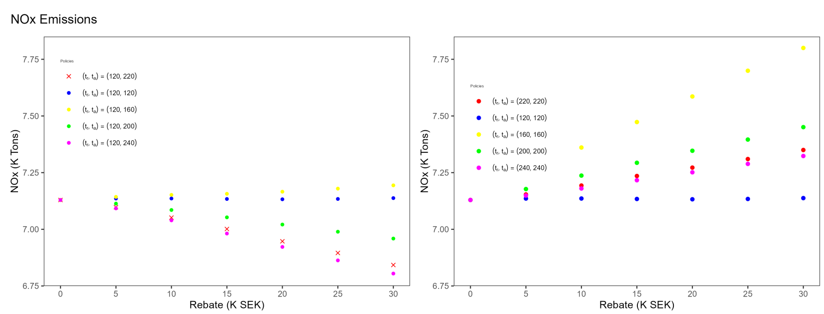

Panel A

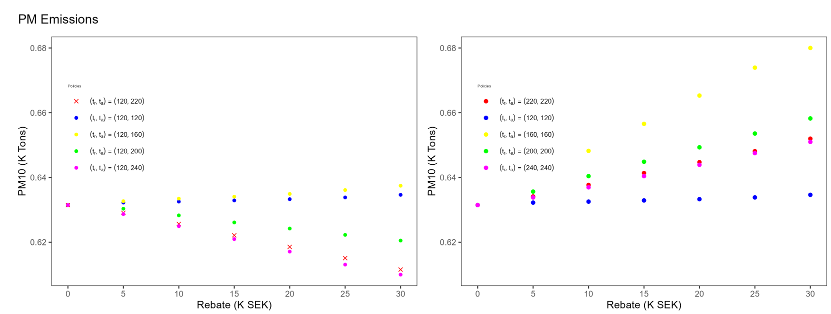

Panel B

Figure 9: **Unintended Consequences: CO2 vs. Local Pollutants**

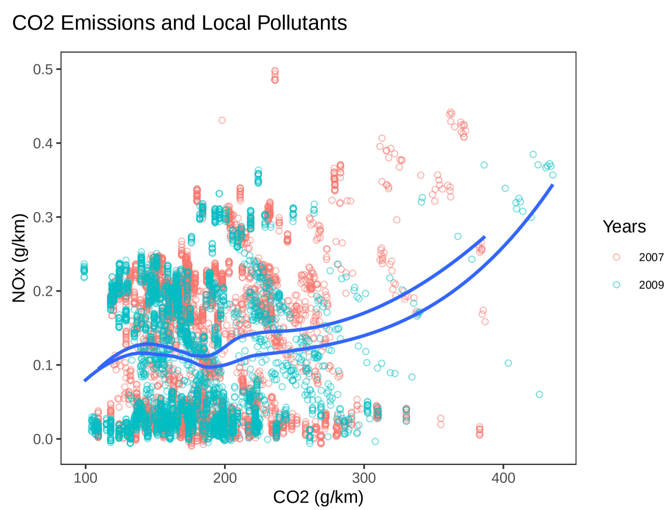

Panel A

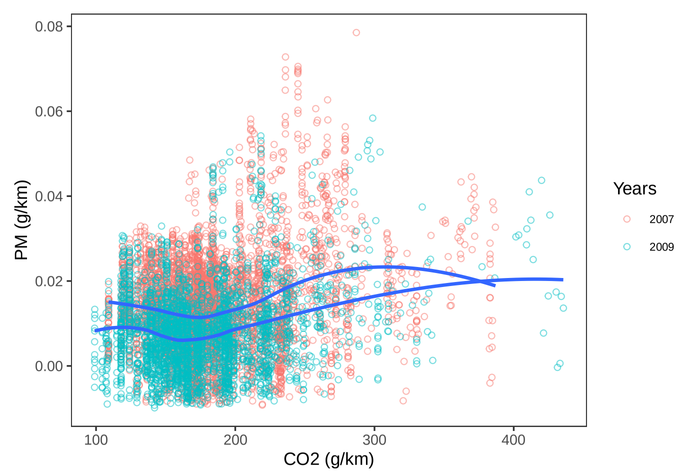

Panel B

Figure 10: **Program Cost versus Change in Total Environmental Damage**

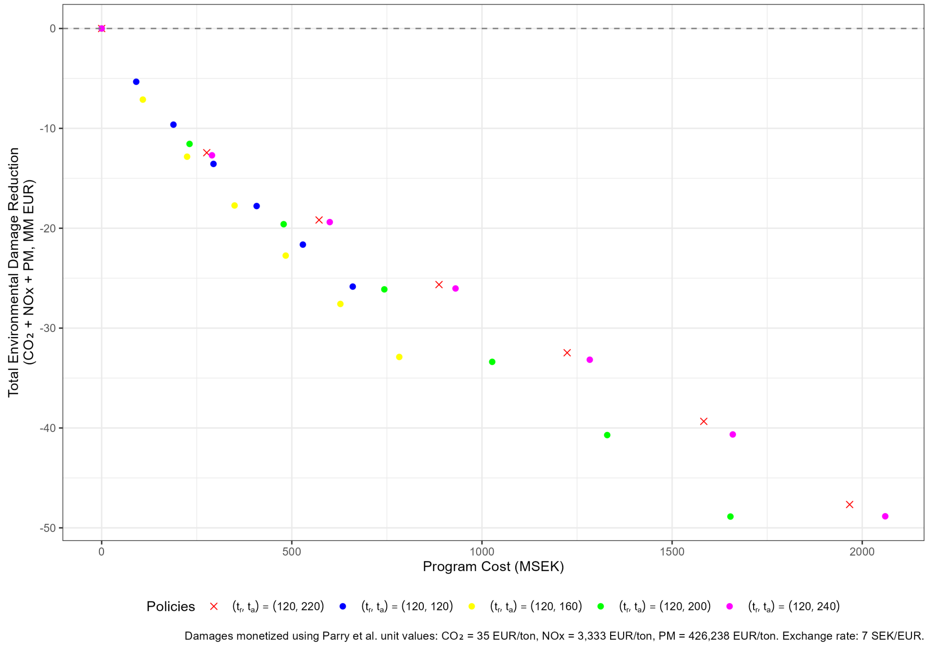

Note: Each point is one policy (τ_r, τ_a, R). The vertical axis is the change in monetized environmental damage (CO2 + NOx + PM) relative to the no-rebate baseline within each (τ_r, τ_a) pair; negative values denote *greater* damage. Damages are monetized using the unit values of Parry et al. (2014).

