FINAL REPORT

Modeling Household Vehicle and Transportation Choice and Usage Part B: Empirical Estimation of Household Vehicle Purchase and Usage Decisions

Contract Number 11-322

Prepared by:
David Rapson, Kenneth Gillingham (Yale), and Chris Knittel (MIT)

University of California, Davis

April 1, 2017

Prepared for the California Air Resources Board and
the California Environmental Protection Agency

Disclaimer

The statements and conclusions in this Report are those of the contractor and not necessarily those of the California Air Resources Board (ARB). The mention of commercial products, their source, or their use in connection with material reported herein is not to be construed as actual or implied endorsement of such products.

Acknowledgment

The authors thank Melanie Zauscher, Emily Wimberger, and Belinda Chen of the Air Resources Board for their comments, feedback, and encouragement in overseeing this project. The authors also gratefully acknowledge the research assistance of Stephen Sun and James Archsmith of UC Davis, and Hao Deng of Yale University. We thank the California Department of Motor Vehicles (DMV) for making the registration data available for research.

This Report, along with "Factors Related to Voluntary Choice of Low Vehicle Ownership and Usage" was submitted in fulfillment of Contract 11-322 by UC Davis, Department of Economics under the sponsorship of the California Air Resources Board. Work was completed on April 14, 2017.

Contents

List of Figures

List of Tables

Abstract

This project used detailed vehicle-level data to explore the entire process of household vehicle choice, holdings, and usage for the light duty fleet in California–the largest contributor to transportation emissions. By using cutting-edge statistical approaches grounded in economic modeling, we have advanced the understanding of how consumers make decisions that influence the evolution of the vehicle fleet in California. We find that there are important differences across households in the sensitivity of travel and vehicle choice to the price of gasoline. We are the first to estimate a "portfolio effect" in household vehicle purchases, where attributes of one car affect the revealed desired choice of another. Results show that this effect influences the turn-over of the fleet, and potentially erodes energy savings from fuel economy standards. Finally, we develop a new model of vehicle choice and driving that incorporates key features of the decision-making process of forward-looking consumers. These innovations provide important insights into the effects of policies to further reduce transportation emissions in California.

Executive Summary

Background. California has long been a leader in transportation policies to reduce smog-forming and particulate matter emissions, and more recently has positioned itself as a leader in policies to reduce greenhouse gas emissions from transportation. These policies have made dramatic progress in reducing ambient air pollution and improving human health, and have the potential to help address the imminent threat of global climate change. Yet much remains to be accomplished. California has set a goal of an 80 percent reduction in greenhouse gases by 2050, a bold target, with a significantly longer outlook than the federal government uses in federal policy development. To reach such an ambitious target requires a suite of regulations and other strategies such as incentives to reduce multiple pollutants and induce innovation in vehicle technology, while at the same time being targeted to be as cost-effective as possible. Each of the research projects described in this report tests hypotheses that will enhance our understanding of the dynamics of the vehicle market and directly inform the next iteration of transportation emissions abatement policies in California.

Objectives and Methods. The empirical aspect of the project advances our understanding of vehicle choice, holdings, and usage for the light duty vehicle fleet in California. The objective is to test three hypotheses. First, we hypothesize that households that own multiple vehicles have a preference for diversity in vehicle attributes that are correlated with fuel economy. Second, we hypothesize that an increase in gasoline prices (whether from market forces or through public policy) will affect vehicle scrappage decisions, new vehicle purchase decisions, and miles traveled. Finally, we hypothesize that the persistence of observed vehicle ownership is due to the transaction frictions associated with buying and selling vehicles.[1] If true, the vehicle choice becomes a dynamic optimization problem that includes not only what car to buy, but when to buy it.

In order to test the household-level hypotheses, we first needed to generate a dataset that tracks household vehicle ownership over time. We developed a detailed algorithm for converting the VIN-level DMV registration record dataset with names and addresses (housed at ARB) into a household- and premise-level dataset that was free from identifying household information. Cleaning and generating the data for this study was a major contribution of this project, and the dataset can now be used by other researchers at ARB.

[1]"Transaction frictions" are any unobserved costs incurred by buyers and sellers when trying to consumate a transaction. These could include search costs, hassle costs, disutility of negotiating with used car dealerships, etc.".

To test the first and second hypotheses we use a variety of regression-based approaches. To estimate household portfolio preferences[2], we narrow our focus to two-car households that replace one of their cars with another during the period 2000-2007 in California. This allows us to isolate a thought exercise that then informs our empirical strategy. We deploy a novel approach to a) exploit random variation that leads to one of the two cars in the portfolio being dropped/sold, and to b) exploit random variation in the fuel economy of the car that is kept. The richness of our data also allow us to identify these portfolio preferences using within-household variation.

To further test the second hypothesis we examine the relationship between gasoline prices and local criteria pollutant reductions in California from 1996-2010. We hypothesize that emissions from vehicles decrease differently in response to high gasoline prices across vehicles and locations. If true, an optimal gasoline tax that accounts for differences between households would differ from one that does not.

For the final contribution, we construct a cutting-edge dynamic discrete choice model of vehicle purchases that allows us to retrieve consumer preference and market parameters. The model is solved by finding parameters that allow our model predictions to match what we observe in the data. The model then serves multiple purposes. The estimated transaction friction parameter provides information about the search, matching, behavioral and other costs of buying and selling cars in the marketplace. It is possible then to perturb one or more variables to examine how key statistics evolve over time under different counterfactual scenarios. For example, we can increase fuel economy or the gasoline price and examine how the composition of the vehicle fleet and derived demand for gasoline change.

Results. Results show that emissions from high-emission vehicles decrease more when gasoline prices are high, highlighting important differences across households. Under these conditions, an optimal gasoline tax that accounts for differences across households will be higher than one that does not, and will also be far more effective (though still falling short of an optimal tax imposed directly on the pollutants themselves). The household portfolio work documents the importance of portfolio interactions when considering overall household gasoline consumption. When a kept car is more fuel efficient, demand in the rest of the portfolio for less fuel efficient cars increases.

Taken together, these results demonstrate how variants in policies intended to reduce tailpipe

[2]We define the household vehicle "portfolio effect" as the impact that changing fuel economy of one car has on the desired fuel economy of another car in that household's portfolio.

emissions can lead to very different outcomes. In the context of multi-car households, the portfolio preferences that we uncover imply that fuel economy standards will generate countervailing forces in subsequent years. In the context of gasoline taxes, it is important to account for household differences in the sensitivity of travel and vehicle choice to the price of gasoline. Bringing together rich micro-datasets and cutting-edge methodologies has allowed us to reach fresh insights on these important topics.

Conclusions. We have developed a new dataset that tracks the household vehicle portfolio over several years. We have developed methods to estimate the size of the household portfolio effect. We find that households exhibit a preference for fuel-economy diversity in their vehicle portfolio, which has implications for the effectiveness of fuel economy standards. We also describe how different responses to changes in gasoline prices can inform optimal gasoline taxes. Finally, we have designed and built a dynamic discrete choice model of vehicle purchase and sale decisions at the household level. There are several next steps. Perhaps the most exciting and important will be to enrich the dynamic discrete choice model and deploy it to forecast the California vehicle fleet evolution under alternate policy scenarios.

1 Introduction

California has long been a leader in transportation policies to reduce smog-forming and particulate matter emissions, and more recently has positioned itself as a leader in policies to reduce greenhouse gas emissions from transportation. These policies have made dramatic progress in reducing ambient air pollution and improving human health, and have the potential to help address the imminent threat of global climate change. Yet much remains to be accomplished. California has set a goal of an 80 percent reduction in greenhouse gases by 2050, a bold target, with a significantly longer outlook than the federal government uses in federal policy development. To reach such an ambitious target requires a suite of regulations and other strategies to reduce multiple pollutants and induce innovation in vehicle technology, while at the same time being targeted to be as cost-effective as possible.

The project objectives are to develop a dataset and analysis to test a series of hypotheses and develop new methodologies for understanding the evolution of the vehicle fleet in light of different policy approaches to reduce emissions and improve human health. Our three hypotheses are the following. First, we hypothesize that an increase in gasoline prices (as one would experience when exposed to an increase in the gasoline tax) will affect vehicle scrappage decisions, new vehicle purchase decisions, and miles traveled. Testing this hypothesis is important for understanding how changing gasoline prices influences the vehicle fleet and environmental outcomes. Second, we hypothesize that households that own multiple vehicles have a preference for diversity in vehicle attributes that are correlated with fuel economy. For example, if a household owns an highly fuel-efficient vehicle, will they be more likely to buy another fuel-efficient vehicle in their next purchase or would they diversify and buy a less-efficient vehicle? If the latter, then policies to improve fuel efficiency of new vehicles in the fleet may have the unintended consequence of leading to increased demand for less-efficient vehicles in the future, with implications for gasoline demand and emissions. Finally, we hypothesize that the persistence of observed vehicle ownership is due to the transaction frictions associated with buying and selling vehicles. Understanding how the turnover of vehicles in the fleet occurs–and is influenced by policy–is important for forecasting the future evolution of the vehicle fleet and associated emissions.

The work undertaken in this project relates to an extensive literature in economics and transportation (Mannering and Winston, 1985; Esteban and Shum, 2007; Stolyarov, 2002; Adda and Cooper, 2000; Schiraldi, 2011; Bento et al., 2009; Wakamori, 2016; Cernicchiaro and de Lapparent, 2015; Busse et al., 2013). Because each of the three hypotheses involves a distinctly different, although somewhat overlapping, sets of references, we cover the related literature in each of the sections of this report and the appendices (the two appendices have full introductions each). It should be noted that we follow standard assumptions in the literature throughout, and in fact, go beyond the literature in some parts of the report. For example, our work on the second hypothesis goes beyond the literature by focusing on the entire household vehicle portfolio, rather than on individual vehicles treated as separate decision-makers. In our work on the third hypothesis–the dynamic discrete choice model–we focus on the standard assumption of individual vehicles being controlled by separate decision-makers in order to focus on the effect of transaction costs, pushing the methodological frontier in a different direction.

This project provides key policy-relevant deliverables, including an algorithm for developing a household-level dataset, an analysis demonstrating the importance of the household portfolio for vehicle choice decisions, an analysis examining how the complex relationship between driving and gasoline prices is mediated and its implications for the air quality, and a complete dynamic discrete choice model of vehicle choice and usage of the California fleet. The project uses cutting-edge statistical approaches to shed new light on important questions about the evolution and usage of the vehicle fleet and is providing a new dataset and tool for use in future analyses of the effects of changes in fuel prices or policies.

There are seven discrete, yet inter-related, tasks that make up Part B this project.[3] The tasks build from the beginning, going from dataset development to analysis to writing.

The first task–TASK B.1–involves merging, cleaning, and preparing a full dataset covering the years 1998 to 2011 of vehicle purchases and holdings from the California Department of Motor Vehicles, as well as driving behavior from the vehicle-level odometer readings taken during smog checks managed by the California Bureau of Automotive Repair. One key aspect of this combined dataset is that we appended a detailed vehicle information number (VIN) decoder to determine the characteristics and fuel economy of all vehicles in the dataset (matched by VIN). The VIN decoder was a commercial decoder from DataOne, Inc that was improved upon by dedicated research assistant work. The combined dataset provides a complete picture of the California light duty fleet and the amount it is driven over most of a decade.

The second task–TASK B.2–converts the dataset from a vehicle-level dataset to a household-level dataset. The household-level dataset includes anonymous identifiers for households,

[3]Part A is "Factors Related to Voluntary Choice of Low Vehicle Ownership and Usage".

allowing for rigourous and informed analysis of household-level decisions at the scale of the entire state. To develop the household-level identifier, we worked closely with staff at ARB to develop an algorithm for designing households that accounts for some of the challenges in the data, such as multi-family dwellings. The complete anonymized dataset includes all residential light-duty vehicles and their driving, as well as a household identifier. This unique dataset forms the foundation for the modeling in all of the following tasks.

The third task–TASK B.3–explores summary statistics to analyze the composition and driving behavior of the vehicle fleet in California over the years in our dataset. This task provides the underpinning knowledge necessary before any more complex statistical analysis. Without initial exploration and understanding of the data, correct interpretation of results is impossible. For this task, we examined a variety of statistics to better understand the vehicle fleet. In particular, we realized through our explorations the importance of the household portfolio for vehicle choice decisions. This led to a rigorous and careful analysis, using a novel approach, of the effect of the other vehicles in a household portfolio for the decision of what the next car to buy will be.

The fourth task–TASK B.4–rigorously examines how vehicles in California respond to changes in the gasoline price (or the price of driving) and how this differs with vehicle attributes. This task involved econometric analysis of driving decisions which unearthed remarkable differences across vehicle types. For example, the analysis revealed that much heavier and less-efficient vehicles were substantially more responsive to gasoline price changes than smaller and more efficient vehicles. The work on this task extensively explored the full policy implications of this heterogeneity.

The fifth task–TASK B.5–involves designing, coding, and running a joint vehicle choice and usage model (using Matlab). The model developed for this task uses a computationally intensive dynamic discrete choice modeling approach, which is considered cutting-edge in the economic literature. There are numerous advantages to this more intensive modeling approach, including a more careful modeling of forward-looking consumers who can make decisions with the forecast of their future vehicle options in mind. The approach also notably includes transactions costs, or transaction frictions, which mediate the evolution of the vehicle fleet. Transaction costs or frictions are the very real search and information costs that deter consumers from buying or selling their car every year to get the car that is optimal for them at any given time period. Modeling these transaction frictions allows for a more careful analysis of the evolution of the vehicle fleet over time as consumers buy, sell, and replace vehicles. Such a model is at the research frontier, with very few successful implementations in the literature (Schiraldi, 2011, see, e.g.,), and none examining the California fleet.

The sixth task–TASK B.6–performs robustness checks, simulations, and forecasts using the dynamic discrete choice model. Using the model, we have been exploring the parameter space and testing a variety of assumptions to better understand the model fit. We have also performed some cursory simulations and forecasts. Going forward, we look forward to further conversations with ARB for guidance on which counterfactual policy simulations will be most relevant for decision-making.

The final task–TASK B.7–is to write this final report.

Some of our key conclusions are as follows. We find evidence that households value diversification in the vehicle portfolio. The greater is the fuel economy of the kept car, the lower the fuel economy of the purchased car. Increases in the fuel economy of the kept car reduces the probability the household purchases a car in the lower quartile of gallons per mile, while such increases reduce the probability the household buys a car in the upper quartile. Changes in gasoline prices affect the preferences for diversification in intuitive ways. As gasoline prices increase, the effect of the fuel consumption of kept vehicle and the probability of buying a car in the lower quartile of fuel consumption becomes even more positive. In contrast, as gasoline prices increase, the effect of the fuel consumption of kept vehicle and the probability of buying a car in the upper quartile of fuel consumption becomes even more negative.

To gauge the importance of the portfolio effect, we use our results to estimate the net effect of an increase in the fuel economy of the kept vehicle. We calculate the decrease in the fuel economy of the newly purchased vehicle when we increase the fuel economy of the kept vehicle by 10, 25, and 50 percent, across gasoline prices of $2.00, $3.00, and $4.00. These calculations suggest that the portfolio effect can have large consequences on the net effect a one-time increase in fuel economy. The fuel savings from increasing the fuel economy of the kept vehicle are eroded from the resulting decrease in fuel economy of the newly purchased vehicle.

Our analysis of how VMT changes in response to gasoline prices leads to several policy implications. A uniform gasoline tax that imposes the same tax on all vehicles does a poor job of addressing the market failure from pollution externalities. The dirtiest vehicles are not taxed enough, and many clean vehicles are over-taxed. This is true even when the uniform tax is calculated taking the correlation between emissions and VMT sensitivities into account. The roughly 50 percent increase in the tax level from a uniform gasoline tax correctly abates more emissions from the dirtiest vehicles, but also over-taxes the cleanest vehicles by a larger amount. The welfare benefits of the uniform gasoline tax are around 10% higher than those from a naïve tax, but still fall far short of the benefits from an optimal tax linked to actual vehicle emissions. Furthermore, there are enough differences in the per-gallon damage that even a tax targeting broad groups leaves a substantial portion of societal loss.

The research findings and modeling tool developed will provide insight into how consumers have responded and may respond to current and future California Air Resources Board vehicular policies, such as new emission standards and incentives. Additionally, this tool can be used to improve statewide vehicular and emission inventories used to support of policy development. The remainder of this report goes into detail in describing our methodology and findings from pursuing each of these tasks. Throughout, we refer back to these tasks to clarify how our work helped meet them. From this project, we have two completed working papers, which are attached as appendices. The first explores the influence of the household portfolio on the evolution of the vehicle fleet. The second demonstrates the differences across vehicle types in the driving responsiveness to gasoline prices and its implications for policy and air quality. In addition, we are providing ARB with our code for the working dynamic discrete choice model, which is explained in detail in Section 4.

2 Household Portfolio Preferences

2.1 Data Sources

TASK B.1 and TASK B.2 were to "merge, clean and prepare the full dataset" and to "convert the dataset from a vehicle-level dataset to a household-level dataset". These tasks took much longer than the initial project timeline projected, primarily due to a change in rules about accessing data from the California Department of Motor Vehicles. The rule change delayed receipt of the household-level data by roughly two years, and made necessary additional tasks by the investigator team that were beyond the scope of the project proposal. Nonetheless, we currently have a dataset of household-level vehicle purchase and usage data for California. This is a novel and extremely valuable platform for investigating issues relating to the evolution of the vehicle fleet and transportation emissions.

A majority of vehicle-owning households in California own multiple cars. In this part of the project we study them. In particular, we ask the question "do the attributes of one car in a household vehicle portfolio affect the desired attributes for the car to be purchased next?" The overwhelming majority of academic research on vehicle demand considers purchase decisions to be independent across cars, thus implicitly assuming "no" as the answer to the question. That is, we ignore any potential correlations between vehicle choices in households that own more than one car. The primary reason for such a narrow focus has been a lack of availability of data on household-level vehicle ownership. In this project, we overcome those data obstacles, consider a thought exercise that allows us to ask a well-defined research question, and develop an empirical strategy that addresses key potential confounding factors.

The cornerstone of our dataset is the universe of California vehicle registration records that occurred from 2001-2007. The DMV dataset includes every vehicle registered under the residential designation code. In California every vehicle must be registered annually. Each record includes the registrant's US Census block group identifier, the 17-digit vehicle identification number (VIN) that uniquely identifies the vehicle, that year's registration date, the date when the vehicle was last sold, transaction price, and various other information. This information allows us to construct a household-level panel dataset vehicle ownership.

Basic vehicle attributes (e.g. horsepower, weight, etc) are available via a VIN decoder that we purchased from DataOne Software. We augment the decoder to include vehicle fuel economy, which is available from the US Environmental Protection Agency. Odometer readings are available for each VIN from the Bureau of Automotive Repair (BAR) whenever the vehicle is sold and upon receiving biannual Smog Check certification. We use odometer readings, along with Smog Check dates, to calculate a rate of vehicle usage.[4] We thus have an average measure of miles traveled by each vehicle and, by extension, each household for each year in our sample. The coarseness of these data are not optimal for examining high-frequency effects of VMT-switching between vehicles in response to changes in gasoline prices. Nonetheless, gasoline prices are a variable of interest in this study, since they affect the household's optimal portfolio of vehicle fuel economy. Our gasoline price data are from the Oil Price Information Service (OPIS) and the U.S. Department of Energy, Energy Information Administration. These data are at the county-month level. The full dataset contains many millions of observations, and we use randomly drawn subsamples of the dataset for each of the data explorations and analyses for computational feasibility (the size of the dataset for each will be clear in the results tables).

[4]For example, if the odometer increases by X miles between Smog Check instance one and Smog Check instance two, and the number of days between these Smog Checks is Y, the daily average VMT during that period is X/Y.

We merge these data sources together to have a full dataset at the vehicle level. This data merging was done at the lowest level of disaggregation possible.

2.2 Methodology

TASK B.3 was to "Explore the summary statistics to analyze the composition and driving behavior of the vehicle fleet in California." These explorations led to our realization of the importance of household vehicle portfolio preferences. We thus went even beyond the scope of the task in providing a full analysis and working paper on household portfolio preferences. For this analysis, consider two-car households that replace one of their cars with another. Our research question asks: do households that own multiple vehicles have a preference for diversity in vehicle attributes that are correlated with fuel economy?

The ideal experiment for our research question would randomly assign which vehicle is "kept", perturb its fuel economy randomly, and then observe the relationship between the fuel economy of this kept vehicle and the fuel economy of the newly-acquired vehicle. Since this ideal experiment is obviously not possible, our identification strategy must overcome two sources of endogeneity stemming from the non-random assignment of the kept vehicle. The first is the choice of which vehicle to replace. Since the household preference for particular features of a multi-car portfolio will directly inform the decision of which car to keep or drop, there is an identification challenge in estimating the portfolio itself using observational data. The second is related to the presence of unobserved household preferences. Household fixed effects can address time invariant unobserved preferences, but there would still be a concern if these preferences change over time. We expect that these time-varying preferences would generate a positive correlation between the fuel economy of the kept and newly-acquired vehicles.

We use two sets of instrumental variables[5] to account for these potential sources of bias. The first set of instruments are derived from the observation that changes in the relative price of cars in a portfolio systematically affect the probability that the lowest fuel economy car is dropped. We discuss and present three instruments that rely on this feature of the choice setting, with our preferred instrument being deviations in expectation of the change in relative vehicle prices at the time when the kept car was initially purchased. The second instrument is the gasoline price at the time of the purchase of the kept vehicle. A number

[5] A valid "instrumental variable" is a special variable that provides us with independent variation, thereby allowing us to separate a causal effect from a correlation.

of papers (Klier and Linn (2010), Busse et al. (2013), Gillingham (2011)) have shown that vehicle purchase behavior is influenced by contemporaneous gasoline prices. Given the results of this literature, we would expect that the fuel economy of the kept vehicle is influenced by the gasoline price at the time of that purchase—something we confirm in our own data. We argue that the instrument provides independent variation because, after controlling for current gasoline prices, we would not expect past gasoline prices to influence the choice of the new vehicle.

The model of household choice that underpins our empirical approach allows for any number of household decision-makers. Different households will operate differently with respect to vehicle choice. In some households, each household member may exercise autonomy over the choice of one vehicle but not another. In other households, a single decision-maker may make all vehicle choices. Our empirical approach is robust to either of these models. It would be an interesting research question to be able to identify which of these models dominates, and what implications that may have for transportation policy. However, we do not have sufficient information to examine this feature of the setting.

2.2.1 Regression Specifications

The basic regression strategies examine the relationship that fuel economy of the kept car has on the chosen fuel economy of the bought car. The dependent variable is thus either fuel economy of the bought car itself (f_{it}^b), or quartile indicators of that variable. Regressors of interest include gasoline price at the time of purchase, fuel economy of kept car (f_{it}^k), and their interaction. In addition, we include a term, $\mathbb{1}\{\Delta f^{kd} > 0\}$, that distinguishes which car was dropped from the initial portfolio, the low- or high-GPM car. Specifically, $\Delta f^{kd} = f_{it}^k - f_{it}^d$ in our main specification, such that $\mathbb{1}\{\Delta f^{kd} > 0\} = 1$ indicates that the car with the lower fuel economy in the initial portfolio is kept.

Most of the regression results that follow are retrieved from estimating a linear model of the probability of purchasing vehicles in a given MPG quartile. For ease of exposition of the results, and to allow our focus to rest on what happens in the top and bottom quartile, we combine vehicles in the 2nd and 3rd quartiles are into a single category, “med”. The baseline specification is

$$Pr\left(q(f_{it}^b) = s\right) = \beta_0 + \beta_g p_{it}^{gas} + \beta_f f_{it}^k + \beta_{gf} p_{it}^{gas} \times f_{it}^k + \beta_{df} \mathbb{1}\{\Delta f^{kd} > 0\} + \alpha_X X_{it}^k + \varepsilon_{it} \quad (1)$$

where the dependent variable, $Pr\left(q(f_{it}^b) = s\right)$, equals one if f_{it}^b falls within the range of quartile $s \in \{1, med, 4\}$. We also estimate a continuous model where the dependent variable is f_{it}^b, keeping the rest of the specification as presented in equation 1. Fuel economy of the vehicles bought (b) and kept (k) by household i in time t are denoted f_{it}^b and f_{it}^k; i's contemporaneous gas price in t is p_{it}^{gas}, whereas $P_{it}^{gas\_k}$ is the price of gasoline *at the time household i purchased the car that it keeps in time t*. Control variables, denoted X_{it}, include vehicle attributes (e.g. class, make, value, age), nonparametric time controls (year and month-of-year fixed effects) and household/demographic (household fixed effects and county-level unemployment).

The IV specifications deploy instruments for the indicator of the kept vehicle's rank in fuel economy within the portfolio ($\mathbb{1}\{\Delta f^{kd} > 0\}$), the kept vehicle fuel economy (f_{it}^k), and the interaction of gas price and fuel economy ($p_{it}^{gas} \times f_{it}^k$). In each specification, we instrument using the gas price at the time the kept vehicle was purchased ($P_{it_k}^{gas_k}$) and that gas price interacted with the current gas price ($p_{it}^{gas} \times P_{it_k}^{gas_k}$). We augment this set of instruments with the instruments based on vehicle price differences that were briefly described in Section 2.2 on identification to estimate the following system of endogenous variables:

$$\mathbf{Z}_{it} = \left[\begin{array}{ccc} f_{it}^k & p_{it}^{gas} \times f_{it}^k & \mathbb{1}\{\Delta f^{kd} > 0\} \end{array}\right]'$$

We now describe the vehicle price difference instruments precisely. In "Price Difference" specification, we include the difference in the current resale value of the kept and sold vehicles ($\Delta P_{it}^{kd} = P_{it}^k - P_{it}^d$) as an additional instrument. The "Price Difference-in-Difference" specification uses the change in value for the kept and dropped vehicles between the point the vehicle was purchased and the current time period: $\Delta\Delta P_{it}^{kd} = (P_{it}^k - P_{i0}^k) - (P_{it}^d - P_{i0}^d)$.

The third instrument, which we call "Price Deviation Difference-in-Difference", is constructed from the deviation of the difference between the kept and dropped vehicles relative to their expected deprecation rates at the time of the kept car purchase. For each of the kept and dropped vehicle we estimate the households expectation of annual vehicle depreciation using depreciation of similar vehicles over the previous five years. Specifically, for vehicle make m and model year y, and value $V_{m,y,t}$ in year t, the expected depreciation is[6]

$$\mathbf{E}[Dep_{m,y,t}] = \left(\prod_{s=1}^{5} \frac{V_{m,y-s+1,t-s+1} - V_{m,y-s,t-s}}{V_{m,y-s,t-s}} \right)^{\frac{1}{5}} \tag{2}$$

We can then calculate the deviation from this expected depreciation rate for each car in the portfolio, and construct the "Price Deviation Difference-in-Difference" instrument. Assuming vehicle j has resale value $P_{j,t}$ in year t, this is:

$$\Delta\Delta V_{it}^{kd} = (P_{it}^{k} - \mathbf{E}[Dep_{it}^{k}] \cdot P_{i,t-1}^{k}) - (P_{it}^{d} - \mathbf{E}[Dep_{it}^{d}] \cdot P_{i,t-1}^{d}) \tag{3}$$

The set of three price difference instruments is $W = \{\Delta P_{it}^{kd}, \Delta\Delta P_{it}^{kd}, \Delta\Delta V_{it}^{kd}\}$. The first stage thus consists of the following system of three equations for each of the instruments $w \in W$:

$$\begin{aligned} \mathbf{Z}_{it}^{w} = {} & \mathbf{\Gamma}_0 + \mathbf{\Gamma}_g p_{it_k}^{gas} + \theta_P P_{it}^{gas\_k} + \mathbf{\Gamma}_{gg} p_{it}^{gas} \times p_{it_k}^{gas} \\ & + \sum_{dc \in CLASS} \Big(\mathbf{\Gamma}_{dc} \mathbf{1}[CLASS_{it}^{dropped} = sc] \\ & + \mathbf{\Gamma}_w w + \mathbf{\Theta} \mathbf{X}_{it} + \mathbf{\Xi}_{it} \end{aligned} \tag{4}$$

2.3 Results

Work completed in this section helps to fulfill tasks B.1, B.2, B.3 and B.4.