W_2 is the theoretically appropriate social criterion, since transfers should net out; W_1 answers the narrower question of whether the program is a sound use of public funds.

The two measures carry opposite signs, and the distance between them is the substance of the result. Under W_1, every policy is welfare-reducing: for the actual GCR, $W_1 = -56.7$ MSEK, with the deficit growing to roughly $-1{,}100$ MSEK for the largest symmetric design, because the consumer-surplus gain never covers the gross outlay. Under W_2, every policy is welfare-improving and substantially so: the actual GCR yields $W_2 = +514.9$ MSEK. The gap between them is exactly the rebate outlay, which W_1 treats as lost and W_2 recognizes as a transfer returned to the private sector. Whether the program "improves welfare" thus turns entirely on whether one counts a transfer as a cost.

Two features of the decomposition settle how it should be read, and both point the same way. First, on *either* measure the environmental term ΔEnv is positive in every row: consistent with Figure 10, the program adds to environmental damage rather than subtracting from it. Second, the W_2 gain is overwhelmingly an *inframarginal transfer*. Because incremental VAT is generated only by newly induced sales while the outlay is paid on all eligible units, the ratio ΔVAT/FC – 6 to 10 percent across our policies – bounds the share of rebated sales that are genuinely induced; for the actual GCR this implies that only some 2–3 percent of eligible purchases were marginal, the remaining outlay subsidizing buyers who needed no inducement. The measured welfare gain is thus not a return on an environmental investment but the accounting surplus of a transfer to consumers and producers, most of it flowing to purchases that would have occurred regardless. This is the fiscal counterpart of the niche-market inertia of Proposition 2.

Taken together, the figure and the table deliver a more nuanced read of the effects of the GCR than simply arguing that it "destroyed welfare." The program raises measured welfare only in the sense that any transfer raises the surplus of its recipients; it yields no environmental improvement – indeed it worsens total damage – and it does so while directing the overwhelming majority of its budget to inframarginal buyers. As an instrument for greening the fleet it fails; as fiscal policy it is an expensive and poorly targeted transfer. Neither reading depends on the supply-side response the design omits, which Proposition 3 predicts would only deepen the environmental failure.

Table 4: Welfare Decomposition

(τ_r, τ_a)	R	ΔCS	FC	ΔVAT	ΔPS	ΔEnv	W_1	W_2
Asymmetric policies $(\tau_r \leq \tau_a)$								
(120,220)	**10**	**534.1**	**571.6**	**58.7**	**512.9**	**19.2**	**–56.7**	**514.9**
(120,220)	30	1716.5	1966.8	156.8	1810.0	47.7	–298.0	1668.8
(120,120)	10	173.7	188.7	15.3	173.4	9.6	–24.7	164.1
(120,120)	30	562.8	660.2	37.8	622.4	25.8	–123.3	537.0
(120,160)	10	207.7	224.6	20.3	204.3	12.8	–29.7	194.9
(120,160)	30	671.7	782.6	46.5	736.0	32.9	–143.7	638.8
Symmetric policies $(\tau_r = \tau_a)$								
(160,160)	10	1132.9	1226.9	86.1	1140.8	65.1	–159.0	1067.8
(160,160)	30	3581.8	4084.5	250.0	3834.4	199.3	–702.0	3382.5
(220,220)	10	2472.9	2660.3	200.7	2459.5	96.5	–283.9	2376.4
(220,220)	30	7611.8	8403.0	648.8	7754.0	311.7	–1102.9	7300.1

Notes: All figures in MSEK, changes from the no-rebate baseline within each (τ_r, τ_a) pair. Bold row is the actual GCR. ΔVAT + ΔPS = FC in every row up to rounding, verifying the accounting identity. ΔEnv is the sum of monetized CO2, NOx and PM damages (positive = greater damage). $W_1 = \Delta\text{CS} - \text{FC} - \Delta\text{Env}$; $W_2 = \Delta\text{CS} - \Delta\text{Env}$.

Conclusion

This paper develops and tests a structural framework for why demand-side subsidies for dual-fuel vehicles fail to meet their environmental objectives. The framework, grounded in the random coefficient nested logit structure we estimate, identifies three reinforcing mechanisms: a Tinbergen conflict that prevents a single rebate from serving two targets; demand-side inertia that confines the rebate to within-segment reshuffling; and supply-side frictions that steer manufacturers toward the cheapest, highest-emission compliance pathway. Our ex-ante analysis of the Swedish Green Car Rebate confirms each prediction using only pre-policy data: the program's twin failure was foreseeable from the market's underlying structure.

The welfare analysis sharpens the lesson in a way we did not anticipate at the outset. Since the rebate expands total vehicle sales, it raises total emissions – and total monetized environmental damage – even though it lowers CO2 per vehicle; the designs best at improving the per-vehicle average are also those that most raise local pollutants by tilting demand toward diesel. The program does not buy environmental improvement inefficiently; it buys environmental harm. Whether it nonetheless "improves welfare" depends entirely on accounting convention: counting the rebate outlay as a social cost, welfare falls; recognizing it as a transfer that returns to producers and the treasury, welfare rises. Either way, the environmental account is negative, and the welfare gain that does appear is overwhelmingly an inframarginal transfer – roughly nine-tenths of the outlay subsidizes purchases that would have happened regardless. As environmental policy the rebate fails on its own terms; as fiscal policy it is an expensive and poorly targeted transfer.

These findings speak directly to the contemporary debate over subsidies for plug-in hybrid electric vehicles. The flex-fuel–ethanol setting of the GCR is structurally the PHEV–electricity setting of today's programs: both subsidize dual-fuel vehicles on optimistic assumptions about the share of driving on the cleaner energy source, and both face a growing performance gap as real-world utilization falls short of the regulatory assumption. Our analysis suggests such subsidies will underperform whenever three conditions hold – consumers prefer incumbent vehicle types, attractive green alternatives are scarce, and the instrument cannot distinguish genuine improvement from cheap compliance – conditions that characterized Sweden in 2007 and characterize segments of the European and North American markets today. The volume-versus-composition tension we document is a particular warning: a subsidy that improves per-vehicle ratings while expanding sales can raise total emissions, a possibility that evaluations focused on ratings alone will overlook.