The inquiry into household portfolio preferences is set against a backdrop of fuel economy standards that reduce the fuel intensity of new cars that are introduced into the vehicle fleet. Figure 1 shows deviations in GPM over time in the top (Q4 – most fuel intense) and bottom (Q1 – least fuel intense) quartiles, along with the overall trend in GPM. The overall trend is a decrease in fuel intensity, and this trend is concentrated in gains in the lowest fuel-intensity quartile.

We now present summary statistics that describe the data sample. Table 1 displays the

[6] As a more concrete example, for a household in year $t = 2005$ owning a 2002 Honda Civic, the expected depreciation is the geometric mean annual depreciation rate of 2001 Hondas in 2004, 2000 Hondas in 2003, 1999 Hondas in 2002, etc.

Figure 1: Fuel Intensity (GPM) by Quartile: 2001-2007

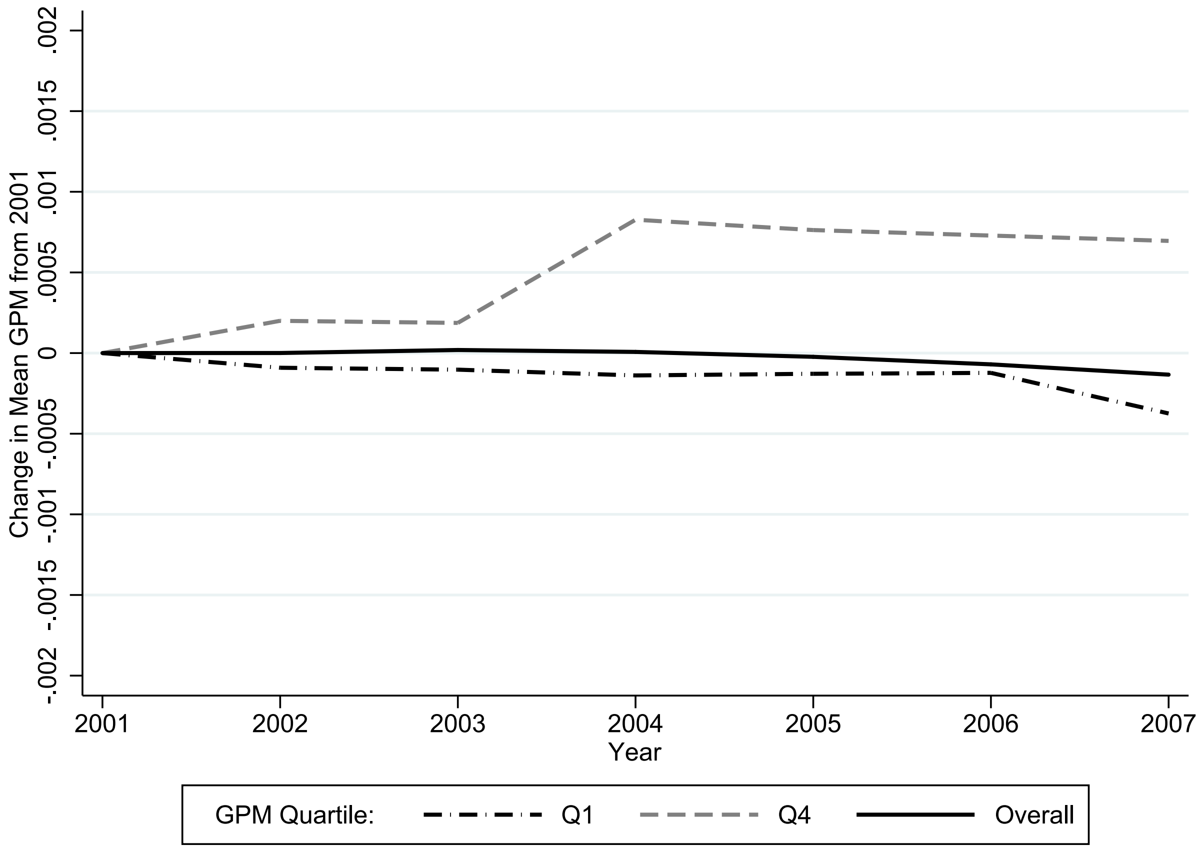

Note: The sample is restricted to 2-car households in California from 2001-2007

number of households that we observe in our data for each portfolio size. As may be expected, roughly 80 percent of households with a vehicle have either one or two cars in 2001. This proportion decreases over time to under 75 percent in 2007. A careful reader may notice that 2003 appears to be somewhat of an outlier. This is due to a data anomaly that occurred between DMV and ARB when the 2003 data were transferred.

Table 2 displays annual vehicle transaction counts and prices, disaggregated into new and used vehicles over the sample period. There are roughly twice as many used vehicle transactions as new, and the price is substantially lower. Used cars sell on average for roughly 40 percent of the price of new cars during the sample period.

The evolution of household vehicle holdings is an important object of interest for this portion of the project. Table 3 shows the distribution of household portfolio transitions. Specifically, rows indicate the number of cars in year t, and columns indicate the number of cars in $t+1$.

Table 1: Household Counts by Portfolio Size (2001-2007)

	Starting Portfolio Size			
Year	1	2	3	4+
2001	4,962,495	3,138,834	1,243,919	687,866
2002	4,764,964	2,926,856	1,221,290	718,505
2003	3,546,316	2,691,691	1,208,954	762,423
2004	4,606,020	2,987,746	1,348,035	929,202
2005	4,820,915	2,897,862	1,427,858	1,127,923
2006	4,665,743	2,950,685	1,453,703	1,203,075
2007	4,547,914	2,956,868	1,471,489	1,259,903

Counts of households by portfolio type in the specified year who are also observed in the following year. Counts limited to households owning 5 or fewer vehicles in the specified year. Years represent DMV reregistration years and run from November 1 of the preceding year to October 31 of the specified year.

Table 2: Vehicle Transaction Count and Mean Price by Year

	Transactions			Mean Price		
	All	New	Used	All	New	Used
2001	5,146,345	1,752,687	3,393,658	15,648	26,273	10,160
2002	5,762,748	1,880,544	3,882,204	16,145	27,127	10,825
2003	5,597,974	1,706,393	3,891,581	16,641	28,643	11,378
2004	3,687,894	1,598,261	2,089,633	19,521	29,669	11,760
2005	5,936,541	1,870,007	4,066,534	17,348	29,164	11,915
2006	5,650,834	1,737,790	3,913,044	17,687	30,200	12,130
2007	5,428,151	1,661,434	3,766,717	18,226	31,450	12,393

Count of the number of vehicle transactions and mean transaction price by year. Years represent DMV reregistration years and run from November 1 of the preceding year to October 31 of the specified year. Vehicles are classified as new if they are the current model year or the preceding model year with less that 500 miles reported on the odometer at the time of registration. Mean prices computed using the reported sale price in the vehicle registration.

The table represents all possible household transitions. The large mass on diagonals indicates that many households do not increase or decrease the number of cars that they register from year to year. It is also clear from Table 3 that a wide range of portfolio transitions occur in any given year.

While the main focus of this paper is to examine the complementarity of vehicle attributes *within* a multi-car household portfolio, one might also be curious about the distribution

Table 3: Number of Unique Households by Portfolio Size

Start Portfolio Size	End Portfolio Size 1	2	3	4+
1	7,262,111	1,360,594	187,558	75,150
2	1,172,278	4,632,425	839,546	259,098
3	168,745	849,703	2,169,948	675,040
4+	35,810	141,618	381,226	1,489,926

Each cell represents the count of unique households from 2001 to 2007 observed to have the starting portfolio size shown in each row and the ending portfolio size shown in the column. These counts provide a measure of the number of households providing identifying variation in each portfolio cell. A single household may appear in multiple cells if their portfolio changes over time but is counted at most once in each cell. For example, two-car household that replaces one car every year would add one to the count of the (2,2) cell. If instead, that household adds a third vehicle in 2004 and returns to a two-car portfolio in 2006 it would add one to the count of the (2,2) cell, one to the count of the (2,3) cell, one to the (3,3) cell, and one to the count of the (3,2) cell. Each household may have zero, one, or multiple vehicle transactions during this time period.

of fuel intensity *across* households.[7] Figure 2 plots a heat map that represents the choice of fuel intensity of kept and bought cars across all households in our two-car replacement sample. Darker shading reflects a higher proportion of households, and it can thus be seen that households tend to locate close to the 45 degree line. This indicates a general desire for similar fuel intensity of cars within a household. This is a strong motivating factor for pursuing the empirical strategy that we do. Changes *within* a household capture the portfolio preferences that are relevant for our policy counterfactuals.

The regressions that follow are estimated using a sample of two-car households that replace one of their cars, a sample which we call "2x2 replacement households".[8] While other transitions are certainly interesting, two-car replacement households provide the cleanest experiment. Households increasing the number of cars in their portfolio are likely to be experiencing an unobserved development that increases their demand for transportation (e.g. having a baby). Furthermore, it is unclear how to characterize the channels through which the portfolio of households with more than two cars affects replacement decisions. Does a portfolio effect for those households operate via the highest-VMT kept car, or the newest?

[7]This is what is referred to in Appendix A as the "type" effect.

[8]We define a household as replacing one vehicle if the starting (in year t) and ending (in year $t+1$ or $t+2$) portfolios differ by one vehicle. The household may conduct multiple vehicle transactions, as long as one of the two vehicles appears in both the starting and ending portfolios. We do not consider households where both vehicles in the two-vehicle portfolio change as the relative timing of each purchase becomes important for defining the portfolio at the time of each vehicle's purchase.

Figure 2: Fuel Intensity of Kept and Bought Cars: 2001-2007

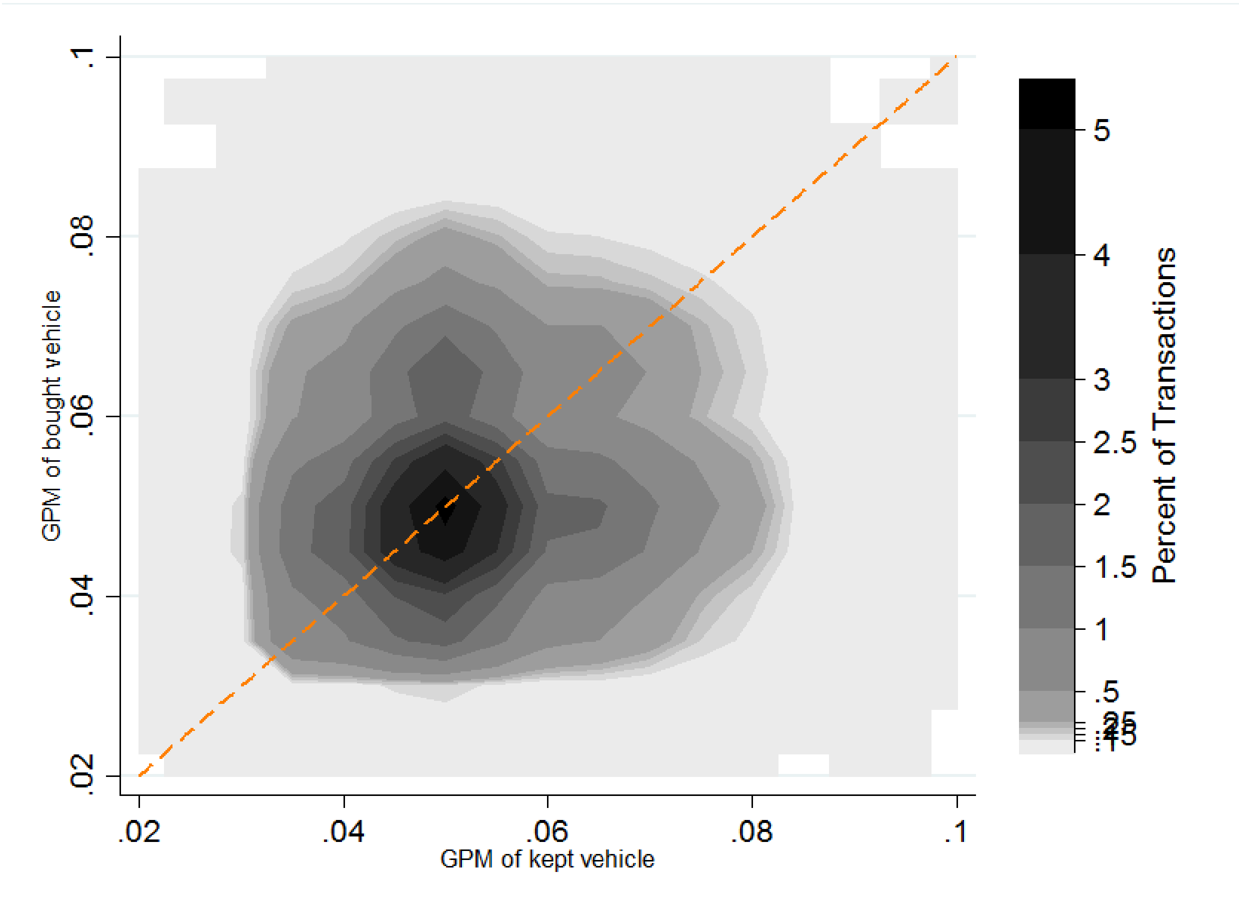

Note: The sample is restricted to 2-car households in California from 2001-2007

Or must the portfolio effect be defined in a higher dimension? Given that no clear answer exists to these questions, we choose the simple path of focusing on the two-car replacement households.

We are also interested in the composition of the household multi-car portfolio. Table 4 shows the frequency of vehicle class pairs in the two-car household portfolio when replacement occurs. Rows indicate the class of the kept vehicle, and rows indicate the class of the bought vehicle. The most common class pairs generally include at least one car, with Car-Truck being the highest frequency. It is also common for Cars to be paired with SUVs. Of non-Car households, Luxury-Luxury and SUV-Truck are the most common. Already, Table 4 reveals that households generally prefer a portfolio comprised of different vehicle classes over one with the same classes. However, we must dig deeper to discover whether the preference for diversification extends to other vehicle attributes.

Table 4: Transaction Counts

		BOUGHT				
		SUV	XUV	Car	Luxury	Truck
KEPT	SUV	43,402	18,666	103,296	58,127	62,167
	XUV	8,586	8,364	31,795	18,782	15,685
	Car	106,628	63,625	376,459	105,209	187,849
	Luxury	64,616	40,864	134,461	101,743	70,719
	Truck	76,317	38,037	223,284	67,378	106,285

Count of transactions for the specified ending portfolio type. Counts limited to the preferred specification estimation sample.

When we deploy the regression methodology with instrumental variables, as described in Section 2.2, we find evidence that households value diversification in fuel economy (or, perhaps more accurately, in vehicle attributes that are correlated with fuel economy). The greater is the fuel economy of the kept car, the lower the fuel economy of the purchased car. We show this using both a continuous measure of fuel economy, as well as by estimating the probability a household purchases a vehicle in the upper and lower quartiles of the fuel economy distribution. Increases in the fuel economy of the kept car reduces the probability the household purchases a car in the lower quartile of gallons per mile, while such increases reduce the probability the household buys a car in the upper quartile. Changes in gasoline prices affect the preferences for diversification in intuitive ways. As gasoline prices increase, the effect of the fuel consumption of kept vehicle and the probability of buying a car in the lower quartile of fuel consumption becomes even more positive. In contrast, as gasoline prices increase, the effect of the fuel consumption of kept vehicle and the probability of buying a car in the upper quartile of fuel consumption becomes even more negative.

Table 5: Regression Estimates - New Vehicle Purchases

(a) Panel A: Bought Vehicle Continuous GPM Measure

	Price Diff (1)	Price DiD (2)	Price Deviation DiD (3)
Gas Price ($/ gal)	0.052	0.031	0.032
	(0.009)**	(0.009)**	(0.010)**
Kept GPM	1.525	0.731	0.349
	(0.398)**	(0.413)	(0.426)
Gas Price × Kept GPM	-0.991	-0.584	-0.596
	(0.177)**	(0.170)**	(0.191)**
Δ GPM > 0	-0.001	0.001	0.000
	(0.001)	(0.001)	(0.001)

(b) Panel B: Bought Vehicle 1st GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	-0.731	-0.228	-0.977
	(0.342)*	(0.322)	(0.447)*
Kept GPM	-3.671	15.711	-22.396
	(15.081)	(14.315)	(18.676)
Gas Price × Kept GPM	13.608	3.929	17.943
	(6.505)*	(6.092)	(8.498)*
Δ GPM > 0	-0.086	-0.119	-0.055
	(0.033)**	(0.031)**	(0.045)

(c) Panel C: Bought Vehicle 2nd/3rd GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	-1.110	-1.570	-0.290
	(0.546)*	(0.585)**	(0.547)
Kept GPM	-63.774	-85.723	-8.027
	(22.248)**	(27.381)**	(23.113)
Gas Price × Kept GPM	21.483	30.147	6.316
	(10.398)*	(11.107)**	(10.386)
Δ GPM > 0	0.129	0.156	0.063
	(0.046)**	(0.050)**	(0.055)

(d) Panel D: Bought Vehicle 4th GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	1.841	1.798	1.267
	(0.488)**	(0.518)**	(0.457)**
Kept GPM	67.445	70.011	30.424
	(19.470)**	(24.383)**	(19.583)
Gas Price × Kept GPM	-35.091	-34.076	-24.259
	(9.290)**	(9.832)**	(8.677)**
Δ GPM > 0	-0.042	-0.037	-0.008
	(0.040)	(0.043)	(0.043)
N	440,809	429,369	348,368
Cragg-Donald stat	145.88	141.44	91.34
Household FE	Yes	Yes	Yes
IV for Kept Vehicle	Base+$\Delta PriceDiD^3$	Base+$PriceDiD^3$	Base+ValueDiD
Subsample	New	New	New

Standard errors robust to heteroskedasticity shown in parentheses.

Table 6: Regression Estimates - Used Vehicle Purchases

(a) Panel A: Bought Vehicle Continuous GPM Measure

	Price Diff (1)	Price DiD (2)	Price Deviation DiD (3)
Gas Price ($/ gal)	0.045	0.027	0.057
	(0.018)*	(0.013)*	(0.022)**
Kept GPM	1.086	0.234	1.532
	(0.806)	(0.504)	(0.909)
Gas Price × Kept GPM	-0.874	-0.523	-1.090
	(0.349)*	(0.245)*	(0.421)**
Δ GPM > 0	0.008	0.008	0.007
	(0.001)**	(0.001)**	(0.001)**

(b) Panel B: Bought Vehicle 1st GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	4.200	3.076	3.081
	(1.137)**	(0.772)**	(1.117)**
Kept GPM	207.785	118.514	109.659
	(50.188)**	(30.916)**	(46.832)*
Gas Price × Kept GPM	-80.034	-58.007	-58.326
	(21.792)**	(14.729)**	(21.382)**
Δ GPM > 0	-0.546	-0.484	-0.489
	(0.061)**	(0.046)**	(0.065)**

(c) Panel C: Bought Vehicle 2nd/3rd GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	-11.152	-7.979	-9.952
	(2.535)**	(1.605)**	(2.671)**
Kept GPM	-485.343	-251.088	-332.787
	(111.819)**	(65.004)**	(111.180)**
Gas Price × Kept GPM	213.213	151.329	189.695
	(48.602)**	(30.613)**	(51.132)**
Δ GPM > 0	0.582	0.432	0.544
	(0.134)**	(0.094)**	(0.151)**

(d) Panel D: Bought Vehicle 4th GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	6.952	4.902	6.871
	(1.618)**	(1.033)**	(1.836)**
Kept GPM	277.558	132.574	223.128
	(71.442)**	(41.549)**	(76.125)**
Gas Price × Kept GPM	-133.179	-93.323	-131.369
	(31.022)**	(19.698)**	(35.146)**
Δ GPM > 0	-0.035	0.052	-0.055
	(0.085)	(0.060)	(0.103)
N	500,882	461,425	364,909
Cragg-Donald stat	39.99	42.48	38.93
Household FE	Yes	Yes	Yes
IV for Kept Vehicle	Base+$\Delta PriceDiD^3$	Base+$PriceDiD^3$	Base+ValueDiD
Subsample	Used	Used	Used

Standard errors robust to heteroskedasticity shown in parentheses.

Tables 5 and 6 present the baseline regression results from new and used car purchases, respectively. The left-most column in each tables corresponds to estimates using the Price Difference IV; column 2 presented estimates using the Price DiD IV; and, finally, Column 3 is our preferred specification, deploying the Price Deviation DiD IV from Equation 4. The four panels correspond to the continuous dependent variable regression (Panel A), the linear probability model (LPM) on the highest fuel economy quartile purchases (Panel B), the linear probability model on second and third fuel economy quartile purchases combined (Panel C), and the linear probability model on purchases of cars in the lowest fuel economy quartile (Panel D).

To more clearly understand the effect of key covariates, we now present and discuss their marginal effects (Table 7). First we discuss marginal effects of gasoline prices on the fuel economy of new and used cars, followed by the marginal effects of kept car GPM on bought car GPM.

Table 7: Marginal Effect of Kept Vehicle MPG on Bought Vehicle GPM - Preferred Specification

(a) Panel A: Bought Vehicle Continuous GPM Measure

Current Gas Price	All (1)	New (2)	Used (3)
Gas Price = 2.00/gal	-0.819	-0.843	-0.648
	(0.140)**	(0.202)**	(0.242)**
Gas Price = 3.00/gal	-1.393	-1.438	-1.738
	(0.235)**	(0.283)**	(0.453)**
Gas Price = 4.00/gal	-1.966	-2.034	-2.828
	(0.369)**	(0.438)**	(0.840)**

(b) Panel B: Bought Vehicle 1st GPM Quartile

	(1)	(2)	(3)
Gas Price = 2.00/gal	-12.832	13.490	-6.992
	(6.902)	(8.860)	(13.288)
Gas Price = 3.00/gal	-32.442	31.433	-65.318
	(11.113)**	(12.649)*	(23.234)**
Gas Price = 4.00/gal	-52.051	49.376	-123.643
	(17.249)**	(19.646)*	(42.632)**

(c) Panel C: Bought Vehicle 2nd/3rd GPM Quartile

	(1)	(2)	(3)
Gas Price = 2.00/gal	71.753	4.605	46.603
	(13.300)**	(11.011)	(30.426)
Gas Price = 3.00/gal	156.553	10.922	236.297
	(21.807)**	(15.438)	(55.240)**
Gas Price = 4.00/gal	241.353	17.238	425.992
	(33.967)**	(23.899)	(102.010)**

(d) Panel D: Bought Vehicle 4th GPM Quartile

	(1)	(2)	(3)
Gas Price = 2.00/gal	-58.920	-18.095	-39.611
	(9.487)**	(9.058)*	(20.474)
Gas Price = 3.00/gal	-124.111	-42.355	-170.980
	(15.808)**	(12.538)**	(37.907)**
Gas Price = 4.00/gal	-189.302	-66.614	-302.349
	(24.757)**	(19.569)**	(70.180)**

Marginal effect of the current gasoline price on the probability a household purchases a vehicle in the GPM quartile specified in table section header. Delta method standard errors robust to heteroskedasticity shown in parentheses. *,**,*** denote results significant at the 10%, 5%, and 1% levels, respectively.

Marginal effects of kept car GPM provide a direct representation of the portfolio effect. We present marginal effects of kept car fuel economy for new and used cars (and pooled) at different gasoline prices in Table 7. Note that both the level of the coefficients within each bought car GPM quartile, and the gradient of these coefficients with respect to the gas price, are important. Positive coefficients reflect an increasing probability of buying in a given quartile as kept car GPM increases (becomes less fuel efficient).

The overall story is clear: households exhibit a preference for GPM diversification in their portfolio, and that preference increases as the gas price rises. This can be seen clearly in the overall results presented in panel A, where negative coefficients imply an increased demand for buying a fuel efficient car as kept-car fuel efficiency decreases. A similar narrative holds when examining the high- and low-GPM bought car quartiles. Households buy new cars in the lowest GPM quartile (high fuel economy) with a higher probability as their kept car fuel economy decreases. Furthermore, this preference increases as the gas price rises. On the other hand, households have a lower probability of buying fuel inefficient new cars as kept car GPM increases, and this preference also increases as the gas price rises.

The situation for used cars is qualitatively similar. The negative coefficient on the marginal effects on 1st GPM quartile implies that the probability of buying a used, highly efficient car decreases in kept-car GPM at this gas price. We presume that this reflects a supply-side effect: households with fuel efficient used cars are less likely to sell them when gas prices are high, thus shifting upwards along the demand curve for this type of car.

In the full paper you can find a full discussion of the implications of this appetite for portfolio diversity. Intuitively, it implies that a policy that exogenously increases the fuel economy of new vehicles will, at the time when those cars are the “kept” cars, induce households to purchase less efficient second vehicles than they otherwise would have. We quantify this prediction in counterfactual simulations. We conjecture that the kept vehicle in a 2-car portfolio becomes 10 percent more efficient, and use the coefficients retrieved from our empirical exercise to simulate how much of the gasoline savings from the efficiency increase are eroded by the portfolio effect. The results for both new and used cars are large. Roughly half (or slightly more) of the anticipated gains from the kept car efficiency increase are eroded by the portfolio effect. While one cannot draw a direct link between this effect and fuel economy standards, the results demonstrate how powerful the portfolio effect can be.

3 Cobenefits of Gasoline Taxes

TASK B.4 was to "rigorously examine how vehicles in California with different characteristics respond differently to changes in gasoline price and vehicle greenhouse gas emissions". We fulfill this task through the work attached in complete paper form in appendix B.[9]

3.1 Data Sources

Our empirical setting is the California personal transportation market. We bring together a number of large data sets. Our analysis is primarily based upon the universe of emissions inspections from 1996 to 2010 from Californias vehicle emissions testing program, the Smog Check Program, which is administered by the California Bureau of Automotive Repair (BAR). A vehicle appears in the data for a number of reasons. First, vehicles more than four years old must pass a Smog Check within 90 days of any change in ownership. Second, in parts of the state (details below) an emissions inspection is required every other year as a pre-requisite for renewing the registration on a vehicle that is six years or older. Third, a test is required if a vehicle moves to California from out-of-state. Vehicles that fail an inspection must be repaired and receive another inspection before they can be registered and driven in the state. There is also a group of exempt vehicles. These are: vehicles of 1975 model-year or older, hybrid and electric vehicles, motorcycles, diesel-powered vehicles, and large natural-gas powered trucks.