Two limitations qualify these conclusions and mark avenues for future work. First, our ex-ante framework holds the supply side fixed. This is the appropriate lens for a short-lived policy in a small, open market, and it is conservative – pass-through and the product-line gaming that Proposition 3 predicts would only deepen the environmental failure – but it does not capture manufacturers' longer-run pricing and product responses, which merit their own structural treatment. Second, our model does not endogenize the fuel-choice margin – whether flex-fuel owners actually burn ethanol – which is central to the renewable-fuel target and to the PHEV analogy alike; incorporating it would sharpen the welfare implications of dual-fuel subsidies.

Ultimately, the analysis demonstrates that financial incentives alone cannot reshape markets defined by strong frictions, and that a demand-side subsidy layered on such a market can spend heavily to make the environment worse. Effective decarbonization of the car fleet requires pairing demand-side "pull" with supply-side "push" — stringent, technology-neutral standards that compel manufacturers to offer attractive low-emission vehicles in the segments consumers actually buy.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2012. "The Environment and Directed Technical Change." *American Economic Review*, 102(1): 131–166.

Adamou, Adamos, Sofronis Clerides, and Theodoros Zachariadis. 2014. "Welfare implications of car feebates: A simulation analysis." *Economic Journal*, 124(578): 420–443.

Allcott, Hunt, and Nathan Wozny. 2014. "Gasoline Prices, Fuel Economy, and the Energy Paradox." *The Review of Economics and Statistics*, 96(5): 779–795.

Anderson, Simon P., Andre de Palma, and Brent Kreider. 2001. "Tax incidence in differentiated product oligopoly." *Journal of Public Economics*, 81(2): 173–192.

Anderson, Soren T., and James M. Sallee. 2011. "Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards." *American Economic Review*, 101(4): 1375–1409.

Bento, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen, and Roger H von Haefen. 2009. "Distributional and Efficiency Impacts of Increased US Gasoline Taxes." *American Economic Review*, 99(3): 667–699.

Berestenau, Arie, and Shanjun Li. 2011. "Gasonline Prices, Goverment Support, and the Demand for Hybrid Vehicles in the United States." *International Economic Review*, 52(1): 161–182.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1999. "Voluntary export restraints on automobiles: Evaluating a trade policy." *American Economic Review*, 89(3): 400–430.

Berry, Steven T., James A. Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–890.

Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. 2013. "Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review*, 103(1): 220–256.

Dubé, Jean-Pierre, Günter J. Hitsch, and Peter E. Rossi. 2010. "State dependence and alternative explanations for consumer inertia." *The RAND Journal of Economics*, 41(3): 417–445.

Fershtman, Chaim, Neil Gandal, and Sarit Markovich. 1999. "Estimating the effect of tax reform in differentiated product oligopolistic markets." *Journal of Public Economics*, 74(2): 151–170.

Fullerton, Don, and Garth Heutel. 2010. "The General Equilibrium Incidence of Environmental Mandates." *American Economic Journal: Economic Policy*, 2(3): 64–89.

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

Goldberg, Pinelopi K. 1996. "Dealer Price Discrimination in New Car Purchases : Evidence from the Consumer Expenditure Survey." *Journal of Political Economy*, 104(3): 622–654.

Goldberg, Pinelopi Koujianou, and Frank Verboven. 2001. "The evolution of price dispersion in the european car market." *Review of Economic Studies*, 68(4): 811–848.

Gowrisankaran, Gautam, and Marc Rysman. 2012. "Dynamics of consumer demand for new durable goods." *Journal of Political Economyolitical Economy*, 120(6): 1173–1219.

Hendel, Igal, and Aviv Nevo. 2006. "Measuring the Implications of Sales and Consumer Inventory Behavior." *Econometrica*, 74(6): 1637–1673.

Huse, Cristian. 2018. "Fuel Choice and Fuel Demand Elasticities in Markets with Flex-fuel Vehicles." *Nature Energy*, 3: 582–588.

Huse, Cristian. 2026. "Fast, Furious, and Dirty: Gaming, Fuel Choice Behaviour, and the Welfare Cost of Dual-Fuel Subsidies." Working paper.

Huse, Cristian, and Claudio Lucinda. 2014. "The market impact and the cost of environmental policy: Evidence from the swedish green car rebate." *Economic Journal*.

Huse, Cristian, and Nikita Koptyug. 2022. "Salience and Policy Instruments: Evidence from the Auto Market." *Journal of the Association of Environmental and Resource Economists*, 9(2): 345–382.

IEA. 2022. *Global EV Outlook.* International Energy Agency.

International Energy Agency. 2025. "Global EV Outlook 2025." International Energy Agency, Paris. Available at https://www.iea.org/reports/global-ev-outlook-2025.

Li, Shanjun, Youming Liu, and Chao Wei. 2022. "The Cost of Greening Stimulus: A Dynamic Analysis of Vehicle Scrappage Programs." *International Economic Review*, 63(4): 1561–1594.

Parry, Ian WH, Mr Dirk Heine, Eliza Lis, and Shanjun Li. 2014. *Getting energy prices right: From principle to practice.* International Monetary Fund.

Plötz, Patrick, Steffen Link, Hermann Ringelschwendner, Marc Keller, Cornelius Moll, Georg Bieker, Jan Dornoff, and Peter Mock. 2022. "Real-World Usage of Plug-in Hybrid Vehicles in Europe: A 2022 Update on Fuel Consumption, Electric Driving, and CO2 Emissions." International Council on Clean Transportation and Fraunhofer ISI White Paper.

Reynaert, Mathias. 2021. "Abatement Strategies and the Cost of Environmental Regulation: Emission Standards on the European Car Market." *The Review of Economic Studies*, 88(1): 454–488.

Riksrevisionen. 2020. "Statliga åtgärder för fler miljöbilar." Riksrevisionen [Swedish National Audit Office] Granskningsrapport RIR 2020:1, Stockholm. ISBN 978-91-7086-549-7; Dnr 3.1.1-2019-0188.

Salvo, Alberto. 2018. "Flexible fuel vehicles, less flexible minded consumers: Price information experiments at the pump." *Journal of Environmental Economics and Management*, 92: 194–221.

Sheldon, T. L. 2022. "Evaluating Electric Vehicle Policy Effectiveness and Equity." *The Annual Review of Resource Economics*, 14: 669–688.

Tinbergen, Jan. 1952. *On the Theory of Economic Policy. Contributions to Economic Analysis*, Amsterdam:North-Holland Publishing Company.