These data report the location of the test, the unique vehicle identification number (VIN), odometer reading, the reason for the test, and test results. We decode the VIN to obtain the vehicles make, model, engine, and transmission. Using this information, we match the vehicles to EPA data on fuel economy. Because the VIN decoding is only feasible for vehicles made after 1981, our data are restricted to these models. We also restrict our sample to 1998 and beyond, given large changes that occurred in the Smog Check Program in 1997. This yields roughly 120 million observations.

The Smog Check data report measurements for NO_x and HCs in terms of parts per million and CO levels as a percentage of the exhaust, taken under two engine speeds. As we are interested in the quantity of emissions, the more relevant metric is a vehicles emissions per

[9]We will refer to this work as "ours", acknowledging that it was performed by a subset of the project Investigators.

mile. We convert the Smog Check emissions readings into emissions per mile using conversion equations developed by Sierra Research for the California Air Resources Board in Morrow and Runkle (2005). The conversion equations are functions of both measurements of all three pollutants, vehicle weight, model year, and truck status.

As part of our simulation exercise, we also use data obtained from CARFAX Inc. to estimate scrappage decisions. These data contain the date and location of the last record of the vehicle reported to CARFAX for 32 million vehicles in the Smog Check data. This includes registrations, emissions inspections, repairs, import/export records, and accidents. Because the CARFAX data include import/export records, we are able to correctly classify the outcomes of vehicles which are exported to Mexico as censored, rather than scrapped, thus avoiding the issues identified in Davis and Kahn (2011).

For a subset of our Smog Check data, we are able to match vehicles to households using confidential data from the California Department of Motor Vehicles (DMV). These data track the registered address of the every vehicle in the state, with one address given for each year. We use the registration information to attach demographic information on income from U.S. Census data. The DMV data were only available for the years 2000 to 2008.

For a portion of our analysis, we use data from the 2009 National Highway Transportation Survey, which contains information on household vehicles, annual VMT, and household income for a sample of households.Finally, we use gasoline prices from EIAs weekly California average price series to construct average prices between inspections.

3.2 Methodology

The basic methodology here is to use California smog check inspection data to estimate how the usage and emissions from different vehicle types responds to changes in gasoline prices. We then use this relationship to estimate the effectiveness of gasoline taxes in reducing local pollutants: carbon monoxide (CO), hydrocarbons (HCs), and nitrogen oxides (NO_x). This exercise assumes a symmetry in response to prices caused by movements in the gasoline price and those caused by changes in a hypothetical gasoline tax.[10] One caveat in this analysis is that while we use the best conversion equations we are aware of (developed by

[10] Li et al. (2014) present evidence that consumers respond more to gasoline taxes than to market-based changes in gasoline prices. However, due to an absence of natural experiments that change gasoline taxes, we rely on price fluctuations as the source of identifying variation.

Sierra Research for ARB), using VMT data to estimate gasoline consumption and emissions does require further assumptions. For example, consumer driving behavior can affect both on-road fuel economy and emissions profiles. This can include factors such as the average speed, braking behavior, and cold starts. We are not aware of a reason why the results should be systematically skewed using the average values that were developed for ARB, but we recognize that on-road real-world driving conditions may differ and additional work could be done to further refine this mapping from VMT to emissions.

In this work we distinguish between "first-best" and "second-best" optimal tax regimes. The first-best earns its name by setting taxes on all pollutants such that the level of the tax equals the level of marginal damages. For example, this would imply setting a tax on carbon dioxide equal to the social cost of carbon, and separate taxes on criteria pollutants (again equal to their marginal damage). A second-best tax in the vehicle setting would be a tax on gasoline. It is not generally first-best optimal, since it does not equate the tax level with marginal social damages.

In the case of the local pollution externalities of driving, the relationship between the second-best optimal gasoline tax and the first-best Pigouvian emissions tax depends on three empirical relationships: the distribution of pollution externalities across vehicles; the extent to which gasoline prices affect the implicit demand for pollution; and the correlation between vehicle-specific demand responses and externality levels. If vehicles do not differ in their sensitivity of vehicle miles traveled (VMT) to gasoline prices—an object that we hereafter call the "VMT elasticity"—the second-best optimal (hereafter "SBO") gasoline tax will simply be the average per-gallon externality across all vehicles. However, if price responsiveness and externalities are correlated, Diamond (1973) shows that the SBO gasoline tax will be the weighted average of vehicle per-gallon externalities, where the weights are the price derivatives of the vehicle-specific gasoline demand curves. In our empirical work, we allow for the VMT elasticity to vary depending on a vehicle's emissions per mile traveled, which we observe.

An important empirical result in our paper is that we find that vehicle-level emissions are correlated with vehicle-specific VMT elasticities; dirtier vehicles are more price responsive. Using detailed vehicle-specific data on miles driven, we show a positive correlation between criteria pollutant emissions and the VMT elasticity (in absolute value) holds for all three pollutants for which we have data: carbon monoxide (CO), hydrocarbons (HCs), and nitrogen oxides (NO_x). VMT elasticities are also positively correlated with greenhouse gas emissions and vehicle weight.

These correlations drive a wedge between the SBO gasoline tax associated with emissions and what we call the "naive" tax, which we define as the the tax based only on the *unweighted*-average externality across vehicles. We show the SBO gasoline tax is larger, on the order of 50 percent, than the naive gasoline tax in each of the years of our sample. However, we also show that even when instituting the SBO gasoline tax, the tax performs poorly in eliminating DWL, and only marginally better than the naive gasoline tax. Across our sample, we estimate the SBO gasoline tax eliminates only 30 percent of DWL associated with the pollutants studied.

We investigate three sources of the documented heterogeneity, which are not necessarily mutually exclusive. First, it may be driven entirely by a vintage effect. That is, older vehicles are both more responsive to changes in gasoline prices and have higher emissions. Second, it might be driven by differences in the incomes of consumers that drive dirtier versus cleaner vehicles.[11] Third, it may result from households shifting which of their vehicles are driven in the face of rising gasoline prices.

To investigate whether it is simply a vintage effect, we redefine the quartiles based on the distribution of emissions within vintage and calendar year bins. We split vehicles into three age categories: 4 to 9 years old, 10 to 15 years old, and 16 to 27 years old. We are able to group a subsample of our Smog Check vehicles into households using our registration data from California DMV. Because of the process used to identify households in these data, this subsample likely draws more heavily from households residing in single-family homes.

3.2.1 Regression Specifications

We estimate how changes in gasoline prices affect decisions about vehicle miles traveled (VMT), and how this elasticity varies with vehicle characteristics. For each vehicle receiving a biennial smog check, we calculate average daily miles driven and the average gasoline price during the roughly two years between smog checks. Obviously vehicle owners with more fuel efficient vehicles will respond less to changes in the per-gallon gasoline price, and to abstract from this we specify the elasticity with respect to the price in dollars per mile (DPM), by dividing the average per gallon price by fuel economy in gallons per mile. Thus, the price faced by each vehicle's owner will vary both with the exact period in between Smog Checks, and with the specific vehicles' fuel economy. We then allow the elasticity to vary based on

[11] West (2004) also documents a positive correlation between income and emissions. She does not separately estimate elasticities, however.

the emissions of the vehicle. We begin by estimating:

$$\ln(VMT_{ijgt}) = \beta \ln(DPM_{ijgt}) + \gamma D_{truck} + \omega time + \mu_t + \mu_j + \mu_g + \mu_v + \varepsilon_{igt} \qquad (5)$$

where i indexes vehicles, j vehicle-types, g geographic locations, t time, and v vehicle age, or vintage. DPM_{ijgt} the average gasoline price per mile faced by vehicle i between time t and the date of the previous smog check, D_{truck} is an indicator variable for whether the vehicle is a truck, $time$ is a time trend, and ε_{igt} is a residual.[12] Our baseline specification assumes that gasoline prices are exogenous to individual driving decisions. Such an assumption is common in the literature, as gasoline prices are largely driven by movements in the world price of crude oil, which saw dramatic changes during the 2000s for reasons unrelated to driving choices in California.[13] However, we have also estimated our main analyses instrumenting for DPM with the Brent Crude oil price, and we obtain very similar results.

We begin by including demographic characteristics by the zip code of smog checks, and year and vintage fixed effects. We then progressively include finer vehicle-type fixed effects by including make, then make/model/model-year/engine, and finally individual vehicle fixed effects. We also differentiate the influence of gasoline prices by vehicle attributes related to the magnitude of their negative externalities—criteria pollutants, CO_2 emissions, and weight.

We allow the VMT elasticity to vary with the magnitude of our externalities in two ways. For both approaches, we begin by ranking vehicles within each calendar year by their emissions per mile of NO_x, HCs, CO, fuel economy, or vehicle weight in pounds. In one set of specifications we split vehicles up by the quartile of these variables and allow each quartile to have a separate β. In another set, we include a linear interaction of the percentiles of these variables and the log of gasoline prices in dollars per mile.

3.3 Results

Results from this portion of the project contribute to fulfilling TASK B.4.

Tables 8 shows our results, focusing on NO_x. Moving from left to right along the columns,

[12] The fuel economy in gallons per mile used to calculate our DPM variable uses the standard assumption that 45 percent of a vehicle's miles driven are in the city and 55 percent are on the highway. This is the standard approach used by the EPA for combined fuel economy ratings.

[13] See, for example, Busse et al. (2013) (BKZ).

Models 1 to 6 correspond to different control variables, as demonstrated by the presence of coefficients or by the fixed effects noted in the table itself. The changes from Models 1 to 5 illustrates the importance of controlling for vehicle-type fixed effects. Initially, the average elasticity falls from -0.265 to -0.117 when including fixed make effects, but then rises when including finer detailed vehicle fixed effects. Our final specification includes individual vehicle fixed effects yielding an average elasticity of -0.147.[14] In Model 6 we examine heterogeneity with vehicle fixed effects. Model 6 includes interactions with quartiles of NO_x, as in Model 3. The VMT elasticity for the cleanest vehicles, quartile one, is positive at 0.041, while the VMT elasticity for the dirtiest vehicles is twice the average elasticity at -0.288. To put these numbers in context, the average per-mile NO_x emissions of a quartile one vehicle is 0.163 grams, while the average per-mile NO_x emissions of a quartile four vehicle is 1.68 grams.

We find similar patterns across the other externalities. The range of the estimated VMT elasticities is somewhat larger when using quartiles of HCs and CO emissions compared to NO_x, with the dirtiest quartiles around -0.30 and the cleanest around 0.05. For CO_2 the cleanest vehicles are those with the highest fuel economy, and here we see the least fuel-efficient vehicles having a VMT elasticity of -0.183, compared to -0.108 for vehicles with fuel economy in the highest quartile. We observe some heterogeneity in the VMT elasticity across vehicle weights as well, although it is smaller than the other externalities.

[14]Our average elasticity is larger than that found in Hughes et al. (2008) reflecting the longer run nature of their elasticity.

Table 8: Vehicle Miles Traveled, Dollars Per Mile, and Nitrogen Oxides (Quartiles by year)

	(1) Model 1	(2) Model 2	(3) Model 3	(4) Model 4	(5) Model 5	(6) Model 6
ln(DPM)	-0.269**	-0.123**	-0.183**	-0.134**		-0.038
	(0.044)	(0.038)	(0.027)	(0.022)		(0.028)
ln(DPM) * NO_x Q1					0.043*	
					(0.021)	
ln(DPM) * NO_x Q2					-0.054*	
					(0.022)	
ln(DPM) * NO_x Q3					-0.152**	
					(0.025)	
ln(DPM) * NO_x Q4					-0.280**	
					(0.028)	
ln(DPM)*NO_x Centile						-0.001**
						(0.000)
NO_x Q2					0.216	
					(0.663)	
NO_x Q3					-1.742	
					(0.881)	
NO_x Q4					-2.417*	
					(1.003)	
NO_x Centile						-0.001
						(0.001)
Truck	0.054	0.057	0.005			
	(0.033)	(0.045)	(0.055)			
Time Trend	-0.244**	-0.314**	-0.278**	-0.035	-0.057	-0.062*
	(0.037)	(0.024)	(0.015)	(0.028)	(0.032)	(0.025)
Time Trend-Squared	0.002**	0.002**	0.002**	0.000	0.000	0.001*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Vintage Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Make Fixed Effects	No	Yes	No	No	No	No
Vin Prefix Fixed Effects	No	No	Yes	No	No	No
Vehicle Fixed Effects	No	No	No	Yes	Yes	Yes
Observations	36387455	36387455	36387455	36387455	29779909	29779909
R-squared	0.210	0.218	0.143	0.121	0.117	0.118

* $p < 0.05$, ** $p < 0.01$

Notes: Each observation is a vehicle's Smog Check inspection. Dependent variable is the log of average daily vehicle miles travelled since the previous inspection. DPM represents the average gasoline price over the period since the previous inspection, converted to dollars per mile by dividing by vehicle fuel economy Quartiles and centiles of NO_x are based on rankings of emissions per mile within the calendar year in which the Smog Check occurs. Standard errors clustered by vehicle make reported in parentheses.

3.3.1 The Source of the Heterogeneity

While the SBO gasoline tax is not (necessarily) affected by the mechanism behind the heterogeneity, it is of independent interest to investigate the mechanism.[15] We investigate three sources, which are not necessarily mutually exclusive. First, it may be driven entirely by a vintage effect. That is, older vehicles are both more responsive to changes in gasoline prices and have higher emissions. Second, it might be driven by differences in the incomes of consumers that drive dirtier versus cleaner vehicles. Third, it may result from households shifting which of their vehicles are driven in the face of rising gasoline prices.

To investigate whether it is simply a vintage effect, we redefine the quartiles based on the distribution of emissions within vintage and calendar year bins. We split vehicles into three age categories: 4 to 9 years old, 10 to 15 years old, and 16 to 27 years old.

Table 9 reports the results for heterogeneity over NO_x emissions.[16] These results suggest that while vintage is a factor in the externality-based heterogeneity, it is not the only source or even the most important source. While middle-aged and older vehicles are more elastic than new vehicles on average, within age bin there is still substantial heterogeneity. For new vehicles, the difference between the dirtiest and cleanest quartiles is two thirds of the range for the whole sample. Middle-aged vehicles have three quarters as much range, and the oldest vehicles, 16 years and older, have a range nearly as large as for the whole sample.

[15] In the presence of other second-best policies, this need not be the case. We abstract away from those issues here.

[16] Results for the other four externality types are quite similar.

Table 9: Vehicle Miles Traveled, Dollars Per Mile, and Nitrogen Oxides (Quartiles by age range)

	(1) 4-9	(2) 10-15	(3) 16-27	(4) 4-9	(5) 10-15	(6) 16-27
ln(DPM)	-0.017	-0.127**	-0.132**			
	(0.011)	(0.027)	(0.041)			
ln(DPM) * NO_x Q1				0.127**	-0.007	0.052
				(0.022)	(0.023)	(0.035)
ln(DPM) * NO_x Q2				0.041**	-0.079**	-0.035
				(0.011)	(0.021)	(0.041)
ln(DPM) * NO_x Q3				-0.014	-0.160**	-0.144**
				(0.013)	(0.030)	(0.044)
ln(DPM) * NO_x Q4				-0.104**	-0.261**	-0.263**
				(0.014)	(0.030)	(0.039)
NO_x Q2				0.315	-1.699**	-2.291**
				(0.633)	(0.388)	(0.801)
NO_x Q3				-0.416	-3.119**	-5.145**
				(1.105)	(0.607)	(0.929)
NO_x Q4				-2.824	-5.562**	-6.269**
				(1.497)	(0.678)	(1.104)
Time Trend	0.207**	0.156**	-0.089	0.373**	0.060	-0.315**
	(0.041)	(0.033)	(0.081)	(0.067)	(0.031)	(0.087)
Time Trend-Squared	-0.001**	-0.001**	0.001	-0.003**	-0.000	0.003**
	(0.000)	(0.000)	(0.001)	(0.001)	(0.000)	(0.001)
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Vintage Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Vin Prefix Fixed Effects	No	No	No	No	No	No
Vehicle Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	15466449	15352009	5568997	12143358	12684397	4952154
R-squared	0.120	0.097	0.093	0.120	0.099	0.091

* $p < 0.05$, ** $p < 0.01$

Notes: Each observation is a vehicle's Smog Check inspection. Dependent variable is the log of average daily vehicle miles travelled since the previous inspection. DPM represents the average gasoline price over the period since the previous inspection, converted to dollars per mile by dividing by vehicle fuel economy Quartiles and centiles of NO_x are based on rankings of emissions per mile within the calendar year in which the Smog Check occurs. Standard errors clustered by vehicle make reported in parentheses.

We are able to group a subsample of our Smog Check vehicles into households using our registration data from California DMV. Because of the process used to identify households in these data, this subsample likely draws more heavily from households residing in single-family homes. Given this selection and the fact that the sample period differs from our base specification, it is not surprising that we find average elasticities that differ from those presented above.

Table 10 presents the results from this subsample. For this sample, we construct two additional variables meant to capture the household stock of vehicles. The variable "Higher MPG in HH" equals one if there is another vehicle in the household whose MPG rating places it in a higher quartile than the vehicle in question. Likewise, the variable "lower MPG in HH" equals one if there is another vehicle in the household whose MPG rating places it in a lower quartile than the vehicle in question.

If households shift usage from low-MPG vehicles to high-MPG vehicles, we would expect "Higher MPG in HH" to be negative and "Lower MPG in HH" to be positive. Column 2 of Table 10 adds these variables to our base specification. The point estimates suggest that a vehicle in the highest fuel economy quartile belonging to a household that also has a lower fuel economy vehicle has an elasticity greater than a third lower. We cannot reject the null hypothesis that the sum of the interactions with quartile four and "Higher MPG in HH" is zero.[17]

[17]The sum of the two vehicle-stock variables is positive, but because lower fuel efficient vehicles are driven more earlier in the sample, the elasticities are not comparable in terms of what they imply for total miles driven.

Table 10: VMT Elasticity for a Sample of Households, 2000-2008

	(1)	(2)	(3)
ln(DPM) * MPG Q1	-0.0881***	-0.0902***	-0.0970***
	(0.0140)	(0.0133)	(0.0174)
ln(DPM) * MPG Q2	-0.0882***	-0.0910***	-0.0976***
	(0.0142)	(0.0132)	(0.0173)
ln(DPM) * MPG Q3	-0.0324	-0.0385*	-0.0448*
	(0.0170)	(0.0158)	(0.0194)
ln(DPM) * MPG Q4	-0.0263	-0.0339	-0.0401
	(0.0181)	(0.0174)	(0.0228)
ln(DPM) * Higher MPG in HH		-0.0290***	-0.0296***
		(0.00634)	(0.00632)
ln(DPM) * Lower MPG in HH		0.0628***	0.0626***
		(0.00715)	(0.00717)
Higher MPG in HH		-0.0739***	-0.0753***
		(0.0138)	(0.0137)
Lower MPG in HH		0.150***	0.149***
		(0.0186)	(0.0186)
ln(DPM) * HH Income Q2			-0.00514
			(0.00858)
ln(DPM) * HH Income Q3			0.00763
			(0.00860)
ln(DPM) * HH Income Q4			0.0224
			(0.0125)
Year Fixed Effects	Yes	Yes	Yes
Vintage Fixed Effects	Yes	Yes	Yes
Demographics	Yes	Yes	Yes
Vehicle Fixed Effects	Yes	Yes	Yes
Observations	7549359	7549359	7549359
R-squared	0.113	0.113	0.113

* $p < 0.05$, ** $p < 0.01$

Notes: Each observation is a vehicle's Smog Check inspection. Dependent variable is the log of average daily vehicle miles travelled since the previous inspection. DPM represents the average gasoline price over the period since the previous inspection, converted to dollars per mile by dividing by vehicle fuel economy Quartiles MPG and Household Income are based on rankings within the calendar year in which the Smog Check occurs. Higher (lower) MPG in household indicates that there is another vehicle in the household whose fuel economy is in a higher (lower) quartile than the vehicle in question. Standard errors clustered by vehicle make reported in parentheses.

For this same sample of vehicles, we also use U.S. Census information based on zip-code of residence to categorize owners into income quartiles. We interact these quartiles with the log of DPM to see if differences in elasticities exist. Column 3 of Table 10 adds these interaction terms. There is some evidence that higher-income consumers are less elastic. However, the emissions quartile effects persist; vehicles in the bottom quartile remain nearly three times more sensitive even after accounting for income differences.

Our Smog Check data report the zip code of the testing station the vehicle visited. For our more general sample, we also use this information to construct measures of income. Table 11 compares these results with the DMV data. We find similar differences in the elasticities, despite the smaller average elasticity.

Table 11: VMT Elasticity by Income Quartile, 2000-2008

	40% Sample of HHs	HHs, with HH FE	10% Sample of VINs
ln(DPM) * HH Income Q1	-0.0659**	-0.0560***	-0.0524*
	(0.0197)	(0.00560)	(0.0216)
ln(DPM) * HH Income Q2	-0.0706***	-0.0591***	-0.0559*
	(0.0142)	(0.00555)	(0.0237)
ln(DPM) * HH Income Q3	-0.0588***	-0.0451***	-0.0679*
	(0.0165)	(0.00552)	(0.0261)
ln(DPM) * HH Income Q4	-0.0461**	-0.0316***	-0.0527
	(0.0151)	(0.00553)	(0.0274)
Year Fixed Effects	Yes	Yes	Yes
Vintage Fixed Effects	Yes	Yes	Yes
Demographics	Yes	Yes	Yes
Vehicle Fixed Effects	Yes	No	Yes
Observations	7549359	7549359	2489373
R-squared	0.112	0.165	0.107

* $p < 0.05$, ** $p < 0.01$

Notes: Each observation is a vehicle's Smog Check inspection. Dependent variable is the log of average daily vehicle miles travelled since the previous inspection. DPM represents the average gasoline price over the period since the previous inspection, converted to dollars per mile by dividing by vehicle fuel economy Quartiles MPG and Household Income are based on rankings within the calendar year in which the Smog Check occurs. Standard errors clustered by vehicle make reported in parentheses.

3.3.2 Efficiency of the second best optimal gasoline tax

In this section, we consider the efficiency of using a SBO gasoline tax to abate the externalities caused by driving, specifically those resulting from emissions of NO_x, HCs, and CO. We begin by calculating both the naive and SBO gasoline tax, and then compare the remaining DWL left over from these second-best taxes to the optimal outcome obtained by a vehicle-specific Pigouvian tax on emissions.

Second best optimal gasoline tax

We calculate the naive tax per gallon of gasoline as the simple average of the externality per gallon caused by all vehicles on the road in California in a particular year. We value the externalities imposed by NO_x and HCs using the marginal damages calculated by Muller and Mendelsohn (2009), based on the county in which each vehicle has its smog check.[18] The damages calculated by Muller and Mendelsohn (2009) are ideal for this purpose, as they use an integrated assessment model to capture how a marginal unit of NO_x or HCs emitted in one location causes damages throughout the United States, both directly and through the formation and removal of ozone and particulate matter. For CO, we use the median marginal damage estimate from Matthews and Lave (2000).

Let the marginal damage per gram of pollutant p in county c be θ_c^p, with emissions rates in grams per mile by vehicle i of ϵ_i^p. Then the externality per mile of vehicle i, E_i is:

$$E_i = \theta_c^{HC} \cdot \epsilon_i^{HC} + \theta_c^{NO_x} \cdot \epsilon_i^{NO_x} + \theta_c^{CO} \cdot \epsilon_i^{CO} \tag{6}$$

The naive tax in year y will then be:

$$\tau_{naive}(y) = \frac{1}{N^y} \sum_{i=1}^{N^y} \frac{E_i}{MPG_i}, \tag{7}$$

where N^y denotes the number of vehicles on the road in year y, and MPG_i denotes the fuel economy rating of vehicle i. In practice, since the stock of vehicles represented in the Smog Check data in any given year will be less than the total stock of vehicles in the vehicles fleet, we weight each Smog Check observation by the frequency with which vehicles of the same

[18]Note that the values used in this paper differ from those used in the published version of Muller and Mendelsohn (2009). The published values were calculated using incorrect baseline mortality numbers that were too low for older age groups. Using corrected mortality data increases the marginal damages substantially. We are grateful to Nicholas Muller for providing updated values, and to Joel Wiles for bringing this to our attention.

vintage and class appear in the California fleet as a whole.

We calculate the SBO gasoline tax, taking into account the heterogeneity in both levels of the externality and the responsiveness to gasoline prices. We estimate a regression similar to Equation (5), but allowing the elasticity of VMT with respect to DPM to vary over all our dimensions of heterogeneity. Let the group-specific elasticity for vehicle i be β_i^q, where q indexes cells by HC emissions, NO_x emissions, CO emissions, MPG, weight, and age, with the externalities again in quartiles by year. Further, let the average price per gallon and the quantity of gasoline consumed per year in gallons in year y be P_i^y and Q_i^y, respectively. Then the optimal tax in year y will be

$$\tau^*(y) = \frac{-\sum_h \sum_{i \in h} \frac{\partial U^h}{\partial \alpha_i} \alpha'_i}{\sum_h \alpha'_h}, \tag{8}$$

with

$$\alpha'_i = -\beta_i^q \cdot \frac{Q_i^y}{P_i^y}, \tag{9}$$

and

$$\frac{\partial U^h}{\partial \alpha_i} = \frac{E_i}{MPG_i}. \tag{10}$$

Table 12 shows the naive and SBO taxes for each year from 1998 to 2008. The naive tax would be 61.2 cents per gallon of gasoline consumed in 1998, while the SBO tax is 86 cents, 39 percent higher. The ratio of the naive and SBO gasoline tax increases even as the level of the externalities declines over time. From 2002 on, the SBO gasoline tax is at least 50 percent larger than the naive tax in each year.

We also account for vehicle owners' decisions to scrap their vehicles to the extent these are affected by gasoline prices. The full paper appendix discusses the details and results of this exercise. To summarize, we allow gasoline price to affect scrappage decisions, and allow this to vary over emissions profiles and vintages. We find that the main source of heterogeneity occurs across vintages; specifically, increases in gasoline prices increase the hazard rate of very old vehicles, but decrease the hazard rate of middle-aged vehicles. Because emissions of criteria pollutants are positively correlated with age, this has the effect of decreasing criteria pollutants, although the aggregate effect is small.

Welfare with Uniform Taxes

We have shown that because of the correlation between elasticities and externality rates, the SBO gasoline tax is much higher than the naive tax calculated as the average of per-gallon

Table 12: Average and Marginal Pollution Externality

	Average Externality (¢/gal)	Marginal Externality (¢/gal)
1998	61.48	91.27
1999	54.78	81.62
2000	48.55	74.31
2001	40.96	64.29
2002	34.18	54.09
2003	28.77	46.89
2004	24.31	39.26
2005	21.25	33.95
2006	18.61	29.52
2007	16.23	25.81
2008	14.36	22.84

Notes: Average Externality is the simple average of damages from emissions of criteria pollutants produced by each car in each year, divided by fuel usage. We refer to a tax on the average externality as the "naive tax". The marginal externality is computed as the weighted average of externality per gallon, using the negative slope of the vehicle's demand curve as the weight. A tax on the marginal externality is the SBO gasoline tax. Both calculations also weight vehicles by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole. Dollar figures inflation adjusted to year 2008.

externalities. We now turn to the question of how much the SBO gasoline tax improves welfare beyond what is achieved by the naive tax. We note again that even the optimal uniform tax is still a second-best policy.[19] Because of the heterogeneity in externality levels, the most polluting vehicles will be taxed by less than their external costs to society, leaving remaining DWL. Vehicles that are cleaner than the weighted average will be taxed too much, overshooting the optimal quantity of consumption and creating more DWL.

In each of the following analyses, we compare the remaining DWL resulting from the local pollution externality with both the naive and SBO gasoline tax to the DWL without any additional tax.

3.3.3 Simulation Results

We begin by approximating the ratios of DWL with and without the taxes using our data to simulate the change in miles driven and thus in gasoline consumption from a tax. Let $miles_i^y$ be the actual average miles per day traveled by vehicle i between its last smog check

[19] We have also repeated the analysis under the assumption that policy makers adopt the second-best optimal VMT tax. The degree of DWL that remains is only slightly reduced.

and the current one, observed in year y, and let $\hat{miles}_i^y(\tau)$ be the miles per day that a vehicle would travel if the average price of gasoline were raised by a tax of $\$\tau$ per gallon that is fully passed through to consumers. We approximate DWL as a triangle, such that the ratio of interest is:

$$r(\tau) = \frac{\sum_i \left(\frac{1}{2} \cdot \frac{miles_i^y - \hat{miles}_i^y(\tau)}{MPG_i} \cdot \frac{E_i}{MPG_i} - \tau \right.}{\sum_i \left(\frac{1}{2} \cdot \frac{miles_i^y - \hat{miles}_i^y(\frac{E_i}{MPG_i})}{MPG_i} \cdot \frac{E_i}{MPG_i} \right.}$$

The fully optimal tax would have a ratio of 0, while a tax that actually increased the DWL from gasoline consumption would be greater than 1. Table 13 shows these ratios for various taxes. The first two columns show ratios for a statewide tax based on the average and marginal externalities, respectively, of all vehicles in California in each year. Deadweight loss from the naive tax averages 72.8 percent of DWL with no additional tax over the sample period, and rises over time as the fleet becomes cleaner. The SBO gasoline tax is little better, averaging 69.8 percent of DWL with no tax during our sample period.[20]

Is it even possible to effectively abate local pollution externalities using a tax on gasoline, or is there too much idiosyncratic variation in externality levels for this to be possible? That is, if hypothetically the tax were allowed to vary by groups observable to policymakers, would the SBO uniform tax perform better? Obviously, this may be politically infeasible depending on how the groups are defined and impractical to implement. The purpose of this analysis is to explore the nature of the failure of the uniform gasoline taxes.

The remaining columns of Table 13 show remaining deadweight loss from the naive and SBO forms of taxes by groups. The marginal damages from Muller and Mendelsohn (2009) are designed to vary at the county level, and within California they vary substantially across counties, due to both baseline emissions levels and the extent to which population is exposed to harmful emissions. As such, a county-specific tax on emissions might be expected to target externality levels more precisely. The third and fourth columns of Table 13 shows the DWL ratios for a naive and SBO gasoline tax computed this way, and it turns out there is relatively

[20]In unreported results, we also used Census income to analyze the regressivity of the two gasoline taxes. We find, consistent with Poterba (1991), that the taxes are initially progressive through the lowest income deciles, but then become regressive. We also find that, while the tax expense for lower income consumers is a larger share of their income, lower income consumers pay a smaller share of the externality that they generate. High income consumers pay more in taxes than the externality that they generate, while low income consumers pay less in taxes due to fact that lower income consumers, on average, drive vehicles with higher emissions rates and thus reduce their miles traveled more in response to a gasoline tax.

Table 13: Ratios of DWL with Tax to DWL With No Tax

	Statewide Tax		County-Level Taxes		Vintage Tax		County/Vintage Tax		Total DWL
	Naive	SBO	Naive	SBO	Naive	SBO	Naive	SBO	$
1998	0.616	0.568	0.573	0.523	0.348	0.341	0.296	0.290	196466745.4
1999	0.636	0.577	0.592	0.529	0.330	0.325	0.269	0.265	158104024.8
2000	0.635	0.583	0.587	0.532	0.320	0.317	0.253	0.251	131221907.6
2001	0.690	0.627	0.649	0.582	0.348	0.345	0.281	0.279	100426398.8
2002	0.700	0.675	0.652	0.625	0.348	0.346	0.284	0.283	76704235.2
2003	0.716	0.699	0.661	0.643	0.316	0.314	0.248	0.247	58869860.6
2004	0.746	0.740	0.699	0.693	0.313	0.312	0.246	0.245	42633365.5
2005	0.766	0.762	0.723	0.718	0.319	0.318	0.250	0.250	27431776.9
2006	0.801	0.796	0.762	0.757	0.338	0.337	0.272	0.271	20756466.1
2007	0.817	0.817	0.780	0.780	0.328	0.327	0.259	0.357	15589665.8
2008	0.838	0.836	0.805	0.802	0.331	0.331	0.264	0.264	12340287.7
Average	0.724	0.698	0.680	0.653	0.331	0.329	0.266	0.273	76413157.7

Notes: DWL with no tax calculated based on the difference in emissions from imposing a tax equal to the actual externality per gallon consumed by a particular car. SBO tax computed as the weighted average of externality per gallon, using the negative slope of the vehicle's demand curve as the weight. All taxes also weight vehicles by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole.

little improvement. In other words, county-by-county variation in emissions and elasticities does not explain the failure of a single, uniform tax to remove a substantial amount of deadweight loss. The average ratios over our sample are 0.680 for the naive tax and 0.653 for the optimal uniform tax.

Since emissions rates are highly correlated with vintage, another approach would be to allow taxes to vary by the age of the vehicle.[21] The fifth and sixth columns of the table show this, and here we see a substantial improvement: 0.331 for the naive tax and 0.329 for a SBO gasoline tax. Combining these and having the tax vary by both vintage and location, shown in the last two columns, reduces the ratios to 0.266 and 0.273, respectively.

This analysis shows two striking results. First, a SBO gasoline tax does a terrible job of addressing the market failure from pollution externalities. The dirtiest vehicles are not taxed enough, and many clean vehicles are over-taxed. This is true even when the uniform tax is calculated taking the correlation between emissions and VMT elasticities into account. The roughly 50 percent increase in the tax level from a SBO gasoline tax correctly abates

[21]Such a system could be built within the Smog Check Program, with vehicle-specific taxes based on mileage since the previous test.

more emissions from the dirtiest vehicles, but also over-taxes the cleanest vehicles by a larger amount. The welfare benefits of the SBO gasoline tax are around 10% higher than those from a naive tax, but still fall far short of the benefits from a true Pigouvian tax linked to actual vehicle emissions. The number of vehicles for which the uniform tax overshoots is remarkable. Table 14 shows the proportion of vehicle-years over the 11 years of our sample for which each tax overshoots. Because the distribution of emissions is so strongly right skewed, the naive uniform tax overshoots for more than 72 percent of vehicles, and the optimal uniform tax for even more. Second, there is enough heterogeneity in the distribution of the per-gallon externality that even a tax targeting broad groups leaves a substantial portion of DWL.

Table 14: Proportion of Vehicles for which a Uniform Tax Overshoots the Optimal Tax

	Mean
Naive Tax on Fleet Average Externality	0.724
SBO Tax on Fleet Marginal Externality	0.803
Naive Tax on County Average Externality	0.714
SBO Tax on County Marginal Externality	0.793
Naive Tax on Vintage Average Externality	0.708
SBO Tax on Vintage Marginal Externality	0.733
Naive Tax on County/Vintage Average Externality	0.673
SBO Tax on County/Vintage Marginal Externality	0.718
N	36023471

Proportion of vehicles over the period 1998-2008 whose VMT would be lower than optimal under the indicated tax. We assume that the tax is adjusted each calendar year to reflect changes in the average or marginal externality

The variance and skewness in the distribution of externality per gallon causes a uniform tax to be less efficient than might otherwise be expected. Figure 3 shows this clearly, plotting the kernel density of the externality per gallon in 1998 and 2008, with vertical lines indicating the naive tax and the optimal tax, respectively. The long right tail of the distribution requires that either tax greatly exceed the median externality.

We next examine how the optimal uniform tax would compare to the optimal vehicle-specific tax if the distribution became less skewed. That is, how would a uniform tax perform if the right tail of the distribution—the oldest, dirtiest vehicles—were removed from the road? This could be achieved directly from a Cash for Clunkers-style program, or indirectly through tightening emissions standards in the Smog Check Program. Sandler (2012) shows that vehicle retirement programs are not cost-effective in reducing criteria emissions, and possibly grossly over pay for emissions; however the overall welfare consequences of this sort

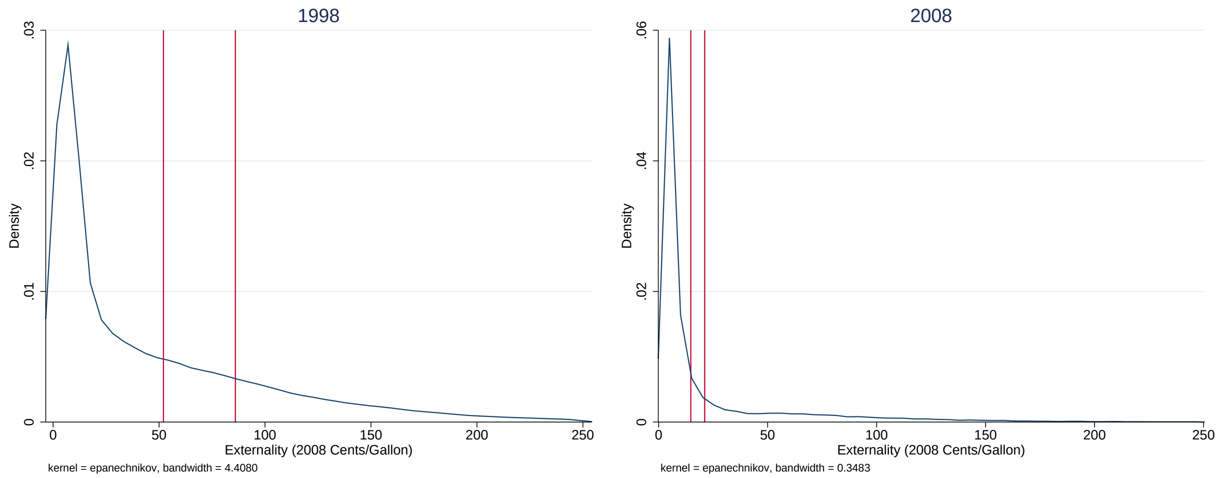

Figure 3: Distribution of externality per gallon—vertical lines indicate naive and marginal uniform tax

of scheme may be more favorable if they improve the efficiency of a uniform gasoline tax. Table 15 shows the ratios of DWL with the SBO gasoline tax to DWL with no tax with increasing proportions of the top of the externality distribution removed. Removing the top 1 percent increases the DWL reduction from 30 percent to 38 percent of the total with no tax. Scrapping more of the top end of the distribution improves the outcome further. If the most polluting 25 percent of vehicles were removed from the road and the SBO gasoline tax was imposed based on the weighted externality of the remaining 75 percent, this would remove 58.3 percent of remaining DWL. Of course, the practical complications of scrapping this large a proportion of the vehicle fleet might make this cost-prohibitive.

Table 15: Ratios of DWL with Tax to DWL With No Tax, Scrapping Most Polluting Vehicles

	Percentile Scrapped					
	None	1%	2%	5%	10%	25%
1998	0.568	0.476	0.451	0.426	0.419	0.439
1999	0.577	0.484	0.467	0.453	0.452	0.469
2000	0.583	0.501	0.486	0.472	0.471	0.478
2001	0.627	0.531	0.515	0.501	0.501	0.456
2002	0.675	0.590	0.578	0.566	0.564	0.509
2003	0.699	0.634	0.625	0.615	0.613	0.488
2004	0.740	0.681	0.672	0.662	0.656	0.458
2005	0.762	0.704	0.693	0.678	0.657	0.390
2006	0.796	0.735	0.725	0.707	0.668	0.363
2007	0.817	0.763	0.753	0.733	0.672	0.378
2008	0.836	0.783	0.772	0.745	0.628	0.373
Average	0.698	0.626	0.612	0.596	0.573	0.436

Notes: DWL with no tax calculated based on the difference in emissions from imposing a tax equal to the actual externality per gallon consumed by a particular car. SBO tax computed as the weighted average of externality per gallon, using the negative slope of the vehicle's demand curve as the weight. All taxes also weight vehicles by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole.

4 Dynamic Model of Vehicle Choice

A major component of this project is the extensive methodological development involved in creating a dynamic model of vehicle choice and driving. This work fulfills TASK B.5, which is to "design, code, and run the joint vehicle choice and usage model." The following lays out the basics of the model and discusses how we coded and implemented it in Matlab. This model follows, and in some respects even extends the cutting edge of dynamic discrete choice modeling applied to vehicles in economics.

Recall that a basic motivation for a dynamic model is that the evolution of the light duty fleet is determined by decisions consumers make about when to buy a vehicle and which vehicle to buy. These decisions are governed by consumer expectations of improvements in vehicle technologies and gasoline prices, as well frictions in the market, such as search costs or other transaction costs. In short, by including the dynamics of the decision process, we can better capture the very real factors that influence consumer decisions of whether to buy and sell, when to buy and sell, and what car to buy and sell. These factors will unquestionably be important for the evolution of the fleet, and are not yet well-understood.

In our framework, the light duty vehicle fleet is comprised of vehicles owned by households. Households decide about whether to keep the vehicles that we currently own or whether to sell or scrap them, and also whether to purchase a different vehicle. The backdrop to these decisions is an ever changing landscape of vehicle technology. On one hand, cars currently existing as part of the fleet are in a process of depreciation, in the sense that we are becoming less valuable to households (and in the market) over time. This change in their value occurs for two primary reasons. First, conditional on all of the other cars in the fleet, using one's vehicle (or even just letting it age without use) causes it to lose productive value. Second, there is a constant stream of new vehicles emerging and available for purchase. Technological change allows these new cars to be of superior quality (or lower price) than older vintages. The interplay between the household preferences, depreciation of the existing car holdings, and the emergence of new vehicle models is what ultimately determines the evolution of the fleet, and is what we will model.

4.1 Model

Consider a household with a single vehicle. In any given time period (e.g., in any given year), consumers have several choices:

- Buy a new car.
- Sell a car.
- Replace a car (i.e., buy and sell).
- Keep their car.

Of course, many households in California have more than one vehicle. However, we would expect most households to have a primary decision-maker for each vehicle. For example, each partner in a couple may have a separate car. There may also be a car for each teenage child. This observation lends itself to a computationally tractable approach where we model each vehicle separately, since each has its own decision-maker. This is an important simplification, for it allows us to model two-car households just like two one-car decision-makers. The other vehicle in the household fleet could influence each decision-maker's utility through a dummy variable following our analysis of the household portfolio; this is an additional feature that could very feasibly be added to the model in the future. With our assumption, we can expeditiously model only one vehicle being sold, replaced, or bought at a time.

This assumption overcomes both computational and conceptual challenges. In fact, virtually every dynamic discrete choice model of car demand we are aware of makes this very assumption (Mannering and Winston, 1985; Esteban and Shum, 2007; Stolyarov, 2002; Adda and Cooper, 2000; Schiraldi, 2011; Cernicchiaro and de Lapparent, 2015). The reason for this assumption is that it dramatically decreases the size of the choice set, and avoids the curse of dimensionality. For example, if a household has three cars and there are 80 cars on the market they could choose from, then there are millions of possible combinations that are possible decisions (e.g., one combination would be replace car 1 with car X and keep the other two, while another combination would be replace car 1 with car X, replace car 2 with car Y, and sell the third without replacement). This would be computationally infeasible to solve for a large dataset in a reasonable amount of time, even with considerable computational power.

One could make further assumptions to attempt to make headway on allowing for there to be some relationship between the portfolio of a household. One way to do it would be to assume that one of the vehicles (e.g., the oldest or the one with the highest odometer reading) is the one to be replaced or sold, while the other vehicles in a household are kept. We explored this path extensively and determined that there is enough heterogeneity in which vehicle is sold that the assumption would be a poor one and would not help the fit of the model substantially. We decided that we would actually be better off treating each vehicle as a separate decision-maker. Another possible approach would be to treat each vehicle as a separate decision-maker, but allow the utility of that decision-maker for any given vehicle type to be determined by the class of the other vehicles in the household. We deem this as a more promising approach. This approach would also align nicely with our evidence above demonstrating a "portfolio effect" guiding the evolution of the vehicle fleet (in fact, the empirical results on the portfolio effect above indicate the degree of the bias that might come about from ignoring any portfolio effects). However, it would increase the computational complexity of the model, albeit much more modestly. While this increase in the computational complexity led us to hold off on this addition for this report. This path of exploration is on the methodological frontier and while we made significant strides forward, it is a challenging endeavor and we see this as a valuable direction for future research to push the frontier beyond what any one has done.

4.1.1 Household utility

Each household will make a choice among those options above based on what provides the greatest expected utility. In each time period, a household i receives utility from the vehicles we have. In our model, cars can be viewed as bundles of attributes. Some examples of attributes are vehicle price, class, vintage (age), fuel economy, horsepower, weight, and brand (make). We refer to the characteristics of car j in time t as X_{jt}. Note that X_{jt} must be indexed by time in order to include vintage and price (which changes due to depreciation) in the set of relevant characteristics. In our modeling, we focus on vehicle age and vehicle class as the key characteristics. Each vehicle class and model year is associated with a different fuel economy , which is then applied in our calculations.

Household i also has characteristics that are relevant for the decision. For example, there is heterogeneity across households in income y_i and driving VMT_i. Of course, households also differ in the current vehicle stock that we hold, which is captured in X_{jt} for each vehicle j. Let MPG_{jt} denote the fuel economy of vehicle j at time t. We then denote the total amount spent on gasoline over the time period as $g_{ijt} = VMT_{it} \cdot p_t^{gas}/MPG_{jt}$, where p_t^{gas} is the price of gasoline at time t. Repair costs are not currently included due to a lack of reliable data on the costs of repairs.[22] Should such data be available, they could be readily included.

A key component of the dynamic model is what is called the "flow utility", or the period-t utility. This is the utility received during time period t by the household for making a certain decision. We model the period-t utility for a household i with car k in time period t as follows:

$$\begin{aligned}
u_{ikkt} &= X_{kt}\theta_i^x - g_{ikt}\theta_i^p + \epsilon_{ikkt} && \text{Keep, including } k = 0 \\
u_{i0jt} &= X_{jt}\theta_i^x - g_{ijt}\theta_i^p + \epsilon_{i0jt} && \text{Buy} \\
u_{ik0t} &= p_{kt}\theta_i^p - \tau_i + \epsilon_{ik0t} && \text{Sell, } k \neq 0 \\
u_{ikjt} &= X_{jt}\theta_i^x - g_{ijt}\theta_i^p - p_{jt}\theta_i^p + p_{kt}\theta_i^p - \tau_i + \epsilon_{ikjt} && \text{Replace, including } k = 0,\ \forall j \in J_t
\end{aligned} \tag{11}$$

The first term, u_{ikkt}, indicates the flow utility when the household begins with vehicle k and chooses to keep vehicle k. It also refers to the situation where the household does not own a vehicle and chooses to continue not owning a vehicle. The equation states that the household receives utility from the attributes of the car (through $X_{kt}\theta_i^x$), loses utility from paying for gasoline (other maintenance costs could be added into this with little difficulty), and has a random idiosyncratic component to utility that is unobserved to the econometrician and depends on the car, time period, and household.

The second term, u_{i0jt}, indicates the flow utility when a household chooses to add a vehicle j to the household portfolio. Households receive utility based on the attributes of the vehicle, lose utility from the cost of the gasoline to run the vehicle, and an unobserved idiosyncratic component.

The third term, u_{ik0t}, indicates the flow utility when the household begins with vehicle k and chooses to sell that vehicle and not buy another one. This is obviously only possible when the decision-maker has a vehicle to start. The household will no longer receive utility from the attributes of the vehicle, but will receive utility from the payoffs from selling the vehicle (p_{kt}). θ_i^p can be thought of as the marginal utility of money and we model this as a quadratic in income $\theta_i^p = \sigma_{1i} + \sigma_{2i}y_i + \sigma_{3i}y_i^2$. The household will also bear a transaction cost t_i, which captures the cost to the household of finding a buyer. Just as before, there is also an idiosyncratic component.

[22] To the extent that repair costs are correlated with fuel costs, e.g., because more driving means more repairs, our fuel cost coefficients will capture both fuel costs and repair costs.

The last term, u_{ikjt}, indicates the flow utility when the household begins with vehicle k and replaces that vehicle with vehicle j. Here there is utility from the attributes of vehicle j, a loss in utility from the gasoline purchased to run vehicle j, the cost of purchasing vehicle j, and the payoff from selling vehicle k. Again, there is a transaction cost and idiosyncratic component.

4.1.2 Value function

A key aspect of a dynamic discrete choice model is that the household makes decisions based on both the current period utility (flow utility) and the expectation of what utility will be in the future based on the decision made today. Another way to think about it is that dynamic discrete choice models incorporate the "option value" (which we will use interchangeably with "continuation value") associated with the choices that made today. If a household delays its purchase this year, it may not own the "best" car today, but it preserves the option to buy a "better" car next year. On the other hand, it can choose to replace its car today, upgrading to the current model that best meets its needs; but it is likely committing to own that car for several years (since subsequent technologies will not provide as much incremental benefit beyond this year's new car than we would have beyond the car that was replaced).

The "value function" captures the current period utility and the future utility, recognizing that the consumer will continue to make decisions to maximize their utility. This expected future utility is often called the "continuation value," so the value function is the sum of the current period utility and the continuation value. For simplicity, we will drop the t subscript here on out and denote the next period's value with a prime, but note that each of the equations going forward applies in each time period. To illustrate, the value function for household i that has chosen to keep vehicle k is given by:

$$\hat{V}_i(k, \epsilon) = u_{ikk} + \beta E[\hat{V}_i(k', \epsilon')|k, \epsilon].$$

This shows that the value function is the sum of the current period utility plus the continuation value, which is the expectation of future utility in the case where the household chooses to keep. Of course, the household can also choose to replace their vehicle with another vehicle or purge the vehicle from their household portfolio without replacing. Thus, for a household that currently has vehicle k, the value function can be written as:

$$
\begin{aligned}
\hat{V}_i(k,\epsilon) = \max \Big\{ & u_{ikk} + \beta E[\hat{V}_i(k',\epsilon')|k,\epsilon], \\
& 1[k \neq 0] \cdot (u_{ik0} + \beta E[\hat{V}_i(0,\epsilon')|0,\epsilon]), \\
& \max_{j \in J} \{ u_{ikj} + \beta E[\hat{V}_i(j',\epsilon')|j,\epsilon] \} \Big\}
\end{aligned}
$$

In the above equation, the first term shows the decision to keep the vehicle k, the second shows the decision to purge the vehicle, and the third shows the decision to replace vehicle k with vehicle j. If we define $k \in \{0, 1, ...\}$, then this equation also subsumes the "buy" option for a household adding a car, since the replace decision can be thought of replacing no car with a car.

For analytical convenience, we can also transform the value function to pull out the cost of purchasing vehicle k:

$$\bar{V}_i(k,\epsilon) = \hat{V}_i(k,\epsilon) - p_k \theta_i^p.$$

Then plugging in from above, we can rewrite the value function as follows:

$$
\begin{aligned}
\bar{V}_i(k,\epsilon) = \max \Big\{ & X_k \theta_i^x - (p_k - \beta E[p_k'])\theta_i^p - g_{ik}\theta_i^p + \epsilon_{ikk} + \beta E[\bar{V}_i(k',\epsilon')|k,\epsilon], \\
& 1[k \neq 0] \cdot (-\tau_i + \epsilon_{ik0} + \beta E[\bar{V}_i(0,\epsilon')|0,\epsilon]), \\
& \max_{j \in J} \{ X_j \theta_i^x - (p_j - \beta E[p_j'])\theta_i^p - g_{ij}\theta_i^p - \tau_i + \epsilon_{ijk} + \beta E[\bar{V}_i(j',\epsilon')|j,\epsilon] \} \Big\}
\end{aligned}
\tag{12}
$$

4.1.3 Logit inclusive value

An increasingly popular approach used in the literature on dynamic discrete choice modeling is to define a single "logit inclusive value" that captures the expected utility of a given choice in a parsimonious manner, which eases the computational burden and yet still includes the key features that influence decision-making. As a preliminary, we define the "net augmented utility" in order to simplify notation:

$$\phi_{ij} = X_j\theta_i^x - (p_j - \beta E[p_j'])\theta_i^p - g_{ij}\theta_i^p$$

The net augmented utility is the net utility in a given time period for household i with vehicle j that includes the expected depreciation of the vehicle $(p_j - \beta E[p_j'])$. Once this is defined, we can move on to the logit inclusive value:

$$\delta_i = \ln\left\{\sum_{j\in J} \exp(\phi_{ij} - \tau_i + \beta E[\bar{V}_i(j', \epsilon')|j, \epsilon]\right\}$$

By plugging ϕ and δ into equation (12) and integrating over the Logit error ϵ we then can rewrite the value function in a much more tractable form:

$$V_i(\phi_{ik}, \delta_i) = \int \bar{V}_i(k, \epsilon)\mathrm{d}\epsilon \tag{13}$$

An implication of this formulation of the value function and the assumption of Type I extreme value logit errors is that we have an analytical formulation for the probability of each decision:

$$\begin{aligned}
P(d_i = \text{Keep }|\phi_{ik}, \delta_i) &= \frac{\exp(\phi_{ik} + \beta E[V_i(\phi_{ik}', \delta_i')|\phi_{ik}, \delta_i])}{\exp(V_i(\phi_{ik}, \delta_i))} \\
P(d_i = \text{Sell }|\phi_{ik}, \delta_i, k \neq 0) &= \frac{\exp(-\tau_i + \beta E[V_i(\phi_{i0}', \delta_i')|\phi_{i0}, \delta_i])}{V_i(\phi_{ik}, \delta_i)} \\
P(d_i = \text{Replace }|\phi_{ik}, \delta_i) &= \frac{\exp(\phi_{ij} - \tau_i + \beta E[V_i(\phi_{ij}', \delta_i')|\phi_{ij}, \delta_i])}{V_i(\phi_{ik}, \delta_i)}
\end{aligned}$$

Recall that the replace decision can subsume the "Buy" decision, so all possible decisions are accounted for here.

4.1.4 Transitions

This dynamic model has two key state variables that evolve over time: the logit inclusive value and the net augmented utility. In fact, one of the primary advantages of using the logit inclusive value approach is to reduce the state space. The transitions for these two variables are given by:

$$\delta_i' = \gamma_{1i} + \gamma_{2i}\delta_i + \nu_i$$
$$\phi_{ij}' = \rho_{1i} + \rho_{2i}\phi_{ij} + \rho_{3i}\delta_i + \mu_{ij}.$$

Here ν_i and μ_{ij} are error terms.