TNO, IEEP, and LAT. 2006. "Review and analysis of the reduction potential and costs of technological and other measures to reduce CO2-emissions from passenger cars. Report for the European Commission." European Commission.

U.S. Department of Energy. 2008. "Flexible Fuel Vehicles: Providing a Renewable Fuel Choice." U.S. Department of Energy.

Appendix A: Proofs

Proof of Proposition 1

Let δ_j^0 be the mean utility at zero rebate – i.e., $\delta_j^0 = x_j\beta - \alpha p_j + \xi_j$. Then for a green vehicle the actual mean utility is $\delta_j = \delta_j^0 + \alpha R, j \in G_g$ because $\tilde{p}_j = p_j - R$ contributes $-\alpha(p_j - R) = -\alpha p_j + \alpha R$. The within-nest share of product j in the green nest is:

$$s_{j|g} = \frac{\exp(\delta_j/(1-\sigma))}{\sum_{k\in G_g}\exp(\delta_k/(1-\sigma))} = \frac{\exp((\delta_j^0+\alpha R)/(1-\sigma))}{\sum_{k\in G_g}\exp((\delta_k^0+\alpha R)/(1-\sigma))} = \frac{\exp(\delta_j^0/(1-\sigma))}{\sum_{k\in G_g}\exp(\delta_k^0/(1-\sigma))}$$

The common factor $\exp(\alpha R/(1-\sigma))$ cancels from numerator and denominator. Thus $s_{j|g}$ is independent of R. The within-nest composition is unchanged: both $S_{AFV} \equiv \sum_{j\in AFV} s_{j|g}$ and $\bar{E}_g \equiv \sum_{j\in G_g} e_j s_{j|g}$ are constants independent of R.

Now, let $\phi_g(R) = s_g(R)/(s_c(R) + s_g(R))$ be the green share of inside-good sales. Because S_{AFV} and $\bar{E}_g$ are constant, both policy targets are affine functions of the single scalar $\phi_g(R)$:

$$\begin{aligned} \text{RFS:} \quad & s_{afv}(R) = S_{AFV}\cdot\phi_g(R), \\ \text{CO2:} \quad & E_{avg}(R) = \bar{E}_c - (\bar{E}_c - \bar{E}_g)\,\phi_g(R). \end{aligned}$$

Both equations are strictly monotone transformations $\phi_g(R)$, hence perfectly collinear: the policymaker has one effective degree of freedom. The RFS target $s_{afv}(R) \geq \bar{s}_{afv}$ requires $\phi_g(R) \geq \bar{s}_{afv}/S_{AFV}$; the CO2 target $E_{avg}(R) \leq \bar{E}$ requires $\phi_g(R) \geq (\bar{E}_c - \bar{E})/(\bar{E}_c - \bar{E}_g)$. Each target pins ϕ_g to a specific threshold. Generically these thresholds differ – one is set by the renewable mandate $\bar{s}_{afv}$ and the technology mix S_{AFV}, the other by the emission targets $\bar{E}, \bar{E}_c, \bar{E}_g$ – so no single R satisfies both. No interior-solution assumption is required. ■

Proof of Proposition 2

The inclusive value of the green nest at rebate level R is $V_g(R) \equiv D_g(R)^{1-\sigma}$, where $D_g(R) = \sum_{k\in G_g}\exp(\delta_k/(1-\sigma))$ and, for $k \in G_g$, $\delta_k = \delta_k^0 + \alpha R$ with δ_k^0 the mean utility at zero rebate. Factoring the uniform rebate out of the inclusive value yields

$$D_g(R) = \exp\left(\frac{\alpha R}{1-\sigma}\right)\sum_{k\in G_g}\exp\left(\frac{\delta_k^0}{1-\sigma}\right) = \exp\left(\frac{\alpha R}{1-\sigma}\right)D_g(0).$$

Therefore the inclusive value becomes

$$V_g(R) = D_g(R)^{1-\sigma} = \left[\exp\left(\frac{\alpha R}{1-\sigma}\right)D_g(0)\right]^{1-\sigma} = e^{\alpha R}\,V_g(0),$$

where the $(1-\sigma)$ exponents cancel exactly. The green nest share then simplifies to a standard upper-level logit probability:

$$s_g(R) = \frac{e^{\alpha R}V_g(0)}{1 + V_c + e^{\alpha R}V_g(0)}.$$

Differentiating with respect to R,

$$\frac{\partial s_g}{\partial R} = \frac{\alpha\, e^{\alpha R} V_g(0)\,(1+V_c)}{\left(1+V_c+e^{\alpha R}V_g(0)\right)^2} = \alpha\, s_g(R)\big(1-s_g(R)\big).$$

The nesting parameter σ does not appear in the derivative. Because the uniform rebate shifts the inclusive value of the entire nest one-for-one, the upper-level logit governs the aggregate response and σ is irrelevant to it. The marginal effect is $\alpha\, s_g(1-s_g)$, which is strictly increasing in s_g for $s_g < \frac{1}{2}$ and approaches zero as $s_g \to 0$: when green vehicles are a niche, the rebate moves the aggregate green share only weakly, regardless of its size. ■

Proof of Proposition 3

The manufacturer chooses pathway $j \in \{afv, rfv\}$ to maximize profits $\Pi_j(R) = (p_j - c_j)D_j(R) - K_j$. The policy applies the identical rebate R to both pathways. As documented in Section 5, available AFV (flex-fuel) variants were priced similarly to their gasoline equivalents, and the limited number of RFV models below 120 gCO2/km were comparably positioned in their segments. We therefore maintain the empirically grounded assumption that gross operating margins are similar: $(p_{afv} - c_{afv})D_{afv}(R) \approx (p_{rfv} - c_{rfv})D_{rfv}(R)$.

The fixed costs are highly asymmetric. Converting an existing internal combustion engine to accept ethanol requires minimal R&D (K_{afv} is low), whereas engineering a vehicle to achieve the strict 120 gCO2/km threshold requires significant lightweighting and powertrain innovation ($K_{rfv} \gg K_{afv}$). Under similar margins, the fixed cost differential dominates the profit condition:

$$\Pi_{afv}(R) - \Pi_{rfv}(R) \approx K_{rfv} - K_{afv} > 0$$

establishing a strict preference for the AFV pathway. Furthermore, if K_{rfv} is sufficiently large that $(p_{rfv} - c_{rfv})D_{rfv}(R) < K_{rfv}$ for all feasible R within the program's budget, then $\Pi_{rfv}(R) < 0$ and the RFV pathway is never chosen. This yields the corner solution in which all compliance vehicle introductions are directed toward the high-emission AFV pathway. ■

Appendix B: Data

B.1 Sales data

Vroom has adjusted new car registration data to better represent the cars that are actually used by an individual and that do not serve as demonstration units or alike. For a registration to be included in the data set, the vehicle has to be acquired by an individual within 30 days of the registration. The sales data is aggregated at the base model level for each fuel type, i.e. the item Audi A3 gasoline contains all versions of the A3 that are primarily driven on gasoline. We consider seven different fuel segments: high- and low-emission gasoline; high- and low-emission diesel; gasoline/ethanol (FFVs); gasoline/gas (CNG); and electric hybrids.

B.2 Vehicle characteristics

The characteristics data is on a more disaggregated level than the sales data, e.g. there are 18 different Audi A3 gasoline versions. To be able to combine characteristics and sales data, we have aggregated the characteristics over sub-models based on the baseline model, thus following the literature, e.g. Berry, Levinsohn and Pakes (1995). Following the Swedish Consumer Agency definition, we define market segments according to vehicle weight, with the five segments defined by the thresholds 1100, 1250, 1400 and 1600 kg.

B.3 Combining sales data and characteristics

When combining the sales data and the characteristics data, a small fraction of the observations did not have a match. We thus expanded our search as follows, checking the following manually. First, we checked for the same brand, model and fuel type for the following year, since models for a given year are released late in the prior fall. Second, we checked for the same brand and fuel type for the same year. Third, for the same brand and fuel type for the following year. Finally, for the same fuel type and same year (the standard deviation is lower within a population consisting of cars of the same fuel type but different brands than within a population of a certain brand but with different fuel types).

B.4 Fuel economy and CO2 emissions

In the consumer guides, the emission data for FFVs is solely based on gasoline driving. According to The Swedish Consumer Agency (2008), there are no official emission values for ethanol driving. However, in their report on the climate effects of new cars, the Swedish Environmental Protection Agency (2008a) develops a way to calculate emission reductions. First, due to its lower energy content, E85 consumption is approximately 35 percent higher than gasoline consumption. Second, carbon dioxide emissions for E85 are 688.3 g/l, regardless of whether it made of sugar cane ethanol or sulphite pulp ethanol. Using this and the data on gasoline consumption from the guides, we can calculate ethanol consumption in l/100km (gasoline consumption*1.35) and carbon dioxide emissions in g/km (ethanol consumption*688.3/100). Thus, carbon dioxide emissions for E85 could be calculated by multiplying the gasoline consumption in l/100km by 9.29205 $(\frac{1.35 \times 688.3}{100})$. For winter months, the ethanol blend used is E75, which has a higher consumption (approximately 42 percent higher). For these months, the carbon dioxide emissions are multiplied by $(\frac{1.42 \times 688.3}{100})$.

To estimate CO2 emissions, we estimate yearly mileage driven (in km), together with a measure of CO2 emissions (in g/km). For captive cars, mileage estimates are based on the results reported in Appendix C whereas for FFVs they also depend on the fraction of vehicle owners using each fuel. We consider four cases, namely 0, 25, 50 and 75 percent gasoline usage. Although we do not take a stand on which level looks more appropriate, Huse (2026) documents a steep decrease in monthly ethanol sales following the 2008 oil price drop, which suggests that fuel switchers correspond to a non-negligible share of FFV owners.

The emission data for gas (CNG) is based on what is called certification gas, which is the same as fossil gas (Din Bil Stockholm/Hammarby, 2008). Carbon dioxide emissions from fossil gas are evaluated to be 2120 g/m3 whereas for biogas these are evaluated to 390 g/m3. The supply of vehicle gas in Sweden consists of both fossil gas and biogas, as well as a mixture of the two. According to Din Bil, the supply is evenly split, which is consistent with the report by the Swedish Environmental Protection Agency (2008a) which states that, in 2007, 53 percent of the vehicle gas sold was biogas and 47 percent was fossil gas. The emission data for gas cars is hence not correct since it assumes all cars are driven on fossil gas, thus the general emission levels for gas cars are exaggerated. We therefore re-estimate these to be equal to gas consumption per km*(2120*0.47 + 390*0.53), based on the numbers above.

B.5 Purchasing a Car

The registration of a vehicle with The Swedish Transport Agency (Transportstyrelsen) must take place within ten working days of a change in vehicle ownership.

B.6 Potential market

To go from observed quantities to observed market shares we need to define the size of the potential market for each time period. One way to obtain the potential market variables would be by estimating them, as suggested in Reiss and Wolak (2007). Alternatively, one could follow the criterion used in BLP (1995), where the total number of households constitutes the potential market. According to Reiss and Wolak (2005), this definition has some shortcomings. First, not all households can afford a new car and other entities than households can purchase cars. Since we only examine car sales to individuals, only the former poses a possible problem. It is not realistic that all households can afford to purchase a new car, therefore this would overestimate the potential market. Therefore we define the market as the number of individuals (instead of households, as Sweden has a high number of single person households) above the age of 20 with a yearly income of 300,000 SEK (about USD 42,000) or more. These are the potential purchasers of a new car. It is however unlikely that they can consider buying a new car each month. We therefore assume that consumers generally consider buying a new car every fifth year, thereby dividing the numbers by 60.

Appendix C: Local Pollutant Data Construction

Table C.1: Hedonic Regressions for Local Pollutants

	PM10	NOx	CO2
(Intercept)	−0.039*	−0.129***	63.579***
	(0.017)	(0.037)	(4.119)
CC	0.000 008	0.000 01	0.019***
	(0.000 006)	(0.000 009)	(0.001)
HP	−0.000 02	−0.0005***	−0.098***
	(0.000 07)	(0.0001)	(0.016)
Car Weight	0.000 02***	0.0003***	0.002
	(0.000 003)	(0.000 02)	(0.002)
Fuel Consumption	0.002	−0.016***	13.080***
	(0.001)	(0.003)	(0.337)
Aut. Transmission	0.008***	0.055***	1.708**
	(0.002)	(0.006)	(0.651)
Fueltype = Diesel	0.012**	0.136***	−0.441
	(0.004)	(0.008)	(0.867)
Fueltype = Diesel(Green)	0.018***	0.137***	−4.467**
	(0.005)	(0.013)	(1.393)
Fueltype = Gasoline(Green)		−0.001	3.944+
		(0.018)	(2.113)
Fueltype = Ethanol		0.0006	3.427*
		(0.012)	(1.335)
Fueltype = Gasoline		0.033**	−10.381***
		(0.012)	(1.394)
Num.Obs.	2016	7344	8453
R2	0.762	0.722	0.956

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
Note: All models include brand, model and year dummies.