4.1.5 Summary of key equations

To summarize, these are the following key equations of the dynamic discrete choice model:

$$
\begin{aligned}
\phi_{ijt} &= X_{jt}\theta_i^x - (p_{jt} - \beta E[p_{jt+1}])\theta_i^p - g_{ijt}\theta_i^p \\
g_{ijt} &= \frac{VMT_{it} \cdot p_t^{gas}}{MPG_{jt}} \\
\theta_i^p &= \sigma_{1i} + \sigma_{2i}y_i + \sigma_{3i}y_i^2 \\
\delta_{it} &= \ln\Big\{\sum_{j\in J_t}\Big(\exp(\phi_{ijt} - \tau_i + \beta E[V_i(\phi_{ijt+1}, \delta_{it+1})|\phi_{ijt}, \delta_{it}]\Big)\Big\} \\
\delta_{it+1} &= \gamma_{1i} + \gamma_{2i}\delta_{it} + \nu_{it} \\
\phi_{ijt+1} &= \rho_{1i} + \rho_{2i}\phi_{ijt} + \rho_{3i}\delta_{it} + \mu_{ijt} \\
V_i(\phi_{ikt}, \delta_{it}) &= \ln\Big\{\Big(\exp(\phi_{ikt} + \beta E[V_i(\phi_{ikt+1}, \delta_{it+1})|\phi_{ik}, \delta_{it}]) + \\
&\quad 1[k \neq 0] \cdot \exp(-\tau_i + \beta E[V_i(\phi_{i0t+1}, \delta_{it+1})|\phi_{i0t}, \delta_{it}]) + \\
&\quad \exp(\delta_{it})\Big)\Big\} \\
P(d_i = \text{Keep } |\phi_{ikt}, \delta_{it}) &= \frac{\exp(\phi_{ikt} + \beta E[V_i(\phi_{ikt+1}, \delta_{it+1})|\phi_{ik}, \delta_{it}])}{\exp(V_i(\phi_{ikt}, \delta_{it}))} \\
P(d_i = \text{Sell } |\phi_{ikt}, \delta_{it}, k \neq 0) &= \frac{\exp(-\tau_i + \beta E[V_i(\phi_{i0t+1}, \delta_{it+1})|\phi_{i0t}, \delta_{it}])}{V_i(\phi_{ikt}, \delta_{it})} \\
P(d_i = \text{Replace } |\phi_{ikt}, \delta_{it}) &= \frac{\exp(\phi_{ijt} - \tau_i + \beta E[V_i(\phi_{ijt+1}, \delta_{it+1})|\phi_{ijt}, \delta_{it}])}{V_i(\phi_{ikt}, \delta_{it})}
\end{aligned}
$$

4.1.6 Solving the model

Since this model has probabilities for the different decisions households can make and we observe the actual decisions, we can solve the model using maximum likelihood with a nested fixed point algorithm, following the standard in the literature (e.g., see Rust (1985)). We have coded the model in Matlab and are using the standard Matlab solvers (e.g., fminunc) to solve the model.

To take the model to the data, we must convert the data to a suitable form. We make a series of further refinements to the data for this purpose. To further clarify, we model the following transitions for different households:

- Decision-makers that started with No Car
  - $0 \to 0$
  - $0 \to 1$
- Decision-makers that started with 1 Car
  - $1 \to 0$
  - $1 \to 1'$
  - $1 \to 1$

This parsimonious summary captures all possibilities. For example, consider a two car household that wants to buy a third car. The third car must go to a new decision-maker (e.g., a teenager) and thus there will be two decision-makers that started with one car and decided to keep, and a third decision-maker that started with zero cars and added a car. It thus follows that we determine households that made the $0 \to 0$ decision based on whether the household later or earlier in the data sample purchases a car. We consider households that never had a car as outside of the relevant sample. Some households that later entered the sample or earlier entered the sample may have moved in or out of California. We have acquired data from R.L. Polk that allow us to determine who these households are and can account for them.

In the data sample, some households have the same vehicle portfolio for most years, but seem to have gap years when we are missing from the sample. In these (relatively unusual) cases, we deem this a data issue and impute the missing observations based on those from the adjacent years.

4.2 Results

Recall that the purpose of this model is to improve our understanding of the evolution of the vehicle fleet and how transaction frictions can affect this evolution. The results presented here should be considered illustrative. They are estimated on the Department of Motor Vehicles and Bureau of Automotive Repair dataset covering the years 2001-2007 and modified for the purposes of this analysis. Specifically, to ease the computational burden, we took a 10% subsample randomly drawn from all years and from throughout California. We also arranged the data so that they are in the form of a repeated cross-sectional dataset, with each observation being a vehicle that is treated as a separate decision-maker. We can see the basic features of the data we are using in the following table of transitions that occur in each year (Table 16).

Table 16: Transitions that occur over a year

Transition	Count	Percent
$0 \rightarrow 1$ (Buy)	565,443	10.04
$1 \rightarrow 0$ (Sell)	173,115	3.07
$1 \rightarrow 1$ (Keep)	3,502,168	62.21
$1 \rightarrow 1'$ (Replace)	319,367	5.67
$0 \rightarrow 0$ (Stay out)	1,069,857	19.00
Total	5,629,950	100.00

This table shows that most transitions are cars that are kept (i.e., retained to the next period). But there are a large number of households that stay out of the market or buy in any given year. We determine the households who "stay out" based on whether household has a vehicle at any point in time over 2002-2008.

To demonstrate that we have successfully completed the coding and debugging of the dynamic model, there are hundreds of possible results that could be presented. For brevity, we only present a few of the results. We will begin simply by presenting the coefficients from estimating the model (standard errors can be calculated either from analytical/numerical derivatives or using a bootstrap approach). In order to characterize heterogeneity in both income and driving behavior, we categorized the dataset into four bins based on Census block group household income and VMT and then run the model separately on each bin. The first bin has below-median income and below-median VMT, the second has below-median income and above-median VMT, and so on. The average income for the below-median income is $67,123 and for the above-average income is $182,930. The average VMT for the below-median VMT is 25 miles per day, while the average VMT for the above-median VMT is 75 miles per day.

To illustrate, we present the coefficients from the first bin (below-median income and below-median VMT) and the fourth bin (above-median income and above-median VMT). These results are shown in Table 17.

Table 17: Model Estimation Coefficient Results

Variable	Bin 1	Bin 4
Variable	low-low	high-high
Vehicle age	-0.09	-0.12
Car dummy	3.40	4.03
Luxury car dummy	4.90	5.50
SUV dummy	4.67	5.29
Luxury SUV dummy	4.64	4.69
Truck dummy	4.65	4.91
Van dummy	4.75	5.35
Vehicle cost	-0.13	-0.05
Gasoline price	-3.22	-1.30
Transaction cost	-7.41	-7.87
Note: XUVs are the omitted class dummy.		

One can interpret these results as indicating how each of these variables influences utility, so the magnitudes are not important here, but are very important for model fit and counterfactuals. However, the relative magnitudes of the vehicle class dummies and the signs of all of the variables can be examined and these results are quite sensible. Consumers get less utility from older vehicles. Consumers get more utility from SUVs and trucks than cars, which is consistent with the much higher purchase prices of these vehicles. Consumers also get more utility from luxury cars and SUVs than the non-luxury versions. Higher prices for vehicles reduce utility, as do higher gasoline prices. The transaction cost is negative (implying a cost) and is larger for wealthier, high-VMT consumers than for poorer, low-VMT consumers. The vehicle cost and gasoline price are also less important for wealthier, high-VMT consumers, as one might expect. These exact coefficients are less important than the conditional choice probabilities (CCPs), which tell us what the probability of a transition is given the current state. These CCPs will be examined below.

But first, to demonstrate the importance of the transaction cost–a key feature in our model–we also estimate a version of the model where we fix the transaction costs to be zero. Table 18 shows results from the model runs without transaction costs, which can be compared to the results in the previous table.

Table 18: Model Estimation Coefficient Results

No Trans Costs		
Variable Bin 1	Bin 4	
Variable	low-low	high-high
Vehicle age	0.033	0.019
Car dummy	-6.473	-6.776
SUV dummy	-7.692	-7.426
Truck dummy	-7.890	-8.367
Van dummy	-8.291	-8.093
Vehicle cost	-0.437	-0.862
Gasoline price	-2.350	-0.621
Transaction cost	–	–
Observations	3.2m	

Comparing the results without transaction costs in the last two columns of Table 18 to those with transaction costs demonstrations just how important including transaction costs is to estimating sensible coefficients. Without transaction costs we see that many of the signs flip and relative magnitudes make much less sense. Thus, it is not even worth interpreting the results without transaction costs. These runs without transaction costs were performed without the luxury vehicle dummies, but results with the luxury vehicle dummies would be similarly problematic.

A common way to explore the model fit in dynamic discrete choice models is to compare the modeled CCPs to the observed choices being made. For example, if we see 60% of consumers keeping their vehicle, how well does the model predict this?

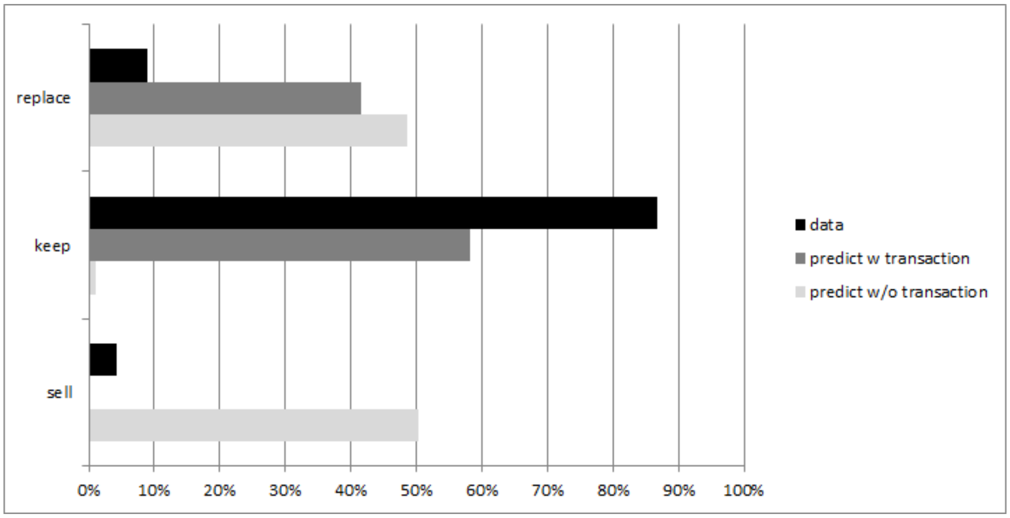

Figure 4: Conditional Choice Probabilities for Transitions for Bin 1

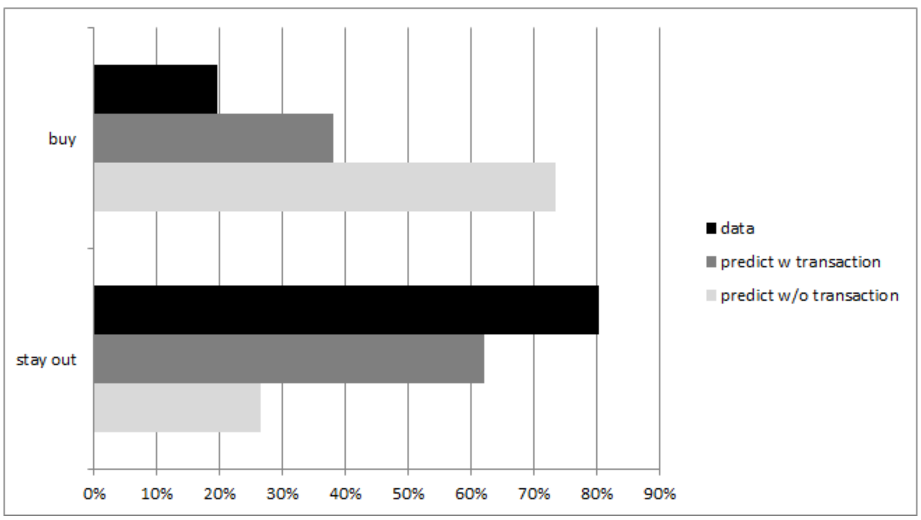

Figure 5: Conditional Choice Probabilities for Transitions for Bin 1

We next present a series of figures that illustrate the model fit over the key transitions in the data (note that it is called a "transition" because we are referring to shifting or transitioning from one state–such as the state of having vehicle X–to another state–such as having another vehicle Y). In each of these figures, we present what we see in the actual data, our model fit and the model fit without transaction costs. The first figure, Figure 4, shows the CCPs for the replace, keep, and sell option for Bin 1 (below-median income and below-median VMT). The second, Figure 5, shows the CCPs for the buy and stay out of the market ($0 \rightarrow 0$) options for Bin 1. The third, Figure 6, shows the CCPs for the replace, keep, and sell

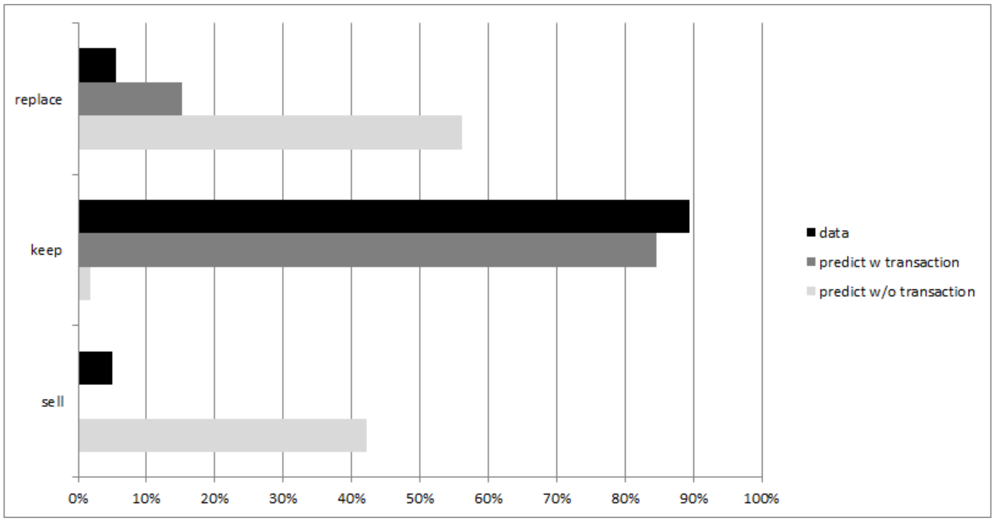

Figure 6: Conditional Choice Probabilities for Transitions for Bin 4

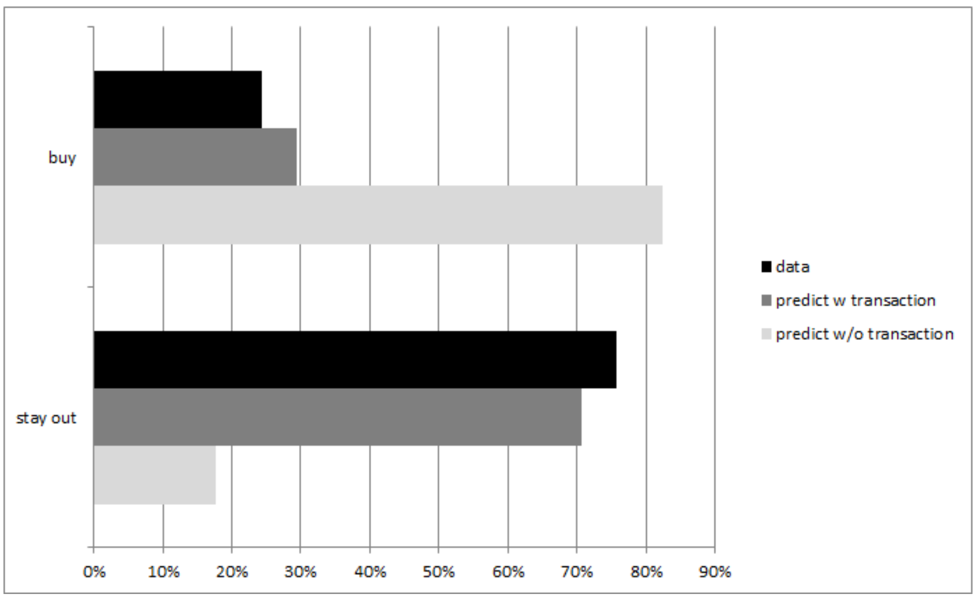

Figure 7: Conditional Choice Probabilities for Transitions for Bin 4

option for Bin 4 (above-median income and above-median VMT). The last, Figure 7, shows the CCPs for the buy and stay out ($0 \rightarrow 0$) options for Bin 4. The figures for the other two bins are similar to these.

The results in these figures are encouraging and demonstrate a working model that does a reasonable job at matching the CCPs for some of the most important decisions. Where it performs worst is for selling decisions–it slightly under-estimates selling and over-estimates buying. Yet it largely captures the primary features of the data. Even more notably, the model fit is quite a bit better when transaction costs are included than when they are not included. Consider the buy decision for consumers in Bin 1. The prediction with transaction costs is not too far from the actual data, while the prediction without transaction costs is far off. The same is true for stay out and is very true for sell. The reason for this is intuitive: transaction costs provide a friction in the market that captures real search and other costs that prevent consumers from simply turning over their car every year in order to get a car that is a better match for their preferences.

These figures from our simulation strongly demonstrate the importance of including transaction costs and also demonstrate that the model is working and that there is a reasonable model fit. Of course, the model fit can still be further improved. One promising next step is to include a transaction cost whenever a consumer is selling a vehicle. This would help the model better fit the sell and replace decisions, which are two of the decisions that are the furthest off from the data (the buy and stay out decisions very closely match the data). Further refinements of the class variable may also help improve the model fit. However, the bottom line is that we have a working model, estimating sensible coefficients, that can be further refined in order to run policy-relevant policy counterfactuals.

4.3 Policy Counterfactual

For the purposes of this final report, we run one illustrative policy counterfactual. The counterfactual that we run increases the price of low fuel economy vehicles relative to high fuel economy vehicles. This has similarities to a feebate policy (i.e., a fee on low fuel economy vehicles and rebate on high fuel economy vehicles). It also mimics the effect of a fuel economy standard for under a fuel economy standard there will be a shadow price on low fuel economy vehicles that is much greater than on higher fuel economy vehicles. In both cases, the relative prices of the vehicles are adjusted by the policy. See Gillingham (2013) for further discussion of the relationship between the two policy instruments.

The steps for running the counterfactual are as follows. First we run the model, estimating the coefficients based on the data in the years 2001-2006. Then we use those coefficients and change the observed vehicle prices to a new set of vehicle prices that are adjusted by the policy. Specifically, the exact policy counterfactual we run increases the price of low fuel economy vehicles by $1,000 (for this counterfactual, high fuel economy vehicles are not changed in absolute price, but become relatively more attractive). Then we calculate the

CCPs going forward.[23] With the CCPs, we can then examine the decisions that consumers make, including decisions that govern the evolution of the light duty vehicle fleet and emissions. We predict forward to 2007 and run the counterfactual that year.

After these steps, we have a set of results to be explored. The CCPs themselves are one set of results, and these are seen below:

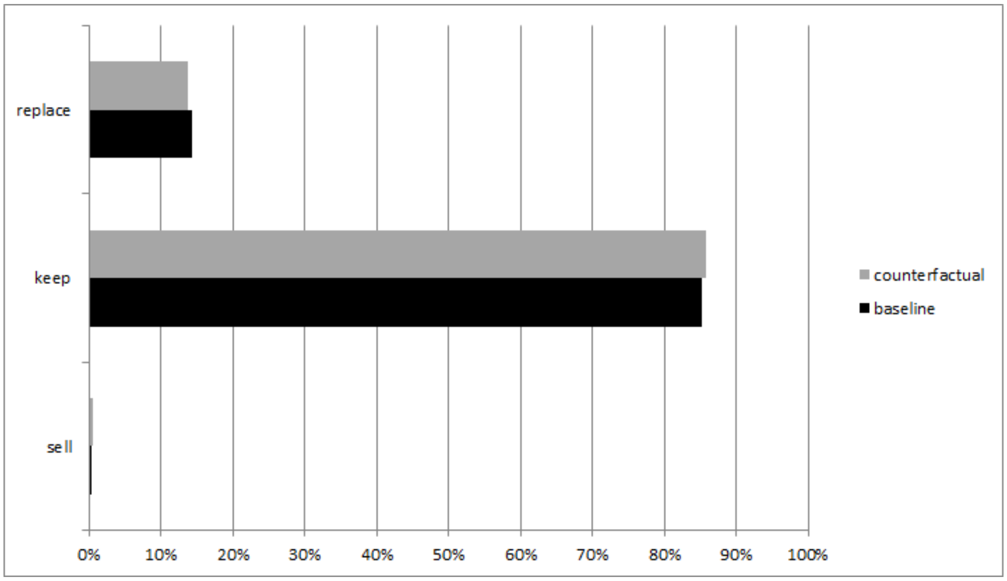

Figure 8: Conditional Choice Probabilities for Counterfactual

These CCPs demonstrate that the counterfactual works in the direction expected: after an increase in new vehicle prices, there is a decrease in households choosing to replace a vehicle and an increase in vehicles being kept. Similarly, there is an increase in households choosing to keep their vehicle. When the price of new vehicles increases, this is exactly the effect one would anticipate.

Another set of important results is the effect of the policy on purchases of new low fuel economy vehicles. The results indicate that these purchases decline by 13%, as evidenced by a 13% lower transition probability to these vehicles in each year. With an average vehicle price for low fuel economy vehicles of $26,183 and an increase in price of $1,000, this implies an elasticity of 3.4, which is within the range in the literature for substitution elasticities for different classes of vehicles (Berry et al., 1995). One could then make calculations of the effect on emissions if this was desired.

[23]For these results we do not re-compute the nested fixed point in order to save computational time, so these results may not be precisely correct, but they should be very close.

5 Summary and Conclusions

Each of the three main contributions of "Empirical Estimation of Household Vehicle Purchase and Usage Decisions" (Part B of this project) significantly advances our understanding of household vehicle choice. They each draw upon multiple datasets and deploy novel methodological techniques that place this work on the frontier of academic and policy research. The first two contributions are each complete manuscripts. The third is a complete model of the household vehicle choice that draws upon methods that are at the frontier of economic modeling.

The paper titled "The Household Vehicle Portfolio: Implications for Emissions Abatement Policies" (co-authored by James Archsmith, Ken Gillingham, Chris Knittel, and David Rapson) implements a thought experiment that examines the two-car household replacement decision. The instrumental variables approach closely mimics a setting in which one car is randomly kept, and its fuel economy randomly increased by some small amount. This allows for a rigorous estimate of the effect of kept-car fuel economy on the subsequent purchase of the replacement vehicle. This thought experiment is tested on a novel dataset that tracks household vehicle ownership in California 2001-2007. The dataset itself is a significant contribution to current and future research, and the results are both economically and statistically meaningful.

Several key findings emerge. Households exhibit an overall preference for diversity in two-car vehicle portfolios, both when acquiring a new car or a used car. This channel operates in the opposite direction to improvements in fuel economy of cars that are kept by households during replacement events. The portfolio effect can erode the final fuel savings from a one-time increase in fuel economy; our simulations indicate that a significant portion of the fuel savings from increasing fuel economy can be lost by a resulting decrease in a lower fuel economy of the newly purchased car due to the portfolio effect. The portfolio effect would also have strong effects on the used car market. For example, if there is higher demand for lower fuel economy cars after a one-time change in new vehicle fuel economy, the portfolio effect would increase demand for used gas guzzlers, slowing the retirement of these vehicles.

We also explore how this portfolio effect interacts with gasoline prices. As gasoline prices increase, the probability of buying a highly efficient car (lowest quartile of fuel consumption) increases, as expected. But we also find that the effect of the fuel consumption of the kept vehicle is also stronger, so if a household currently has a low-efficiency vehicle, it is more likely to buy a high-efficiency vehicle when gasoline prices increase. In contrast, when gasoline prices increase, the probability of buying one of the least-efficient cars decreases and so does the portfolio effect. So if a household currently has a high-efficiency vehicle, it is less likely to buy a low-efficiency vehicle when gasoline prices increase. More generally, we also find that owners of fuel-efficient cars are less likely to buy cars in the middle quartiles of fuel economy as gasoline prices rise, whereas owners of low fuel economy cars are more likely to buy a middle-quartile car.

The paper titled "The Welfare Impact of Second Best Uniform-Pigouvian Taxation: Evidence from Transportation" (co-authored by Chris Knittel and Ryan Sandler) examines local criteria pollutant reductions that accompany gasoline taxes. The analysis shows that vehicle-level emissions are correlated with vehicle-specific VMT elasticities. An important implication follows: emissions from dirtier vehicles are more responsive to gasoline prices than emissions from fuel efficient vehicles. Policymakers should therefore consider the presence of heterogeneity when determining optimal tax levels. The second-best optimal gas tax is approximately 50 percent larger than the naive gas tax that neglects this heterogeneity. However, the authors also find that the second-best optimal tax falls significantly short of the benefits that would be achieved by the first best (a direct tax on the pollutants themselves). The SBO eliminates only 30 percent of the welfare loss associated with tailpipe emissions.

This work also shows that a hypothetical gasoline tax could vary based on vehicle vintage and county to better account for the externalities. This is because emissions rates are highly correlated with vintage and differ by location. Further it shows that to increase the fairness of gasoline taxes to tackle criteria air pollution, one approach might be to remove highly-polluting vehicles from the road either through a vehicle scrappage program or by tightening emissions standards in the Smog Check program. For example, the findings indicate that if the most polluting 25 percent of vehicles were removed from the road and the standard uniform gasoline tax was imposed on the remaining, this would reduce net welfare losses by 58.3 percent. This is particularly relevant because the data clearly show that poorer households have more-polluting cars and pollute more in total, despite driving fewer miles than richer households. More broadly, older and polluting vehicles were seen to be driven less. Such vehicles may be more likely to be driven by poorer households, but they are found in households throughout all income brackets. One key conclusion emerging from this work is that standard uniform gasoline taxes are not the most efficient policy tool to reduce vehicle emissions in a place like California where the most polluting vehicles that are driven the most disproportionately contribute to emissions. Another is that gasoline taxes used to control emissions are actually more regressive than a tax that directly addresses the emissions.

Finally, a major contribution of this work comes in the form of a new tool for understanding vehicle choice and the evolution of the vehicle fleet in California. We describe and estimate a dynamic discrete model of vehicle choice. The model more closely captures the actual decision-making process of consumers who are forward-looking and recognize that their choices today impact future choices. The novel feature that operationalizes the forward-looking nature of the consumer choice is the presence of transaction frictions. Economists widely believe that search and matching frictions, along with other behavioral considerations, will make it costly for buyers to find their desired product and for sellers to find a good match for cars that they have in stock. Discrete choice models of vehicle choice typically do not include this feature. These frictions are important in the California vehicle market since they alone can rationalize the strong persistence of vehicle ownership. The model that we have developed may be used to forecast changes in the vehicle fleet in response to a variety of common policies, and we present one illustrative counterfactual along these lines.

The research findings and modeling tool developed will provide insight into how consumers have responded and may respond to current and future California Air Resources Board vehicular policies, such as new emission standards and incentives. Additionally, this tool can be used to improve statewide vehicular and emission inventories used to support of policy development.

6 Recommendations

There are several natural next steps on this research agenda. One opportunity is to refine the dynamic model and use it to simulate additional counterfactual policy scenarios. There are two extensions of the household portfolio preferences work that may be seen as natural follow-on work. Given the strength of the household preference for diversity, a natural question is how households relate to electric vehicles and fuel cell vehicles as part of their portfolio. How relevant are range considerations when household select their complementary vehicles? To what extent do households shift VMT across cars in their portfolio? In order to answer questions relating to VMT, higher-frequency data on VMT must become available. It would be beneficial to researchers if there were a standard way to collect VMT annually, or even electronically at higher frequency.

7 References

References

Adda, J. and Cooper, R. (2000). Balladurette and juppette: A discrete analysis of scrapping subsidies. *Journal of Political Economy*, 108(4):778–806.

Bento, A., Goulder, L., Jacobsen, M., and von Haefen, R. (2009). Distributional and efficiency impacts of increased us gasoline taxes. *American Economic Review*, 99(3):667–699.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Busse, M., Knittel, C., and Zettelmeyer, F. (2013). Are consumers myopic? evidence from new and used car purchases. *American Economic Review*, 103(1):220–56.

Cernicchiaro, G. and de Lapparent, M. (2015). A dynamic discrete/continuous choice model for forward-looking agents owning one or more vehicles. *Computational Economics*, 46(1):15–34.

Davis, L. and Kahn, M. (2011). International trade in used vehicles: The environmental consequences of nafta. *American Economic Journal: Economic Policy*, 2(4):58–82.

Diamond, P. A. (1973). Consumption externalities and imperfect corrective pricing. *The Bell Journal of Economics*, 4(2):526–538.

Esteban, S. and Shum, M. (2007). Durable-goods oligopoly with secondary markets: The case of automobiles. *RAND Journal of Economics*, 38(2):332–354.

Gillingham, K. (2011). The consumer response to gasoline price changes: Empirical evidence and policy implications. *Stanford University*, Ph.D. Dissertation.

Gillingham, K. (2013). The economics of fuel economy standards versus feebates. *Yale University Working Paper*, New Haven, CT.

Hughes, J., Knittel, C., and Sperling, D. (2008). Evidence in a shift in the short-run price elasticity of gasoline demand. *Energy Journal*, 29(1):113–134.

Klier, T. and Linn, J. (2010). The price of gasoline and the demand for fuel efficiency: Evidence from monthly new vehicle sales data. *American Economic Journal: Economic Policy*, 2(3):134–53.

Li, S., Linn, J., and Muehlegger, E. (2014). Gasoline taxes and consumer behavior. *American Economic Journal: Economic Policy*, 6(2):302–342.

Mannering, F. and Winston, C. (1985). A dynamic empirical analysis of household vehicle ownership and utilization. *RAND Journal of Economics*, 16(2):215–236.

Matthews, H. S. and Lave, L. (2000). Applications of environmental valuation for determining externality costs. *Environmental Science & Technology*, 34:1390–1395.

Morrow, S. and Runkle, K. (2005). April 2004 evaluation of california's enhanced vehicle inspection and maintenance (smog check) program. *California Air Resources Board Report to the Legislature,*, Sacramento, California.

Muller, N. Z. and Mendelsohn, R. (2009). Efficient pollution regulation: Getting the prices right. *American Economic Review*, 99(5):1714–1739.

Poterba, J. M. (1991). Is the gasoline tax regressive? *Tax Policy and the Economy*, 5:145–164.

Rust, J. (1985). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033.

Sandler, R. (2012). Clunkers or junkers? adverse selection in a vehicle retirement program. *American Economics Journal: Economic Policy*, 4:253–281.

Schiraldi, P. (2011). Automobile replacement: A dynamic structural approach. *RAND Journal of Economics*, 42(2):266–291.

Stolyarov, D. (2002). Turnover of used durables in a stationary equilibrium: Are older goods traded more? *Journal of Political Economy*, 110(6):1390–1413.

Wakamori, N. (2016). Portfolio considerations in differentiated product purchases: An application to the japanese automobile market. *Working Paper*.

West, S. (2004). Distributional effects of alternative vehicle pollution control technologies. *Journal of Public Economics*, 88:735.

8 Appendices

The Household Vehicle Portfolio: Implications for Emissions Abatement Policies

James Archsmith, UC Davis

Kenneth Gillingham, Yale

Christopher Knittel, MIT

David Rapson*, UC Davis

November 16, 2016

PRELIMINARY AND INCOMPLETE.
Please do not cite or circulate.

Abstract

This paper quantifies the extent to which multi-car households exhibit preferences for a diversified vehicle portfolio. We deploy a novel identification strategy to examine how an exogenous change in the fuel economy of a kept vehicle affects a household's choice of a second vehicle purchased. This has potential implications for the relative long-run effectiveness of greenhouse gas abatement initiatives such as fuel economy standards and a price on carbon. Results indicate a strong preference for a diverse portfolio in fuel-economy. They highlight the importance of our instrumental variables approach and the pitfalls that arise when using between-household, rather than within-household, variation. Our results imply that the portfolio effect will exert a strong force that may erode a substantial portion of expected gasoline savings from higher fuel economy standards.

*Archsmith: PhD Candidate in Economics, UC Davis, jesmith@ucdavis.edu. Gillingham: Assistant Professor, Yale University and NBER, kenneth.gillingham@yale.edu. Knittel: William Barton Rogers Professor of Energy Economics, Sloan School of Management, and Director of the Center for Energy and Environmental Policy Research, MIT, and NBER, knittel@mit.edu. Rapson: Associate Professor of Economics, UC Davis, dsrapson@ucdavis.edu.

This paper has benefited from conversations with seminar participants at the Energy Institute @ Haas Summer Energy Camp. We thank James Archsmith for excellent research assistance. We gratefully acknowledge financial support from the California Air Resources Board. The statements and conclusions presented are those of the authors and not necessarily those of the California Air Resources Board. The mention of commercial products, their source, or their use in connection with material reported herein is not to be construed as actual or implied endorsement of such products.

1 Introduction

Understanding the demand patterns for vehicles is important for a number of policy issues, most notably fuel economy standards or gasoline taxes on vehicle choice. Empirical models used to analyze the costs and benefit of such policies have typically assumed away the multi-vehicle nature of many households. That is, these models assume that consumers choose only only vehicle; or, alternatively, that the choice of each vehicle in a household is independent of the others.

There are likely to be two sources of interdependence. The first is that households vary in their "type." This source of dependence is *implicitly* captured in empirical models that allow for variation in the willingness to pay for vehicle attributes. For example, the choice of fuel economy across vehicles within a household will be positively correlated for a household that particularly values horsepower. An alternative source of interdependence exists if households have a taste for diversification. In this case, a household endowed with a high horsepower vehicle will favor a more fuel-efficient second vehicle.

The presence of this second form of interdependence can alter the predictions from policy counterfactuals. For example, suppose a policy were to increase the chosen fuel economy of vehicles for a given household at time t. Counterfactual predictions based on models that assume independence would generate a naive measure the fuel economy benefits. If households prefer diversification, the actual fuel economy benefit from the policy across the entire household portfolio would be smaller, as the fuel economy of the next vehicle purchased would fall.

In this paper, we use panel data on the portfolio of vehicles within a household to estimate how a household's choice of vehicle depends on the other vehicles owned by the household. In particular, our data track the portfolio of vehicles, for a particular household, over time. Our empirical strategy estimates how the fuel economy of a newly-added vehicle depends on the fuel economy of the vehicle already held by the household. Our identification strategy relies both on the richness of the panel, which allows us to control for household-level fixed effects, as well as a novel instrumental variables approach to control for the endogeneity of the fuel economy of the existing vehicle.

The ideal experiment for our research question would randomly assign the "kept" vehicle to households in the market for a new or used vehicle and then observe the relationship between the fuel economy of this kept vehicle and the fuel economy of the newly-acquired vehicle. Since this ideal experiment is obviously not possible, our identification strategy must overcome two sources of endogeneity stemming from the non-random assignment of the kept vehicle. The first is the choice of which vehicle to replace. Since the household preference for particular features of a multi-car portfolio will directly inform the decision of which car to keep or drop, there is an identification challenge in estimating the portfolio itself using observational data. The second is related to the presence of unobserved household preferences. Household fixed effects can address time invariant unobserved preferences, but there would still be a concern if these preferences change over time.

We expect that these time-varying preferences would generate a positive correlation between the fuel economy of the kept and newly-acquired vehicles.

We employ two sets of instruments to account for these potential sources of bias. The first set of instruments are derived from the observation that changes in the relative price of cars in a portfolio systematically affect the probability that the lowest fuel economy car is dropped. We discuss and present three instruments that rely on this feature of the choice setting, with our preferred instrument being deviations in expectation of the change in relative vehicle prices at the time when the kept car was initially purchased. The second instrument is the gasoline price at the time of the purchase of the kept vehicle. A number of papers (Klier and Linn (2010), Busse et al. (2013), Gillingham (2011)) have shown that vehicle purchase behavior is influenced by contemporaneous gasoline prices. Given the results of this literature, we would expect that the fuel economy of the kept vehicle is influenced by the gasoline price at the time of that purchase—something we confirm in our own data. We argue that the instrument satisfies the exclusion restriction because after controlling for current gasoline prices, we would not expect past gasoline prices to influence the choice of the new vehicle. This, we grant, rests on the assumption that consumers are not using this past gasoline price in the formation of the expected gasoline prices and we have adequately controlled for serial correlation in the residuals.

We find evidence that households value diversification. The greater is the fuel economy of the kept car, the lower the fuel economy of the purchased car. We show this using both a continuous measure of fuel economy, as well as by estimating the probability a household purchases a vehicle in the upper and lower quartiles of the fuel economy distribution. Increases in the fuel economy of the kept car reduces the probability the household purchases a car in the lower quartile of gallons per mile, while such increases reduce the probability the household buys a car in the upper quartile. Changes in gasoline prices affect the preferences for diversification in intuitive ways. As gasoline prices increase, the effect of the fuel consumption of kept vehicle and the probability of buying a car in the lower quartile of fuel consumption becomes even more positive. In contrast, as gasoline prices increase, the effect of the fuel consumption of kept vehicle and the probability of buying a car in the upper quartile of fuel consumption becomes even more negative.

To gauge the importance of the portfolio effect, we use our results to estimate the net effect of an exogenous increase in the fuel economy of the kept vehicle. We calculate the decrease in the fuel economy of the newly purchased vehicle when we increase the fuel economy of the kept vehicle by 10, 25, and 50 percent, across gasoline prices of \$2.00, \$3.00, and \$4.00. These calculations suggest that the portfolio effect can have large consequences of the net effect a one-time increase in fuel economy; between 75 to 95 percent of the fuel savings from increasing the fuel economy of the kept vehicle are eroded from the resulting decrease in fuel economy of the newly purchased vehicle.

The remainder of the study proceeds as follows. The next section describes the household vehicle choice problem and outlining a simple theoretical model (Section 2). We then describe our datasets, the restrictions that determine the sample used for our empirical tests, our identification strategy and empirical approach (Section 3). We then present our results and their economic importance (Section 4). We conclude with a brief discussion of the implications for policymakers and empiricists (Section 5).

2 Context and Model

We begin by developing a simple economic framework to fix ideas and motivate our empirical work. Consider a household who has a vehicle and is considering adding a second vehicle. They may have just sold their second car, or they may be adding a new car to the household's vehicle portfolio. For simplicity, ignore the outside option of not purchasing a second vehicle. Consider a standard discrete choice framework with a random utility model. The household is the decision-maker in this framework; we abstract from issues of within-household bargaining.

Let the characteristics of a vehicle be given by the vector θ_V, where $V \in \{A, B, ...\}$ is the vehicle type. Vehicle types may be defined broadly, such as the class of vehicle (e.g., SUV or small car), or at a finer level, such as at the make-model level. Suppose the household has a vehicle of type A to start.

The household receives utility based on the characteristics of each type of vehicle, and may also receive utility from having a diversity of vehicles, which allows them to optimize their use of the vehicles (e.g., use the larger one for hauling goods and the more efficient one for the longer commutes). Let the contribution to utility from the diversity of the portfolio be given by Γ_{V_1,V_2}, where V_1 and V_2 are the vehicles types for the first and second car respectively.

The indirect utility for household i starting with a vehicle of type A and purchasing a vehicle of type B is thus given as:

$$u_i^{AB} = f(\theta_A) + f(\theta_B) + \Gamma_{AB} - \alpha(p_{A1} + p_{B2}) + \epsilon_i,$$

where $f(.)$ is a function that converts characteristics into consumer utility, and p_{Vj} is the remaining "present value lifetime ownership cost" for a vehicle of type V and order in the household j, where $j = 1$ refers to the vehicle the household already holds and $j = 2$ refers to the new vehicle. To understand the present value lifetime ownership cost, note that for p_{B2} the ownership cost includes the purchase price in addition to the future costs of fuel and maintenance, while for p_{A1} the ownership cost is just the future fuel and maintenance costs. These can be thought of as expectations based on the future expected driving of each of the vehicles, which need not necessarily be the same across vehicle types. α is the marginal utility of money.

In making the choice of which vehicle type to buy for the second vehicle, the household will compare u_i^{AB} to the utility from all other options. For example, suppose the household with vehicle of type A already is considering the option of buying another vehicle of type A. The utility u_i^{AA} would then be given by:

$$u_i^{AA} = f(\theta_A) + f(\theta_A) + \Gamma_{AA} - \alpha(p_{A1} + p_{A2}) + \varepsilon_i.$$

This equation allows for Γ_{AA}, although this might plausibly be assumed to be 0, for there may be no added utility from having a diverse portfolio if a household buys two of the same type of vehicle. One story for why $\Gamma_{AA} > 0$ would be that by having two of the same type of vehicle, the household receives additional utility by showing peers that they identify with a particular type (e.g., two hybrids showing the household is eco-friendly or two pickups showing the household is "tough").

2.1 Implications for vehicle choice

This paper is about how consumers choose their portfolio of vehicles. In other words, in this simple setting, it is about how the kept vehicle (i.e., vehicle 1 in this setting) influences the choice of the second vehicle.

Continuing the thought experiment where the consumer can only choose options A and B for the second vehicle, we can consider the conditions under which A or B is chosen. Specifically, the household chooses portfolio AB rather than AA if $u_i^{AB} > u_i^{AA}$, which is equivalent to (assuming $\Gamma_{AA} = 0$)

$$f(\theta_B) + \Gamma_{AB} - \alpha p_{B2} > f(\theta_A) - \alpha p_{A2}. \tag{2.1}$$

This simple inequality indicates that the household will choose B when the net utility of the vehicle characteristics, lifetime ownership cost, and portfolio effect from B dominate the net utility of the vehicle characteristics and lifetime ownership cost from A. Rewritten differently, we have

$$\Gamma_{AB} > f(\theta_A) - f(\theta_B) + \alpha(p_{B2} - p_{A2})$$

This states that if the added utility from having a portfolio is greater than the difference in utility from the characteristics of the two types of vehicles plus the difference in the lifetime ownership cost of purchasing the vehicles of the two types (converted to utility terms), then the household will choose a vehicle of type B. In other words, B will be chosen if a positive effect from portfolio diversity is larger than the household "type" effect due to the household valuing the characteristics and lifetime ownership cost. This can be rephrased as an empirically testable prediction:

The portfolio effect will dominate if we observe the household choosing vehicle B when the first vehicle is A (as in the setting described so far), but the "type" effect will dominate if we observe the household choosing vehicle A when the first vehicle is also A.

This simple model provides further insight on the decision process. Let $g(.)$ be the distribution of utilities in the population. Then the simple choice between A and B for the new vehicle, the choice probabilities are given as follows:

$$Pr_{AB} = \int_{\mathbf{u}} I(u_{AB} > 0) I(u_{AB} > u_{AA}) dg(\mathbf{u}),$$
$$Pr_{AA} = \int_{\mathbf{u}} I(u_{AA} > 0) I(u_{AA} > u_{AB}) dg(\mathbf{u}).$$

2.2 Changes in choice with gasoline price or fuel economy standards

This simple model also lends itself to a set of policy-relevant comparative statics.

Changes in Gasoline Prices

Consider a permanent increase in gasoline prices, with perfect foresight of this change. This will imply that p_{B2} and p_{A2} will change based on the relative fuel economy of the two vehicle types. Such a change in the market will have two effects:

1. Direct Effect: The probability of the choosing the higher fuel economy vehicle will increase.

2. Indirect Effect: The relative prices of new vehicles in equilibrium will change, so that higher fuel economy vehicles will increase in price relative to others.

The indirect effect will work in the opposite direction as the direct effect, and depending on the marginal utility of income α, may even dominate. We will see the net of these two effects in the empirics.

Changes in Fuel Economy Standards

Consider the implementation of fuel economy standards. These will have a short-run and long-run effect. The long-run effect is through changing the direction of innovative activity by automobile manufacturers. The short-run effect would imply changing the relative prices of vehicles of different fuel economies, consistent with fuel economy standards putting a shadow price on fuel economy. Focusing only on the short-run effect, this would imply that p_{B2} and p_{A2} would change, so that the higher fuel economy vehicle would be relatively less expensive. This would have a similar effect to the "direct effect" from a permanent gasoline price change.

3 Data and Identification

The cornerstone of our dataset is the universe of California vehicle registration records that occurred from 2001-2007.[1] The DMV dataset includes every vehicle registered under the residential designation code. In California every vehicle must be registered annually. Each record includes the registrant's US Census block group identifier, the 17-digit vehicle identification number (VIN) that uniquely identifies the vehicle, that year's registration date, the date when the vehicle was

[1]We thank the California Department of Motor Vehicles (DMV) for making these data available for research.

last sold, and various other information. This information allows us to construct a household-level panel dataset vehicle ownership.

Basic vehicle attributes (e.g. horsepower, weight, etc) are available via a VIN decoder that we purchased from DataOne Software. We augment the decoder to include vehicle fuel economy, which is available from the US Environmental Protection Agency. Vehicle-miles traveled are available for each VIN whenever the vehicle is sold and upon receiving biannual Smog Check certification. We thus have a average measure of miles traveled by each vehicle and, by extension, each household for each year in our sample. The coarseness of these data are not optimal for examining high-frequency effects of VMT-switching between vehicles in response to changes in gasoline prices. Nonetheless, gasoline prices are a variable of interest in this study, since they affect the household's optimal portfolio of vehicle fuel economy. Our gasoline price data are at the county-month level.

3.1 Describing the Sample

In each year households are characterized by the starting and ending number of vehicles in their portfolio. In year t a household's starting portfolio size N^s is the number of vehicles registered in that year. If the household registers exactly N^s vehicles in year $t+1$ or $t+2$ then the ending portfolio size N^e in year t is N^s. If the number of vehicle registered in years $t+1$ and $t+2$ are identical, but not equal to N^s then the ending portfolio size is the number of vehicles registered in the later years.[2]

[Table 1 about here]

Table 1 shows the distribution of household portfolio transitions. Specifically, rows indicate the number of cars in year t, and columns indicate the number of cars in $t+1$. The table represents all possible household transitions. The large mass on diagonals indicates that many households do not increase or decrease the number of cars that they register from year to year. A careful interpretation of "0" is necessary: a household with 0 cars is not in our dataset, so transitions from 0 occur when a Californian household without a car in t registers one in $t+1$, or with observationally-equivalence, a household moves to California from another state. Similarly, transitions to 0 occur either when a household sells all of its registered cars, or if it exits the data via a move to another state or a dissolution of the household.

Many of the regressions that follow are estimated using a sample of two-car households that replace one of their cars, a sample which we call "2x2 replacement households".[3] While other

[2]We examine one and two years in the future as a household that may register more cars in one year than they ever owned simultaneously. For example, consider a household that owns two cars in year t. In year $t+1$ they re-register both previously owned vehicles ans the registrations expire. Then, toward the end of the year, they sell one vehicle and replace it with a new one (which requires registration of the new vehicle). This household has registered three unique vehicles in year $t+1$ but only ever owned two at any given time. In year $t+2$, barring the purchase of yet another new vehicle, the household would return to registering two vehicles.

[3]We define a household as replacing one vehicle if the starting (in year t) and ending (in year $t+1$ or $t+2$) portfolios differ by one vehicle. The household may conduct multiple vehicle transactions, as long as one of the two

transitions are certainly interesting, two-car replacement households provide the cleanest experiment. Households increasing the number of cars in their portfolio are likely to be experiencing an unobserved development that increases their demand for transportation (e.g. having a baby). Furthermore, it is unclear how to characterize the channels through which the portfolio of households with more than two cars affects replacement decisions. Does a portfolio effect for those households operate via the highest-VMT kept car, or the newest? Or must the portfolio effect be defined in a higher dimension? Given that no clear answer exists to these questions, we choose the simple path of focusing on the two-car replacement households.

Table 2 shows summary stats for all 2x2 replacement households, including segmentation based on the fuel economy of the bought car. Households that purchase relatively fuel efficient vehicles (gallons per mile quartile 1) tend to keep relatively fuel efficient cars as well. The converse is true for households buying fuel inefficient vehicles, suggesting that households may have an overall preference for either high- or low-GPM cars.

[Table 2 about here]

Many analyses that follow examine use quartile of fuel economy to describe bought and sold cars. The GPM cutoffs are presented in Table 3, along with their corresponding MPG analogs.

[Table 3 about here]

3.2 Identification

To understand the challenges associated with identifying the portfolio effect, we consider a thought exercise. For a given two-car household that replaces one of its vehicles, we would like to know the effect that randomly dropping one of the cars and exogenously perturbing fuel economy of the "kept" car has on the choice of fuel economy of the "bought" car. That is, we would like to randomly assign one car to be the "kept" car, and to randomly assign it a fuel economy (f^k) to see how changes in f^k affect the household's observed choice of f^b, the GPM of the car purchased. This is what we mean when we refer to the "portfolio effect". There are two identification challenges to operationalizing this thought experiment to retrieve an estimate of the portfolio effect in our observational dataset. We propose instrumental variables to address each.

Identification Challenge 1: Which Vehicle to Keep? As described earlier, our sample isolates two-car households that replace one of their cars with another. In general, the choice of which car to keep is endogenous and many potential stories could be told about preferences and conditions that would lead to one or the other of the cars being kept. Such a choice is inconsistent with our thought experiment of randomly assigning the household its f^k. However, our data offer several appealing instrumental variables.

vehicles appears in both the starting and ending portfolios. We do not consider households where both vehicles in the two-vehicle portfolio change as the relative timing of each purchase becomes important for defining the portfolio at the time of each vehicle's purchase.

A valid instrument will provide exogenous variation in the process that determines which of the household vehicles is kept and which is replaced. The exclusion restriction requires that the instrument affect the household's choice of f^b only indirectly, through the choice of which car to keep. We assert that variation in the price differential between the kept and dropped car contains such identifying variation. There are three functions of the price differential that we use. For exposition, let P_t^k and P_t^d be the average retail value of the kept and dropped cars, respectively, at the time when the car is dropped (t). The first candidate instrument is the price difference at time t: $\Delta P^{kt} = P_t^k - P_t^d$. One might be concerned that attributes of the car that are correlated with both the choice of which car to drop and the price difference, which would violate the exclusion restriction.

The second instrument is the change in price differences between time t and time 0, when the kept car was purchased. That is $\Delta\Delta P^{kd} = (P_t^k - P_t^d) - (P_0^k - P_0^d)$. To the extent that market forces are exogenous to portfolio preferences, this instrument has promise. However, one may be concerned that the change in relative prices was expected by the buyer in time t, and thus potentially correlated with preferences in time t as well.

The third candidate instrument addresses the above concerns by extracting only the portion variation in the price difference-in-difference that deviates from expectations. We assume that households form expectation using lagged 1-, 3-, and 5-year depreciation rates at the make-model-modelyear level, and project these into the future. Deviations from these projections are what we refer to as the "deviation difference-in-difference". We find it difficult to come up with a violation of the exclusion restriction for this instrument. Recall that the concern is that a correlation exists between portfolio preference exhibited in the initial purchase of the kept car and the instrument. Relying on an instrument using market-level changes in relative prices that arise only after the purchase of the kept car would be problematic only if those market level changes were correlated with individual household preferences over the vehicle portfolio.

[Figure 1 about here]

Figures 1 (a)-(c) display the reduced form relationship between these price differentials and the probability of dropping the least expensive car in the initial portfolio. Each of the instruments appears to have power. Since the relationships are non-linear, in our estimation we deploy them as third-order polynomials in their respective first stages.

Identification Challenge 2: Omitted Variables. The household's choice of f^k may be influenced by many factors that are unobservable to the researcher. These may include unobserved car attributes that are correlated with GPM (e.g. safety via weight) or unobserved household attributes (e.g. features of commutes). Of particular interest in our setting is fuel economy, and the confounding effect that unobservables may have on f^k. To address this identification challenge, we follow an instrumental variable approach and control for time-invariant household preferences via household fixed-effects.

Our preferred instrument for addressing omitted variables is the price of gasoline at the time of the kept car purchase, $p^{gas}_{it_k}$. Both theory and evidence (e.g. BKZ) demonstrate that households consider future operating costs of the vehicle in their purchase decision. Changes in California gasoline prices are exogenous with respect to the household choice, vary extensively over the time period of our data, and alter the expected lifecycle cost of vehicles according to each vehicle's fuel efficiency. Based on this logic, when gasoline prices are high at the time of the kept car purchase, we would expect the household to purchase a more fuel efficient car than when gasoline prices are low (as also demonstrated in BKZ). The price of gasoline at the time of the kept car purchase thus provides exogenous variation in the potentially endogenous variable of interest, f^k.

Recall that the relationship between f^k and f^b is theoretically ambiguous. A preference for diversification in the household portfolio will lead to a negative correlation, but complementarity between attributes associated with fuel economy may lead to a positive correlation. By extension, the relationship between $p^{gas}_{it_k}$ and f^b may also appear to be positive or negative.

The reduced form relationship between the instrument and our outcome variable of interest, f^b, is presented in Figures 2a and 2b. Many factors influence a consumer's choice of vehicle attributes, including f^b, so a plot of the raw data reveals little about the underlying relationship between our variables of interest. Instead, we present the variables after partialing out other covariates. The x-axis and y-axis are the residuals retrieved from regressing $p^{gas}_{it_k}$ and f^b, respectively, on covariates.

[Figures 2a and 2b about here]

A clear relationship emerges. For new cars, we observe a *lower* (conditional) f^b at higher levels of (conditional) $p^{gas}_{it_k}$. The relationship is most clear in the Kernel regression line. The relationship is reversed for used cars: a *higher* (conditional) f^b is observed at lower levels of (conditional) $p^{gas}_{it_k}$. For the purposes of our empirical approach, what is important here is that we observe a relationship between the instruments and the outcome variable of interest, which (if you believe the exclusion restriction) implies a strong first-stage.

Further Consideration: Within-Household Variation. We further refine our identification strategy to address potential concerns that first-order household preferences for cars with high (or low) fuel economy may overwhelm our ability to identify the potentially second-order preference for a diversified fuel economy portfolio. The richness of our panel dataset allow us to deploy household fixed effects to control for time-invariant unobservable preferences such as this. The importance of these controls can be seen via a simple example. Suppose that there are two types of households. They both prefer a diverse vehicle portfolio, but one (say household A) has an overall preference for gas guzzlers and the other (household B) for fuel efficiency. Examining each household's portfolio may reveal that household A holds cars that are both in the highest GPM quartile, whereas household B holds cars that are both in the lowest GPM quartile. Were we to randomly remove one of the cars from each portfolio, they would each be left with a car in the GPM quartile of their preference. They would also be likely to purchase a new car that is also in that

GPM quartile. On the surface, it would appear as though the households have a low preference for a diversified vehicle portfolio. However, that may be a false conclusion. Were we to examine GPM *within* the preferred quartile, we may discover that the household prefers diversification within that range. Using household fixed effects as controls allows our empirical approach to reveal the true impact of an exogenous marginal change in the fuel economy of the kept vehicle on the (marginal) fuel economy of the vehicle purchased.

Identifying household fixed effects requires observing at least two transactions per household, which imposes a restriction on our viable sample. Figures 3a - 3b present histograms of the number of transactions per household under various sample restrictions. It reveals that, while many households must be excluded to estimate specifications with household fixed effects, we are still left with approximately 235,000 households in the IV specification that includes household fixed effects.