Table C.2: Predicted and Observed values for Local Pollutants

variable	selector	min	q1	median	mean	q3	max	n
NOx	Imp.	0.00	0.09	0.12	0.15	0.19	0.94	2401
NOx	Obs.	0.00	0.02	0.04	0.13	0.20	0.70	7344
PM10	Imp.	0.00	0.0003	0.02	0.02	0.04	0.11	7729
PM10	Obs.	0.000007	0.004	0.02	0.03	0.04	0.10	2016
CO2	Imp.	128.69	173.13	193.57	211.13	240.34	367.00	1292
CO2	Obs.	109.00	161.50	190.86	197.47	224.00	397.00	8453

Appendix D: Details on Welfare Calculations

D.1 Welfare Accounting

We now derive the welfare components reported in Section 7.4: the change in VAT revenue, the change in producer surplus, and the accounting identity that links them to the fiscal cost. Throughout, quantities are measured as changes from the no-rebate ($R = 0$) baseline within each threshold pair (τ_r, τ_a).

Notation

For each product j, let p_j denote the VAT-inclusive list price (invariant across scenarios, since we hold prices fixed), $\tilde{p}_j$ the price the consumer effectively faces, $q_j(R)$ the quantity sold under rebate R, and E the set of rebate-eligible vehicles. The consumer who buys an eligible vehicle pays the full list price p_j at purchase and later receives the rebate r directly from the government, so the effective consumer price is $\tilde{p}_j = p_j - R$ for $j \in E$ and $\tilde{p}_j = p_j$ otherwise. The Swedish VAT rate is $\tau = 0.25$, so the VAT content of a VAT-inclusive price is $\frac{\tau}{1+\tau} p_j = 0.2\, p_j$. Define the simulation aggregates

$$\text{trevenue}(R) = \sum_j p_j\, q_j(R), \qquad \Pi(R) = \sum_j (\tilde{p}_j - mc_j)\, q_j(R), \qquad \text{FC}(R) = R \sum_{j \in E} q_j(R),$$

where $\Pi(R)$ is the variable profit reported by the demand model and mc_j the (unobserved) marginal cost.

VAT revenue

VAT is levied on the full list price the consumer pays at the point of sale; the rebate is a separate transfer, disbursed later and directly to the buyer, and does not enter the tax base. Total VAT revenue is therefore

$$\text{VAT}(R) = 0.2 \cdot \sum_j p_j\, q_j(R) = 0.2 \cdot \text{trevenue}(R).$$

Because list prices are invariant across scenarios, the change in VAT revenue reflects only quantity responses:

$$\Delta\text{VAT} = 0.2 \cdot \sum_j p_j\, \Delta q_j, \qquad \Delta q_j = q_j(R) - q_j(0).$$

Inframarginal units ($\Delta q_j = 0$) contribute nothing; only newly induced sales raise VAT.

Producer surplus without marginal costs

Producer surplus is variable profit valued at pre-VAT prices,

$$\text{PS}(R) = \sum_j \left(\frac{p_j}{1.25} - mc_j \right) q_j(R) = \frac{1}{1.25}\, \text{trevenue}(R) - \sum_j mc_j\, q_j(R),$$

which appears to require the unobserved mc_j. It does not. The demand model reports profit at consumer-facing prices, $\Pi(R) = \sum_j (\tilde{p}_j - mc_j) q_j(R)$, and since $\sum_j \tilde{p}_j q_j = \text{trevenue}(R) - \text{FC}(R)$ (the rebate lowers the eligible consumer price by exactly R), we can solve for the cost aggregate,

$$\sum_j mc_j\, q_j(R) = \text{trevenue}(R) - \text{FC}(R) - \Pi(R).$$

Substituting and simplifying,

$$\text{PS}(R) = -0.2 \cdot \text{trevenue}(R) + \text{FC}(R) + \Pi(R),$$

and, in changes from the $R = 0$ baseline,

$$\Delta\text{PS} = -\Delta\text{VAT} + \text{FC} + \Delta\Pi$$

which uses only ΔVAT, FC, and $\Delta\Pi$ – all direct simulation output. No marginal cost need be recovered.

The accounting identity

Rearranging,

$$\Delta\text{PS} + \Delta\text{VAT} = \text{FC} + \Delta\Pi.$$

Across all ten policies the change in variable profit is negligible, $\Delta\Pi \approx 0$: with prices fixed, a rebate reallocates quantities but leaves the markup structure intact. The identity thus collapses to

$$\Delta\text{PS} + \Delta\text{VAT} \approx \text{FC},$$

i.e. the rebate outlay returns in full to the private sector, as producer surplus on additional sales ($\approx$ 90–94%) and as incremental VAT ($\approx$ 6–10%). Consumers separately gain ΔCS from the lower effective price.

Three welfare measures

The measures differ only in how the rebate outlay is treated:

$$\begin{aligned} W_1 &= \Delta\text{CS} - \text{FC} - \Delta\text{Env} && \text{(government-budget: outlay is a cost)},\\ W_3 &= \Delta\text{CS} + \Delta\text{PS} - \text{FC} - \Delta\text{Env} && \text{(}W_1\text{ augmented by producer surplus)},\\ W_2 &= \Delta\text{CS} - \Delta\text{Env} && \text{(total surplus: transfers cancel)}. \end{aligned}$$

Using the identity, $W_3 = W_1 + \Delta\text{PS}$ and $W_3 = W_2 - \Delta\text{VAT}$, so

$$W_1 \leq W_3 \leq W_2, \qquad W_3 - W_1 = \Delta\text{PS} \geq 0, \qquad W_2 - W_3 = \Delta\text{VAT} \geq 0.$$

W_2 is the theoretically appropriate social criterion; W_1 is the government-budget perspective used to assess whether the program is a sound use of public funds.

Why $\Delta\text{VAT} \ll \text{FC}$: inframarginal sales

The fiscal cost is paid on all eligible units, $\text{FC} = R(Q_E^0 + \Delta Q_E)$, where Q_E^0 would have sold without the rebate and ΔQ_E are newly induced. VAT rises only with new activity, so $\Delta\text{VAT}/\text{FC} \approx (0.2\bar{p}/R)(\Delta Q/Q_E)$. For the actual GCR, $\Delta\text{VAT}/\text{FC} = 58.7/571.6 \approx 10\%$; with an average list price $\bar{p} \approx 200$ KSEK, this implies $\Delta Q/Q_E \approx 2.6\%$. Only a few percent of rebated sales are genuinely induced; the remainder are inframarginal transfers. This is the fiscal counterpart of the niche-market inertia of Proposition 2.

D.2 Detailed Welfare Results

Here we provide a version of the Pareto frontier distinguishing between CO2 and local pollutants and a welfare table with the three welfare measures.

Figure D.1: **Program Cost versus Change in CO2 and Local Pollutant Emissions**

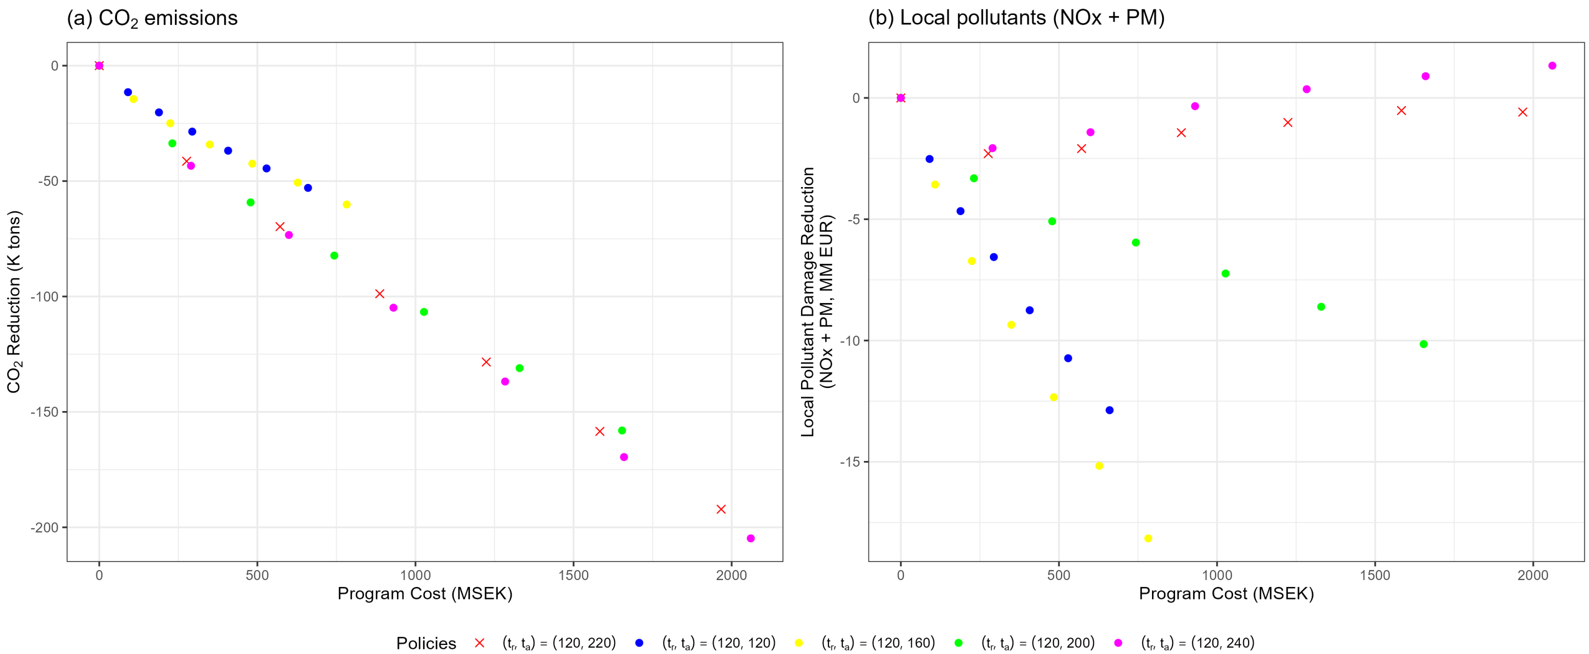

Panel A

Panel B

Table D.1: Welfare Decomposition: Full Breakdown

(τ_r, τ_a)	R	ΔCS	FC	ΔVAT	ΔPS	ΔEnv	W_1	W_2	W_3
Asymmetric policies $(\tau_r \leq \tau_a)$									
(120,220)	**10**	**534.1**	**571.6**	**58.7**	**512.9**	**19.2**	**–56.7**	**514.9**	**456.2**
(120,220)	30	1716.5	1966.8	156.8	1810.0	47.7	–298.0	1668.8	1512.0
(120,120)	10	173.7	188.7	15.3	173.4	9.6	–24.7	164.1	148.8
(120,120)	30	562.8	660.2	37.8	622.4	25.8	–123.3	537.0	499.2
(120,160)	10	207.7	224.6	20.3	204.3	12.8	–29.7	194.9	174.6
(120,160)	30	671.7	782.6	46.5	736.0	32.9	–143.7	638.8	592.3
Symmetric policies $(\tau_r = \tau_a)$									
(160,160)	10	1132.9	1226.9	86.1	1140.8	65.1	–159.0	1067.8	981.7
(160,160)	30	3581.8	4084.5	250.0	3834.4	199.3	–702.0	3382.5	3132.5
(220,220)	10	2472.9	2660.3	200.7	2459.5	96.5	–283.9	2376.4	2175.7
(220,220)	30	7611.8	8403.0	648.8	7754.0	311.7	–1102.9	7300.1	6651.3

Notes: All figures in MSEK, changes from the no-rebate baseline within each (τ_r, τ_a) pair. Bold row is the actual GCR. ΔVAT + ΔPS = FC in every row up to rounding, verifying the accounting identity. ΔEnv is the sum of monetized CO2, NOx and PM damages (positive = greater damage). W_1 = ΔCS − FC − ΔEnv; W_2 = ΔCS − ΔEnv; W_3 = ΔCS + ΔPS − FC − ΔEnv = W_1 + ΔPS.