[Figures 3a to 3b about here]

3.3 Regression Specifications

The basic regression strategies examine the relationship that fuel economy of the kept car has on the chosen fuel economy of the bought car. The dependent variable is thus either fuel economy of the bought car itself (f_{it}^b), or quartile indicators of that variable. Regressors of interest include gasoline price at the time of purchase, fuel economy of kept car (f_{it}^k), and their interaction. In addition, we include a term, $\mathbb{1}\{\Delta f^{kd} > 0\}$, that distinguishes which car was dropped from the initial portfolio, the low- or high-GPM car. Specifically, $\Delta f^{kd} = f_{it}^k - f_{it}^d$ in our main specification, such that $\mathbb{1}\{\Delta f^{kd} > 0\} = 1$ indicates that the car with the lower fuel economy in the initial portfolio is kept.

Most of the regression results that follow are retrieved from estimating a linear model of the probability of purchasing vehicles in a given MPG quartile. For ease of exposition of the results, and to allow our focus to rest on what happens in the top and bottom quartile, we combine vehicles in the 2nd and 3rd quartiles are into a single category, "med". The baseline specification is

$$Pr\left(q(f_{it}^b) = s\right) = \beta_0 + \beta_g p_{it}^{gas} + \beta_f f_{it}^k + \beta_{gf} p_{it}^{gas} \times f_{it}^k + \beta_{df} \mathbb{1}\{\Delta f^{kd} > 0\} + \alpha_X X_{it}^k + \varepsilon_{it} \tag{3.1}$$

where the dependent variable, $Pr\left(q(f_{it}^b) = s\right)$, equals one if f_{it}^b falls within the range of quartile $s \in \{1, med, 4\}$. We also estimate a continuous model where the dependent variable is f_{it}^b, keeping the rest of the specification as presented in equation 3.1. Fuel economy of the vehicles bought (b) and kept (k) by household i in time t are denoted f_{it}^b and f_{it}^k; i's contemporaneous gas price in t is p_{it}^{gas}, whereas $P_{it}^{gas\_k}$ is the price of gasoline *at the time household* i *purchased the car that it keeps in time* t. Control variables, denoted X_{it}, include vehicle attributes (e.g. class, make, value, age), nonparametric time controls (year and month-of-year fixed effects) and household/demographic (household fixed effects and county-level unemployment).

The IV specifications deploy instruments for the indicator of the kept vehicle's rank in fuel economy within the portfolio ($\mathbb{1}\{\Delta f^{kd} > 0\}$), the kept vehicle fuel economy (f^k_{it}), and the interaction of gas price and fuel economy ($p^{gas}_{it} \times f^k_{it}$). In each specification, we instrument using the gas price at the time the kept vehicle was purchased ($P^{gas_k}_{it_k}$) and that gas price interacted with the current gas price ($p^{gas}_{it} \times P^{gas_k}_{it_k}$). We augment this set of instruments with the instruments based on vehicle price differences that were briefly described in Section 3.2 on identification to estimate the following system of endogenous variables:

$$\mathbf{Z}_{it} = \begin{bmatrix} f^k_{it} & p^{gas}_{it} \times f^k_{it} & \mathbb{1}\{\Delta f^{kd} > 0\} \end{bmatrix}'$$

We now describe the vehicle price difference instruments precisely. In "Price Difference" specification, we include the difference in the current resale value of the kept and sold vehicles ($\Delta P^{kd}_{it} = P^k_{it} - P^d_{it}$) as an additional instrument. The "Price Difference-in-Difference" specification uses the change in value for the kept and dropped vehicles between the point the vehicle was purchased and the current time period: $\Delta\Delta P^{kd}_{it} = (P^k_{it} - P^k_{i0}) - (P^d_{it} - P^d_{i0})$.

The third instrument, which we call "Price Deviation Difference-in-Difference", is constructed from the deviation of the difference between the kept and dropped vehicles relative to their expected deprecation rates at the time of the kept car purchase. For each of the kept and dropped vehicle we estimate the households expectation of annual vehicle depreciation using depreciation of similar vehicles over the previous five years. Specifically, for vehicle make m and model year y, and value $V_{m,y,t}$ in year t, the expected depreciation is[4]

$$\mathbf{E}[Dep_{m,y,t}] = \left. \prod_{s=1}^{5} \frac{V_{m,y-s+1,t-s+1} - V_{m,y-s,t-s}}{V_{m,y-s,t-s}} \right)^{\frac{1}{5}} \tag{3.2}$$

We can then calculate the deviation from this expected depreciation rate for each car in the portfolio, and construct the "Price Deviation Difference-in-Difference" instrument. Assuming vehicle j has resale value $P_{j,t}$ in year t, this is:

$$\Delta\Delta V^{kd}_{it} = (P^k_{it} - \mathbf{E}[Dep^k_{it}] \cdot P^k_{i,t-1}) - (P^d_{it} - \mathbf{E}[Dep^d_{it}] \cdot P^d_{i,t-1}) \tag{3.3}$$

The set of three price difference instruments is $W = \{\Delta P^{kd}_{it}, \Delta\Delta P^{kd}_{it}, \Delta\Delta V^{kd}_{it}\}$. The first stage

[4] As a more concrete example, for a household in year $t = 2005$ owning a 2002 Honda Civic, the expected depreciation is the geometric mean annual depreciation rate of 2001 Hondas in 2004, 2000 Hondas in 2003, 1999 Hondas in 2002, etc.

thus consists of the following system of three equations for each of the instruments $w \in W$:

$$\begin{aligned}\mathbf{Z}_{it}^{w} = \boldsymbol{\Gamma}_0 + \boldsymbol{\Gamma}_g p_{it_k}^{gas} + \theta_P P_{it}^{gas\_k} + \boldsymbol{\Gamma}_{gg} p_{it}^{gas} \times p_{it_k}^{gas} \\ + \sum_{dc \in CLASS} \Big(\boldsymbol{\Gamma}_{dc} \mathbf{1}[CLASS_{it}^{dropped} = sc] \\ + \boldsymbol{\Gamma}_w w + \boldsymbol{\Theta} \mathbf{X}_{it} + \boldsymbol{\Xi}_{it}\end{aligned} \tag{3.4}$$

4 Results

The objective of this section is to present and justify our empirical approach, illuminate the effect of key variables on the choice of bought car fuel economy, and present some simple counterfactual analyses that demonstrate the policy-relevance of our findings. The section is comprised of two main parts.

First, we show results from various regression specifications. This allows us to demonstrate the importance of our instrumental variables approach and the inclusion of household fixed effects, both of which qualitatively alter key coefficient estimates. We then present marginal effects of gas prices and kept car fuel economy on bought car fuel economy, which reveals interesting features of household preferences for a diversified portfolio.

In the second subsection, we present results from counterfactual analyses. Specifically, we exogenously perturb the fuel economy of the kept vehicle, which is roughly what a successful fuel economy standard would do.[5] We observe the extent to which potential gasoline consumption reductions are either magnified or eroded due to portfolio considerations will become apparent.

4.1 Regressions and Marginal Effects

Tables 4 and 5 present the baseline regression results from new and used car purchases, respectively. The left-most column in each tables corresponds to estimates using the Price Difference IV from Equation ??. Column 2 presented estimates using the Price DiD IV shown in Equation ??. Finally, Column 3 is our preferred specification, deploying the Price Deviation DiD IV from Equation 3.4. The four panels correspond to the continuous dependent variable regression (Panel A), the linear probability model on the highest fuel economy quartile purchases (Panel B), the linear probability model on second and third fuel economy quartile purchases combined (Panel C), and the linear probability model on purchases of cars in the lowest fuel economy quartile (Panel D).

[Tables 4 and 5 about here]

Looking across the columns of Table 4 and 5 the specifications are robust to the specific choice of instruments. First stage power decreases from left to right as the instruments discard more

[5] A fuel economy standard will cause new car buyers to purchase more fuel efficient cars, on average. The question of interest in this counterfactual is whether and to what extent this effects the desired choice of fuel economy in the next vehicle purchase.

information contained in the relative resale prices of the kept and dropped vehicles. However, even our preferred instrument results in a powerful first stage; the Cragg-Donald minimum eigenvalue statistic from the system of first-stage IV regressions is reported below the table, and it reflects a very strong first stage for both new and used cars.

The inclusion of Δf^{kd} is important, reflecting the difference between dropping the lower- or higher-GPM vehicle in the portfolio. In Panel B, for example, a negative coefficient on Δf^{kd} would indicate that, for households dropping their fuel-efficient (low-GPM) car, a larger the difference between GPM in their initial portfolio leads the household to buy a low-GPM replacement with more likelihood (i.e. the household replaces the dropped car with a car of like kind). In Panel D, the logic is identical, but the sign is reversed: a *positive* coefficient on Δf^{kd} implies the fuel economy of the replacement car is higher than the kept car if the dropped car fuel economy was also higher.

The coefficients on Δf^{kd} are thus important litmus tests for the LPM results. If we were to find that households encouraged to drop their Prius due to unexpected change the relative value of vehicles in their portfolio generally then replaced it with a Hummer, that would be suspicious. This simple example provides a thought exercise for understanding the important role of household fixed effects. When the regression is identified using across-household variation, the evidence indicates that households will tend towards replacing their dropped car with one that is qualitatively similar in GPM to the kept car. In this way, there appear to be some households that prefer fuel efficiency and others that prefer gas guzzlers (presumably due to power, comfort, safety, etc). It is only when we look within the household that the portfolio effect of interest is seen. Despite the possibility that some households prefer low- or high-GPM cars in general, there appears to be a preference for diversity in GPM within that band.

To more clearly understand the effect of key covariates, we now present and discuss their marginal effects (Tables 6 to 7). First we discuss marginal effects of gasoline prices on the fuel economy of new and used cars, followed by the marginal effects of kept car GPM on bought car GPM. For readers preferring a graphical representation, these are shown in Figures 4 to 7.

[Tables 6 and 7 about here]

[Figures 4 - 7 about here]

The effect of gasoline prices on bought car GPM is quite heterogeneous across both new and used cars, and across kept car GPM quartiles. New car GPM is relatively unaffected by changes in gasoline prices, although there is an decrease in the probability of buying a car in the highest-GPM quartile (fuel inefficient) as gas prices rise. Surprisingly, the probability of buying a new car in the top quartile of fuel efficiency does not increase with gasoline prices. There are a multitude of effects on the probability of used car purchases across the distribution, however. This should not be surprising, given the equilibrium effects that gas prices induce on both the the supply- and demand-sides of the market.

Note that the continuous regressions (Panel A) shroud heterogeneous marginal effects that occur across the distribution. There does appear to be a little desire for a diversified portfolio when examining the marginal gas price effect on high- and low-GPM quartile used car purchases. Households with fuel efficient kept cars are more likely to purchase another high fuel-economy car as gas prices rise; similarly, households that keep a car in the lowest fuel economy quartile are less likely to buy a fuel efficient car as gas prices rise. Demand for middle-quartile GPM cars responds in an opposite way. Owners of fuel efficient cars are less likely to buy middle-quartile GPM cars as gas prices rise, whereas owners of low fuel economy cars are more likely to buy a middle-quartile GPM car. The results for low fuel-economy used purchases reflect the general equilibrium effects mentioned above. As gas prices rise, households that keep a fuel efficient car are more likely to buy a gas guzzler, as result that indicates supply-side effects in the used market outweigh demand effects. On the other hand, households that keep a fuel-inefficient car are less likely to buy another one as gas prices rise.

Marginal effects of kept car GPM provide a more direct representation of the portfolio effect. We present marginal effects of kept car fuel economy for new and used cars (and pooled) at different gasoline prices in Table 7. Note that both the level of the coefficients within each bought car GPM quartile, and the gradient of these coefficients with respect to the gas price, are important. Positive coefficients reflect an increasing probability of buying in a given quartile as kept car GPM increases (becomes less fuel efficient).

The overall story is clear: households exhibit a preference for GPM diversification in their portfolio, and that preference increases as the gas price rises. This can be seen clearly in the overall results presented in panel A, where negative coefficients imply an increased demand for buying a fuel efficient car as kept-car fuel efficiency decreases. A similar narrative holds when examining the high- and low-GPM bought car quartiles. Households buy new cars in the lowest GPM quartile (high fuel economy) with a higher probability as their kept car fuel economy decreases. Furthermore, this preference increases as the gas price rises. On the other hand, households have a lower probability of buying fuel inefficient new cars as kept car GPM increases, and this preference also increases as the gas price rises.

The situation for used cars is qualitatively similar. The negative coefficient on the marginal effects on 1st GPM quartile implies that the probability of buying a used, highly efficient car decreases in kept-car GPM at this gas price. We presume that this reflects a supply-side effect: households with fuel efficient used cars are less likely to sell them when gas prices are high, thus shifting upwards along the demand curve for this type of car.

4.2 Counterfactuals (Preliminary)

The results presented above provide a platform for examining how preferences over the household portfolio may affect energy conservation resulting from popular environmental policies. In particular, we view these results as particularly informative about long-run impacts of fuel economy standards, such as Corporate Average Fuel Economy (CAFE) standards in the United States. CAFE affects the suite of cars available for purchase, and their purchase prices. Once these cars become part of a household's vehicle portfolio, the change in attributes (relative to no CAFE standard) may influence the subsequent choice of vehicle purchased. In particular, if households exhibit a preference for diversification in their portfolio, increasing the fuel efficiency of their kept car will lead to a less fuel efficient second-car purchase.

The counterfactuals that we describe here examine the net effect of an exogenous decrease of 10% and 25% in the kept car GPM (i.e. a fuel efficiency increase) on predicted gasoline consumption. The net effect includes changes in gasoline consumption relating to use of the (now more efficient) kept car, but also changes in gasoline consumption relating to the use of the bought car whose GPM is influenced via the portfolio effect that we estimate above. Throughout the exercise, we assign cars the vehicle-miles traveled (VMT) that are observed in our dataset. We do not adjust these to account for a rebound effect, although in reality, one may exist.[6] Notice that this implies the only changes in gasoline consumption that we estimate will occur due to changes in the (extensive margin) choice of bought car fuel economy that are induced via the portfolio effect. It may well be the case that there is a first-order effect of gas prices on bought car GPM, but we seek to isolate only the portfolio effect in this exercise.

[Tables 8 and 9 about here]

Turning to Table 8, first note the observed gas consumption of the kept and bought vehicles. Each is estimated to consume approximately 550 gallons of gasoline per year (the product of VMT and vehicle GPM). Changes in kept car GPM do have a large effect; a 10% decrease in kept car GPM mathematically reduces kept car gas consumption by 10%, as can be seen in the top rows at each gas price level. The increases in bought car gas consumption reflect the portfolio effect. It is immediately apparent that these effects are quite large, and that the majority of gasoline conservation enjoyed by a kept car GPM improvement are eroded by the household's response via the choice of lower bought car fuel economy. Indeed, our estimates predict that the portfolio effect offsets about two thirds of the fuel savings from increasing the kept vehicle's fuel economy.

The story is qualitatively similar when the bought car is used instead of new. Table 9 reveals that households purchasing used cars have, on average, lower gasoline consumption associated with the vehichle added to the portfolio, but slightly higher gasoline consumption with the kept vehicle. The exogenous GPM changes associated with the kept car are 67% offset by the portfolio effect.

These results are quite startling and, if true, have unfortunate implications for the effectiveness of fuel economy standards as a way to reduce gasoline consumption. While it is clear that these portfolio effects cannot be directly applied to a fuel economy standard, the magnitude of the portfolio response implies that strong forces will be at work, particularly in the used car market,

[6] As fuel economy changes it, in turn, alters the cost per mile traveled. Consumers faced with this change in relative prices may choose a different VMT. See Borenstein (2015) and Gillingham et al. (2016) for more on the rebound effect.

which is not covered by CAFE. Increased demand for used, fuel inefficient cars will occur as a result of increased efficiency from CAFE. The increase in demand will lead used gas guzzlers to be more valuable, and thus more slow to be retired from the fleet (similar to the effect documented in Jacobsen and van Benthem (2015).

5 Conclusions

The effects of a number of polices applied the light duty vehicle market depend crucially on vehicle choice patterns. Typically empirical estimates of vehicle choice assume that the vehicle choices within a household are made independently. We provide show evidence that this assumption does not hold.

Using panel data on the portfolio of vehicles within a household, and a novel instrumental variables approach, we find evidence that households value diversification; exogenous increases in the fuel economy of the kept car lowers the fuel economy of the purchased car. We show this using both a continuous measure of fuel economy, as well as by estimating the probability a household purchases a vehicle in the upper and lower quartiles of the fuel economy distribution. Increases in the fuel economy of the kept car reduces the probability the household purchases a car in the lower quartile of gallons per mile, while such increases reduce the probability the household buys a car in the upper quartile.

We also find that gasoline prices affect the preferences for diversification in intuitive ways. As gasoline prices increase, the effect of the fuel consumption of kept vehicle and the probability of buying a car in the lower quartile of fuel consumption becomes even more positive. In contrast, as gasoline prices increase, the effect of the fuel consumption of kept vehicle and the probability of buying a car in the upper quartile of fuel consumption becomes even more negative.

To understand the economic important of this taste for diversification, we use our results to estimate the net effect of an exogenous increase in the fuel economy of the kept vehicle. These calculations suggest that the portfolio effect can have large consequences of the net affect a one-time increase in fuel economy; between 75 to 95 percent of the fuel savings from increasing the fuel economy of the kept vehicle are eroded from the resulting decrease in fuel economy of the newly purchased vehicle.

References

Borenstein, S. (2015). A microeconomic framework for evaluating energy efficiency rebound and some implications. *Energy Journal*, 36:1–21.

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

Gillingham, K. (2011). How do consumers respond to gasoline price shocks? heterogeneity in vehicle choice and driving behavior. Manuscript: Yale University.

Gillingham, K., Rapson, D., and Wagner, G. (2016). The rebound effect and energy efficiency policy. *Review of Environmental Economics and Policy*, 10:68–88.

Jacobsen, M. R. and van Benthem, A. A. (2015). Vehicle scrappage and gasoline policy. *American Economic Review*, 105(3):1312–1338.

Klier, T. and Linn, J. (2010). New vehicle characteristics and the cost of the corporate average fuel economy standard. Resources for the Future Discussion Paper 10-50.

Table 1: Number of Unique Households by Portfolio Size

Start Portfolio Size	End Portfolio Size 1	2	3	4+
1	7,262,111	1,360,594	187,558	75,150
2	1,172,278	4,632,425	839,546	259,098
3	168,745	849,703	2,169,948	675,040
4+	35,810	141,618	381,226	1,489,926

Each cell represents the count of unique households from 2001 to 2007 observed to have the starting portfolio size shown in each row and the ending portfolio size shown in the column. These counts provide a measure of the number of households providing identifying variation in each portfolio cell. A single household may appear in multiple cells if their portfolio changes over time but is counted at most once in each cell. For example, two-car household that replaces one car every year would add one to the count of the (2,2) cell. If instead, that household adds a third vehicle in 2004 and returns to a two-car portfolio in 2006 it would add one to the count of the (2,2) cell, one to the count of the (2,3) cell, one to the (3,3) cell, and one to the count of the (3,2) cell. Each household may have zero, one, or multiple vehicle transactions during this time period.

Table 2: Summary Statistics - 2x2 Replacement Households

	All Households	Bought GPM Qtile 1	Bought GPM Qtile 2 or 3	Bought GPM Qtile 4
N Households	1,452,896	392,168	768,517	413,367
N Transactions	2,004,312	491,010	1,003,044	510,258
Mean Transactions/HH	1.34	1.24	1.24	1.23
Std. dev	(0.676)	(0.504)	(0.612)	(0.491)
Median Transactions/HH	1.00	1.00	1.00	1.00
IQR	[2.00 - 1.00]	[1.00 - 1.00]	[1.00 - 1.00]	[1.00 - 1.00]
Median Bought GPM	0.05	0.04	0.05	0.07
IQR	[0.04 - 0.06]	[0.04 - 0.04]	[0.05 - 0.05]	[0.06 - 0.07]
Median Kept GPM	0.05	0.05	0.05	0.05
IQR	[0.05 - 0.06]	[0.04 - 0.06]	[0.05 - 0.06]	[0.05 - 0.06]
Median Bought MPG	20.23	26.01	19.97	15.19
IQR	[17.20 - 22.60]	[23.80 - 27.80]	[18.80 - 21.20]	[14.20 - 16.20]
Median Kept MPG	19.97	20.59	19.88	19.55
IQR	[17.20 - 22.00]	[17.80 - 23.20]	[17.30 - 21.80]	[16.80 - 21.50]

Table 3: Distribution of observed fuel economy

Percentile	Gallons per Mile (GPM)	Miles Per Gallon (MPG)
25th Percetile	0.045	22.0
Median	0.052	19.3
75th Percentile	0.059	17.0

Table 4: Regression Estimates - New Vehicle Purchases

(a) Panel A: Bought Vehicle Continuous GPM Measure

	Price Diff (1)	Price DiD (2)	Price Deviation DiD (3)
Gas Price ($/ gal)	0.052	0.031	0.032
	(0.009)**	(0.009)**	(0.010)**
Kept GPM	1.525	0.731	0.349
	(0.398)**	(0.413)	(0.426)
Gas Price × Kept GPM	-0.991	-0.584	-0.596
	(0.177)**	(0.170)**	(0.191)**
Δ GPM > 0	-0.001	0.001	0.000
	(0.001)	(0.001)	(0.001)

(b) Panel B: Bought Vehicle 1st GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	-0.731	-0.228	-0.977
	(0.342)*	(0.322)	(0.447)*
Kept GPM	-3.671	15.711	-22.396
	(15.081)	(14.315)	(18.676)
Gas Price × Kept GPM	13.608	3.929	17.943
	(6.505)*	(6.092)	(8.498)*
Δ GPM > 0	-0.086	-0.119	-0.055
	(0.033)**	(0.031)**	(0.045)

(c) Panel C: Bought Vehicle 2nd/3rd GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	-1.110	-1.570	-0.290
	(0.546)*	(0.585)**	(0.547)
Kept GPM	-63.774	-85.723	-8.027
	(22.248)**	(27.381)**	(23.113)
Gas Price × Kept GPM	21.483	30.147	6.316
	(10.398)*	(11.107)**	(10.386)
Δ GPM > 0	0.129	0.156	0.063
	(0.046)**	(0.050)**	(0.055)

(d) Panel D: Bought Vehicle 4th GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	1.841	1.798	1.267
	(0.488)**	(0.518)**	(0.457)**
Kept GPM	67.445	70.011	30.424
	(19.470)**	(24.383)**	(19.583)
Gas Price × Kept GPM	-35.091	-34.076	-24.259
	(9.290)**	(9.832)**	(8.677)**
Δ GPM > 0	-0.042	-0.037	-0.008
	(0.040)	(0.043)	(0.043)
N	440,809	429,369	348,368
Cragg-Donald stat	145.88	141.44	91.34
Household FE	Yes	Yes	Yes
IV for Kept Vehicle	Base+$\Delta PriceDiD^3$	Base+$PriceDiD^3$	Base+ValueDiD
Subsample	New	New	New

Standard errors robust to heteroskedasticity shown in parentheses.

Table 5: Regression Estimates - Used Vehicle Purchases

(a) Panel A: Bought Vehicle Continuous GPM Measure

	Price Diff (1)	Price DiD (2)	Price Deviation DiD (3)
Gas Price ($/ gal)	0.045	0.027	0.057
	(0.018)*	(0.013)*	(0.022)**
Kept GPM	1.086	0.234	1.532
	(0.806)	(0.504)	(0.909)
Gas Price × Kept GPM	-0.874	-0.523	-1.090
	(0.349)*	(0.245)*	(0.421)**
Δ GPM > 0	0.008	0.008	0.007
	(0.001)**	(0.001)**	(0.001)**

(b) Panel B: Bought Vehicle 1st GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	4.200	3.076	3.081
	(1.137)**	(0.772)**	(1.117)**
Kept GPM	207.785	118.514	109.659
	(50.188)**	(30.916)**	(46.832)*
Gas Price × Kept GPM	-80.034	-58.007	-58.326
	(21.792)**	(14.729)**	(21.382)**
Δ GPM > 0	-0.546	-0.484	-0.489
	(0.061)**	(0.046)**	(0.065)**

(c) Panel C: Bought Vehicle 2nd/3rd GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	-11.152	-7.979	-9.952
	(2.535)**	(1.605)**	(2.671)**
Kept GPM	-485.343	-251.088	-332.787
	(111.819)**	(65.004)**	(111.180)**
Gas Price × Kept GPM	213.213	151.329	189.695
	(48.602)**	(30.613)**	(51.132)**
Δ GPM > 0	0.582	0.432	0.544
	(0.134)**	(0.094)**	(0.151)**

(d) Panel D: Bought Vehicle 4th GPM Quartile

	(1)	(2)	(3)
Gas Price ($/ gal)	6.952	4.902	6.871
	(1.618)**	(1.033)**	(1.836)**
Kept GPM	277.558	132.574	223.128
	(71.442)**	(41.549)**	(76.125)**
Gas Price × Kept GPM	-133.179	-93.323	-131.369
	(31.022)**	(19.698)**	(35.146)**
Δ GPM > 0	-0.035	0.052	-0.055
	(0.085)	(0.060)	(0.103)
N	500,882	461,425	364,909
Cragg-Donald stat	39.99	42.48	38.93
Household FE	Yes	Yes	Yes
IV for Kept Vehicle	Base+$\Delta PriceDiD^3$	Base+$PriceDiD^3$	Base+ValueDiD
Subsample	Used	Used	Used

Standard errors robust to heteroskedasticity shown in parentheses.

Table 6: Marginal Effect of Gasoline Price on Bought Vehicle GPM - Preferred Specification

(a) Panel A: Bought Vehicle Continuous GPM Measure

Kept Vehicle GPM	All (1)	New (2)	Used (3)
25th Pctile GPM (0.045)	0.00410	0.00453	0.00717
	(0.00111)**	(0.00147)**	(0.00290)*
Median GPM (0.052)	0.00041	0.00070	0.00015
	(0.00031)	(0.00058)	(0.00054)
75th Pctile GPM (0.058)	-0.00317	-0.00302	-0.00666
	(0.00090)**	(0.00121)*	(0.00254)**

(b) Panel B: Bought Vehicle 1st GPM Quartile

	(1)	(2)	(3)
25th Pctile GPM (0.045)	0.15223	-0.16149	0.42954
	(0.05208)**	(0.06575)*	(0.14765)**
Median GPM (0.052)	0.02598	-0.04597	0.05401
	(0.01608)	(0.02728)	(0.03124)
75th Pctile GPM (0.058)	-0.09651	0.06611	-0.31032
	(0.04205)*	(0.05472)	(0.13038)*

(c) Panel C: Bought Vehicle 2nd/3rd GPM Quartile

	(1)	(2)	(3)
25th Pctile GPM (0.045)	-0.62143	-0.00311	-1.32956
	(0.10266)**	(0.08032)	(0.35277)**
Median GPM (0.052)	-0.07545	0.03756	-0.10822
	(0.03161)*	(0.03294)	(0.07349)
75th Pctile GPM (0.058)	0.45426	0.07702	1.07671
	(0.08235)**	(0.06653)	(0.31155)**

(d) Panel D: Bought Vehicle 4th GPM Quartile

	(1)	(2)	(3)
25th Pctile GPM (0.045)	0.46920	0.16460	0.90003
	(0.07463)**	(0.06636)*	(0.24244)**
Median GPM (0.052)	0.04947	0.00841	0.05421
	(0.02210)*	(0.02574)	(0.04896)
75th Pctile GPM (0.058)	-0.35774	-0.14313	-0.76639
	(0.06005)**	(0.05475)**	(0.21349)**

Marginal effect of the current gasoline price on the probability a household purchases a vehicle in the GPM quartile specified in table section header. Delta method standard errors robust to heteroskedasticity shown in parentheses. *,**,*** denote results significant at the 10%, 5%, and 1% levels, respectively.