Appendix E: Vehicle Sales and Program Reach

E.1 Total Vehicle Sales and Market Share of Eligible vehicles

We examine the impact of alternative policies on new vehicle sales by calculating the quantity sold of each vehicle under each simulated policy for the period April 2007-June 2009, see Figure E.1. This analysis helps us to understand the channels through which our results materialize and to which extent environmental policies aimed at the new vehicle market can stimulate the economy.

Total Sales and Size of the Program. Panel A examines the effect of alternative policies on total vehicle sales. While policies within a graph show an intuitive pattern according to which more generous policies result in increased sales, comparing symmetric and asymmetric policies, i.e. policies across graphs within a panel, shows the limited extent to which these policies (especially asymmetric ones) can affect total vehicle sales. Of particular interest are the "jumps" in sales when going from $(120, 120, R)$ to $(160, 160, R)$ and then from the latter to $(200, 200, R)$, which is consistent with the idea that the larger, best-selling models have relatively high emissions.

To assess the breadth of the alternative policies, Panel B displays the market share of vehicles eligible for the rebate. On the left-hand side, the fact that there are few low-emission RFVs substantially constrains the market share of eligible vehicles yet results in a low-cost program.

In the case of symmetric policies, the levels of the responses are again intuitive, with the most generous $(240, 240, R)$ set of policies resulting in eligible vehicles commanding a market share in excess of 90 percent in the new vehicle market.

Finally, Panel D reinforces the above results, with intuitive responses in sales to changes in emission thresholds and rebate amounts.

E.2 Sales of Gasoline and Diesel Vehicles

We now document the sales of diesel (Panel A) and gasoline vehicles (Panel B) for the policies we simulate.

The overall effects for diesel vehicles are in the range of 10 percent: for asymmetric policies, the responses go from approximately 78,000 units to 72,000 units. For symmetric policies, the changes range from 78,000 units to 87,000 units. In what concerns gasoline vehicles, the changes are smaller in relative terms: for asymmetric policies, the positive responses go from approximately 155,000 units to 160,000 units whereas the negative responses range from approximately 155,000 units to 143,000 units. For symmetric policies, the changes range rom 155,000 to 162,000 units.

Zooming into the effects for the diesel segment, we start with the left-hand side chart of Panel A, which reports the effects of asymmetric policies. The first thing to note is that the sales responses differ: the stricter family of policies (120,160) shifts the sales slightly towards green vehicles whereas the looser families of policies result in a decrease of the sales of gasoline vehicles. In the case of symmetric policies, the effect is monotonic in rebate amounts but not in CO2 emission thresholds, as the strongest response occurs in the case of families (200,200), which is closely followed by (240,240) then (160,160).

Focusing now on the gasoline segment, we see that once again the sales responses differ among asymmetric policies: the stricter family of policies (120,160) shifts the sales slightly towards green vehicles whereas the looser families of policies result in a decrease of the sales of gasoline vehicles. In the case of symmetric policies, the effect is monotonic in rebate amounts but not in CO2

Figure E.1: **Total Vehicle Sales and Market Share of Eligible Vehicles**

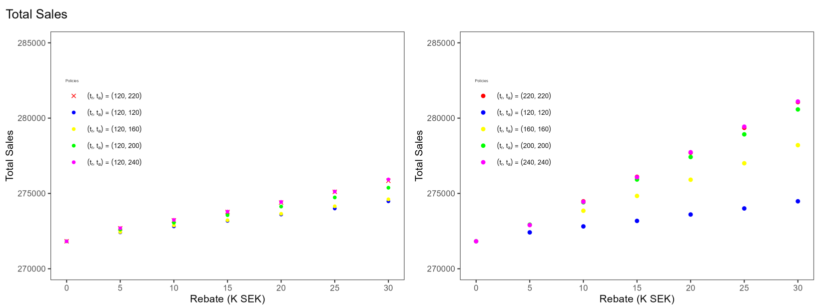

Panel A

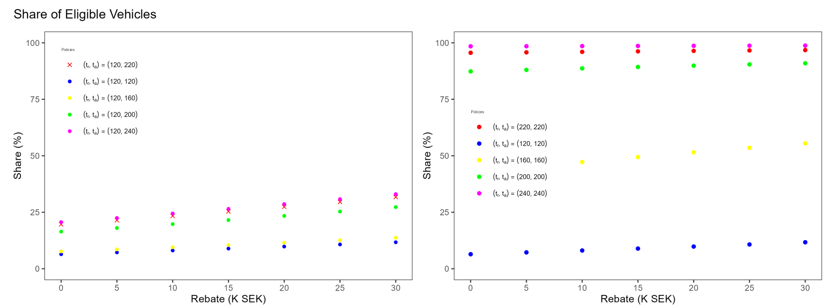

Panel B

emission thresholds, as the strongest response occurs in the case of families (200,200), which is closely followed by (240,240) then (160,160).

Figure E.2: **Vehicle Sales in the Diesel and Gasoline Segments**

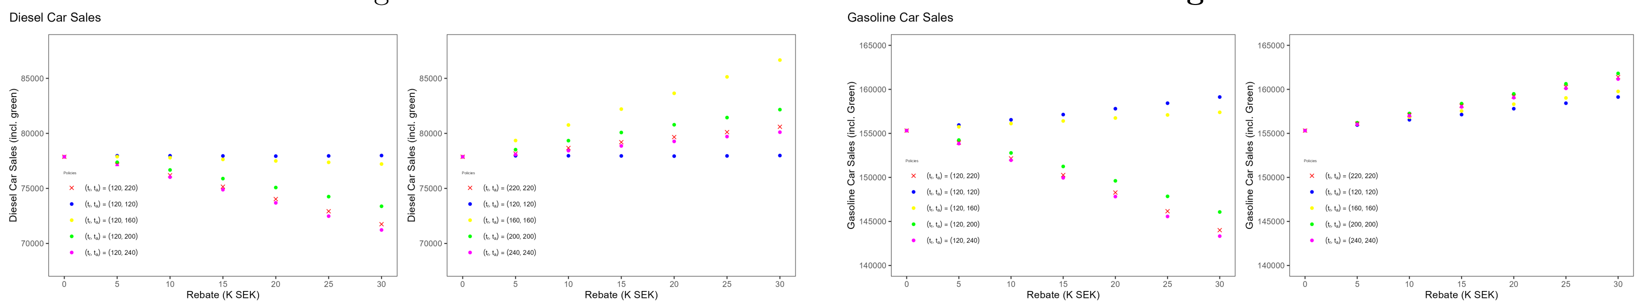

Panel A

Panel B