Table 7: Marginal Effect of Kept Vehicle MPG on Bought Vehicle GPM - Preferred Specification

(a) Panel A: Bought Vehicle Continuous GPM Measure

Current Gas Price	All (1)	New (2)	Used (3)
Gas Price = 2.00/gal	-0.819	-0.843	-0.648
	(0.140)**	(0.202)**	(0.242)**
Gas Price = 3.00/gal	-1.393	-1.438	-1.738
	(0.235)**	(0.283)**	(0.453)**
Gas Price = 4.00/gal	-1.966	-2.034	-2.828
	(0.369)**	(0.438)**	(0.840)**

(b) Panel B: Bought Vehicle 1st GPM Quartile

	(1)	(2)	(3)
Gas Price = 2.00/gal	-12.832	13.490	-6.992
	(6.902)	(8.860)	(13.288)
Gas Price = 3.00/gal	-32.442	31.433	-65.318
	(11.113)**	(12.649)*	(23.234)**
Gas Price = 4.00/gal	-52.051	49.376	-123.643
	(17.249)**	(19.646)*	(42.632)**

(c) Panel C: Bought Vehicle 2nd/3rd GPM Quartile

	(1)	(2)	(3)
Gas Price = 2.00/gal	71.753	4.605	46.603
	(13.300)**	(11.011)	(30.426)
Gas Price = 3.00/gal	156.553	10.922	236.297
	(21.807)**	(15.438)	(55.240)**
Gas Price = 4.00/gal	241.353	17.238	425.992
	(33.967)**	(23.899)	(102.010)**

(d) Panel D: Bought Vehicle 4th GPM Quartile

	(1)	(2)	(3)
Gas Price = 2.00/gal	-58.920	-18.095	-39.611
	(9.487)**	(9.058)*	(20.474)
Gas Price = 3.00/gal	-124.111	-42.355	-170.980
	(15.808)**	(12.538)**	(37.907)**
Gas Price = 4.00/gal	-189.302	-66.614	-302.349
	(24.757)**	(19.569)**	(70.180)**

Marginal effect of the current gasoline price on the probability a household purchases a vehicle in the GPM quartile specified in table section header. Delta method standard errors robust to heteroskedasticity shown in parentheses. *,**,*** denote results significant at the 10%, 5%, and 1% levels, respectively.

Table 8: Net Effect of Kept Vehicle GPM Changes on Gasoline Consumption, New Vehicles

(a) 10% Reduction in Kept Vehicle GPM

Vehicle	Observed Gasoline Consumption (gal/yr)	Change in Gasoline Consumption (%Δ Kept vehicle gallons): Price Diff	Price DiD	Price Deviation DiD
Kept	537.64	-10.00	-10.00	-10.00
Bought	555.34	5.07	4.13	6.94
		(1.13)	(0.95)	(1.67)
Total	1,092.98	-4.93	-5.87	-3.06
		(1.13)	(0.95)	(1.67)

(b) 25% Reduction in Kept Vehicle GPM

Vehicle	Observed Gasoline Consumption (gal/yr)	Change in Gasoline Consumption (%Δ Kept vehicle gallons): Price Diff	Price DiD	Price Deviation DiD
Kept	537.64	-25.00	-25.00	-25.00
Bought	555.34	12.67	10.32	17.34
		(2.82)	(2.37)	(4.19)
Total	1,092.98	-12.33	-14.68	-7.66
		(2.82)	(2.37)	(4.19)

Predicted average change in fuel consumption resulting from an exogenous decrease in kept vehicle GPM of 10% (e.g., from 27.5 MPG to 30.6 MPG) or 25% (e.g., from 27.5 MPG to 36.7 MPG) for new vehicle purchases. Change in fuel economy expressed as percentage of annual gasoline consumption of the kept vehicle. Standard errors computed using the delta method shown in parentheses.

Table 9: Net Effect of Kept Vehicle GPM Changes on Gasoline Consumption, Used Vehicles

(a) 10% Reduction in Kept Vehicle GPM

	Observed Gasoline	Change in Gasoline Consumption (%Δ Kept vehicle gallons)		
Vehicle	Consumption (gal/yr)	Price Diff	Price DiD	Price Deviation DiD
Kept	569.12	-10.00	-10.00	-10.00
Bought	537.29	6.18	6.15	6.66
		(1.42)	(2.12)	(2.28)
Total	1,106.41	-3.82	-3.85	-3.34
		(1.42)	(2.12)	(2.28)

(b) 25% Reduction in Kept Vehicle GPM

	Observed Gasoline	Change in Gasoline Consumption (%Δ Kept vehicle gallons)		
Vehicle	Consumption (gal/yr)	Price Diff	Price DiD	Price Deviation DiD
Kept	569.12	-25.00	-25.00	-25.00
Bought	537.29	15.44	15.38	16.64
		(3.55)	(5.29)	(5.69)
Total	1,106.41	-9.56	-9.62	-8.36
		(3.55)	(5.29)	(5.69)

Predicted average change in fuel consumption resulting from an exogenous decrease in kept vehicle GPM of 10% (e.g., from 27.5 MPG to 30.6 MPG) or 25% (e.g., from 27.5 MPG to 36.7 MPG) for used vehicle purchases. Change in fuel economy expressed as percentage of annual gasoline consumption of the kept vehicle. Standard errors computed using the delta method shown in parentheses.

Figure 1: Instrumental Variables Reduced Form Relationships

(a) Price Difference IV

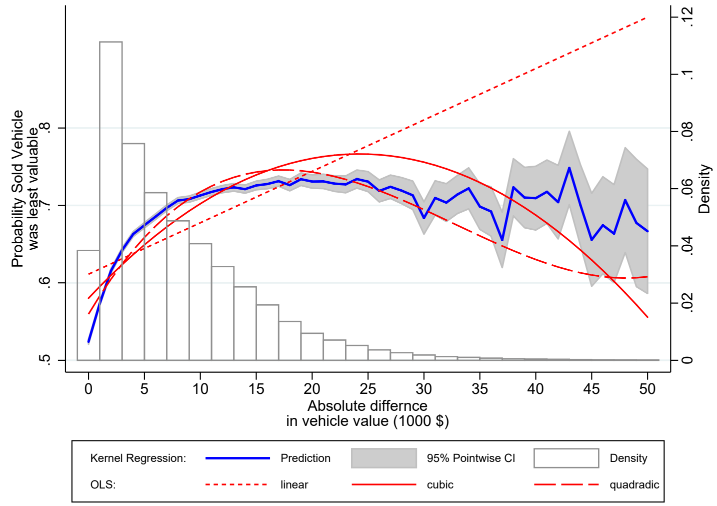

(b) Price DiD IV

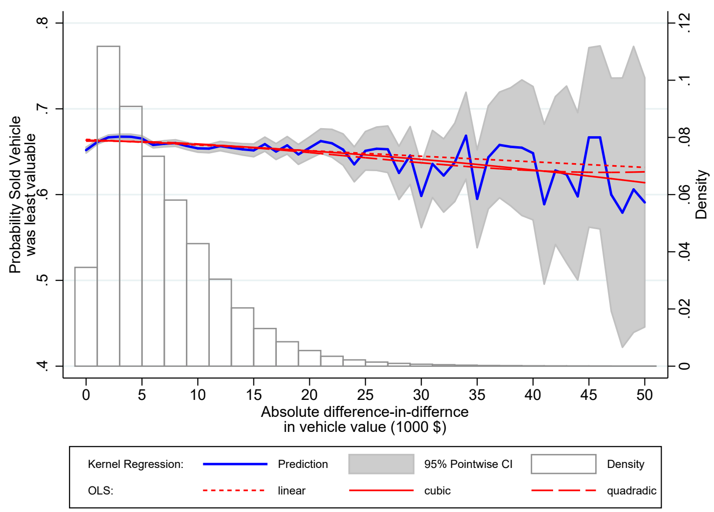

(c) Price Deviation DiD IV

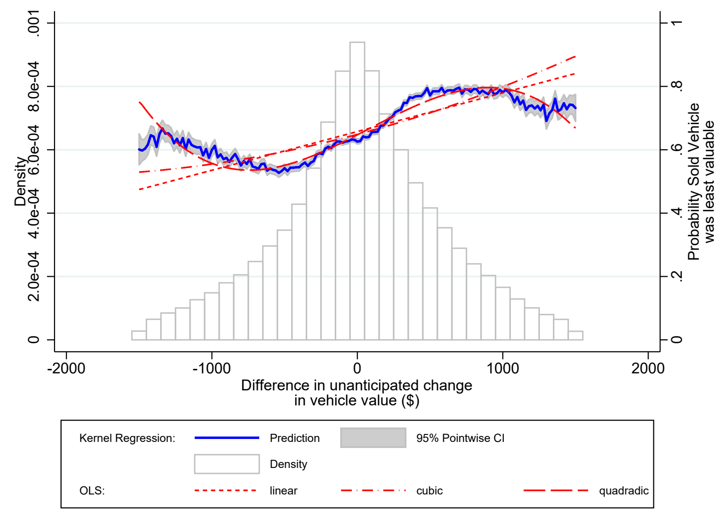

All 2x2 households. Probabilities conditional on a vehicle purchase (new or used) estimated within $1,000 bins. Binomial 95% confidence intervals shown in dashed lines. Values of the instruments in the Price Difference IV and Price DiD IV less than or greater than zero perfectly predict the least valuable vehicle in the portfolio and graphs are shown for the absolute value of these variables.

Figure 2: Reduced form relationship: Gas price at time of kept car purchase

(a) New Vehicle Purchases

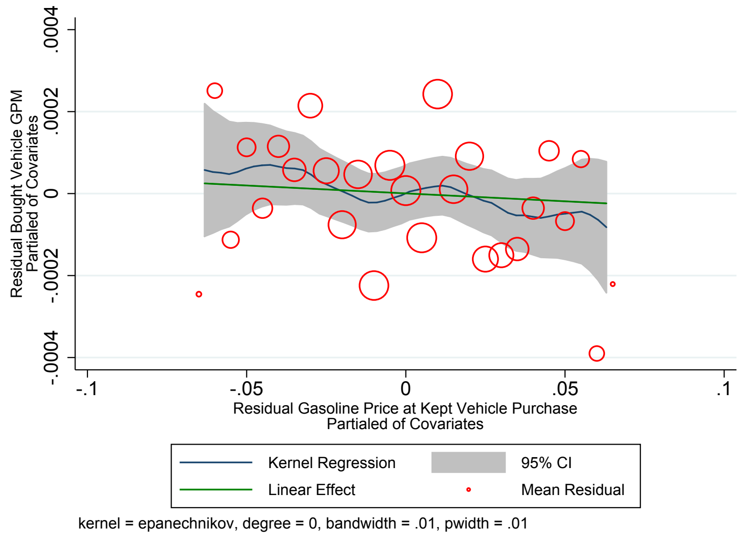

(b) Used Vehicle Purchases

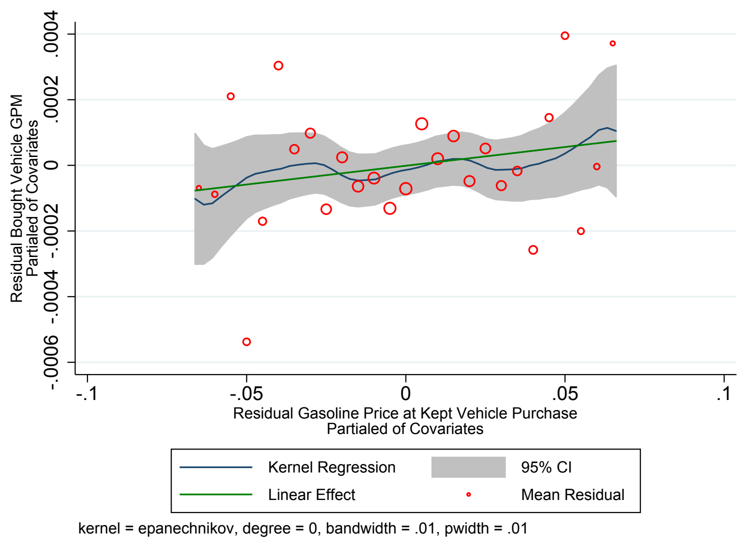

Plot of the reduced-form relationship between gasoline price at the time of kept vehicle purchase and the fuel economy (in GPM) of the purchased vehicle. Both variables are partialed of all other regression covariates. Graphs are limited to the 1st through 99th percentiles of residual kept vehicle gasoline price. Excludes observations where the household fixed effect perfectly predicts the outcome. Blue line is a kernel regression with Epanechnikov kernel and bandwidth 0.1. The gray band is the 95% confidence interval using the same kernel and bandwidth. The green line is the linear relationship estimated using OLS. Red circles are mean residuals for each 0.005 in kept vehicle GPM. The size of each circle is proportional to the number of observations used to compute the mean residual.

Figure 3: Number of Transactions per 2x2 Replacement Household

(a) Full Sample

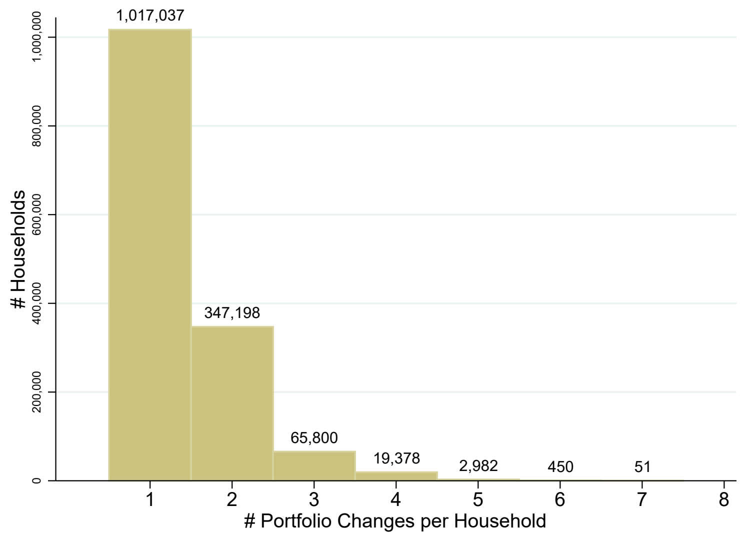

(b) IV Sample

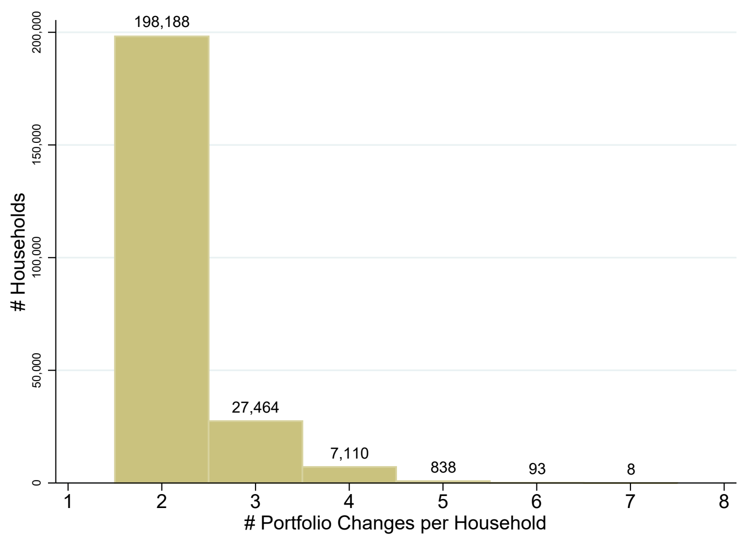

Distribution of the total number of observed vehicle transactions for each household from 2001 to 2007 for (a) the full sample of 2x2 replacement households and (b) households for which the data support deploying our IVs. In specifications including household fixed effects the fixed effect perfectly predicts the decision of a household if it only engages in one transaction. Other model parameters are identified by households engaging in multiple transactions from 2001 to 2007.

Figure 4: Marginal Effect of Gasoline Price, New Vehicles

(a) Price Difference IV

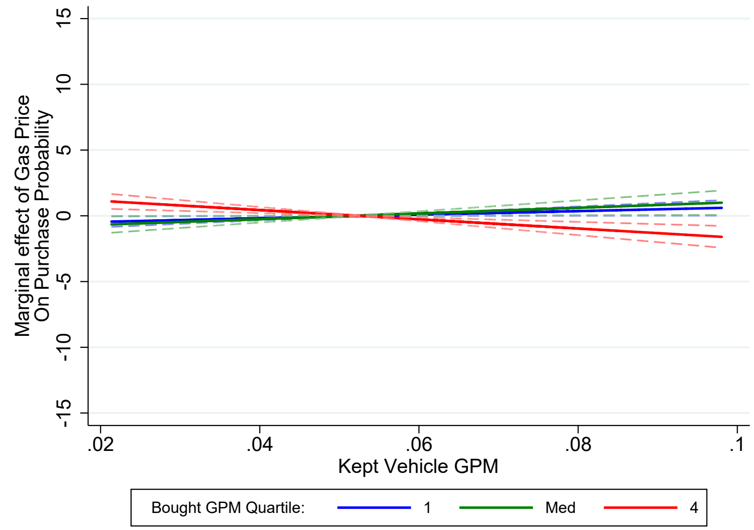

(b) Price DiD IV

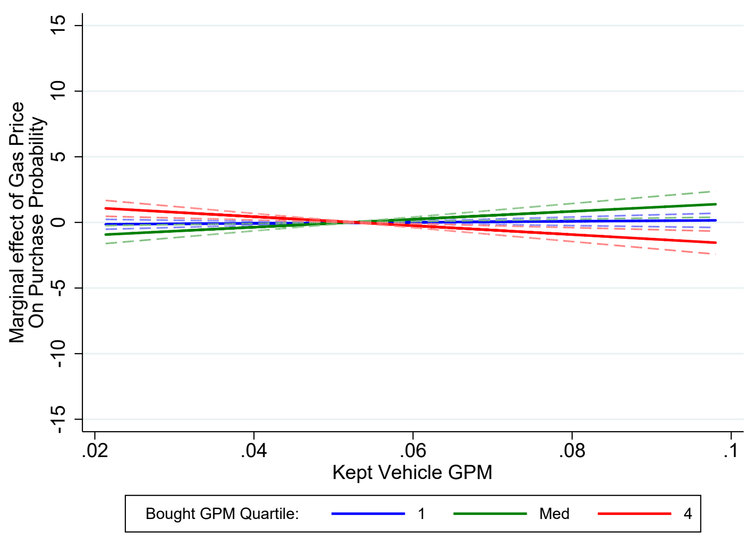

(c) Price Deviation DiD IV

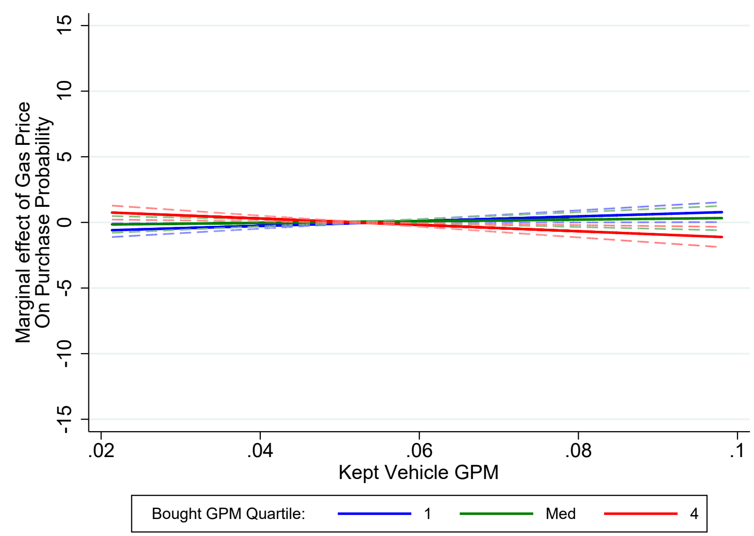

Pointwise 95% confidence intervals robust to heteroskedasticity shown in dashed lines.

Figure 5: Marginal Effect of Gasoline Price, Used Vehicles

(a) Price Difference IV

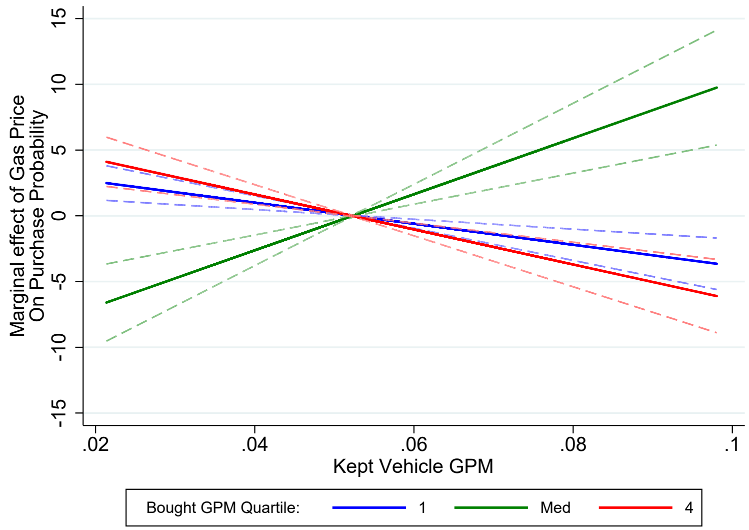

(b) Price DiD IV

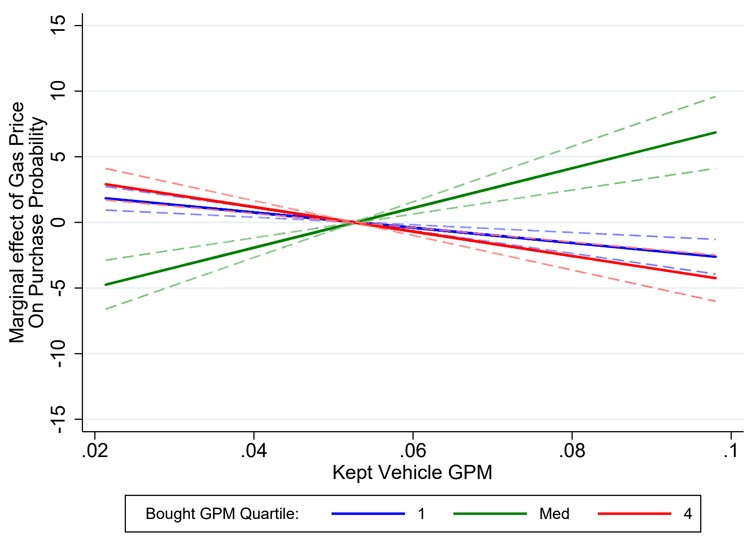

(c) Price Deviation DiD IV

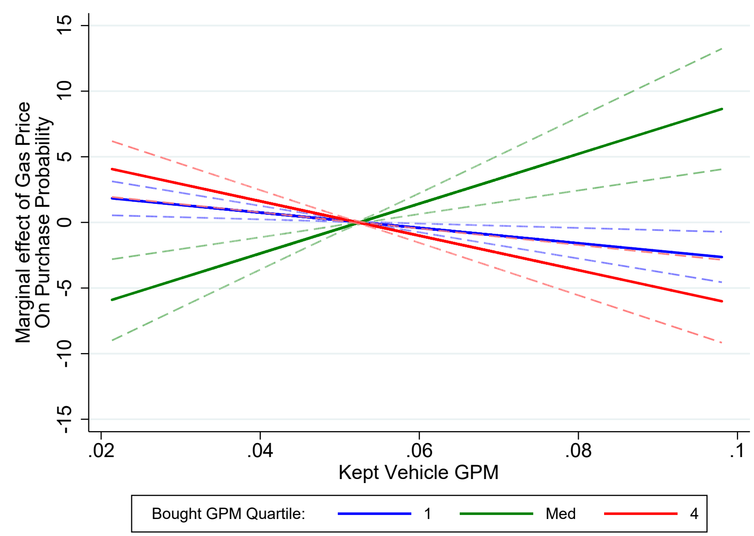

Pointwise 95% confidence intervals robust to heteroskedasticity shown in dashed lines.

Figure 6: Marginal Effect of Kept Vehicle GPM, New Vehicles

(a) Price Difference IV

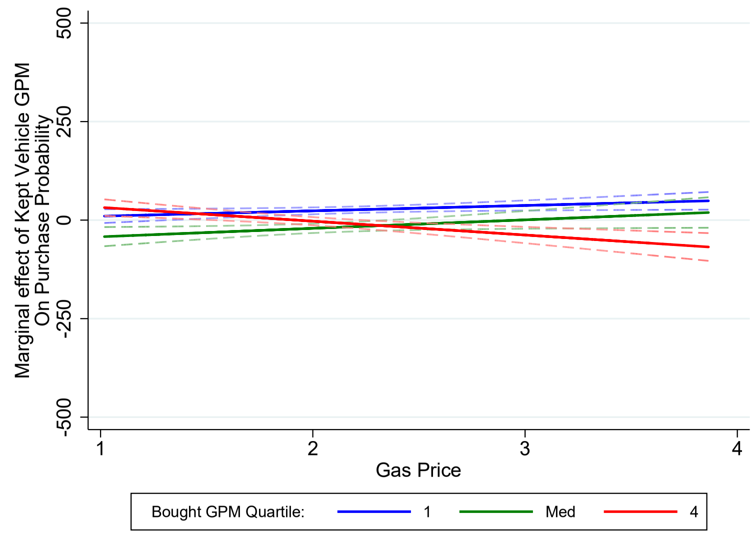

(b) Price DiD IV

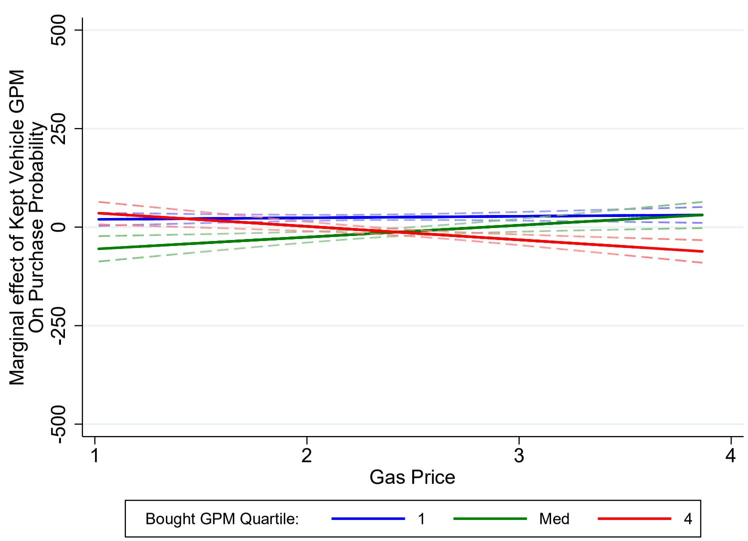

(c) Price Deviation DiD IV

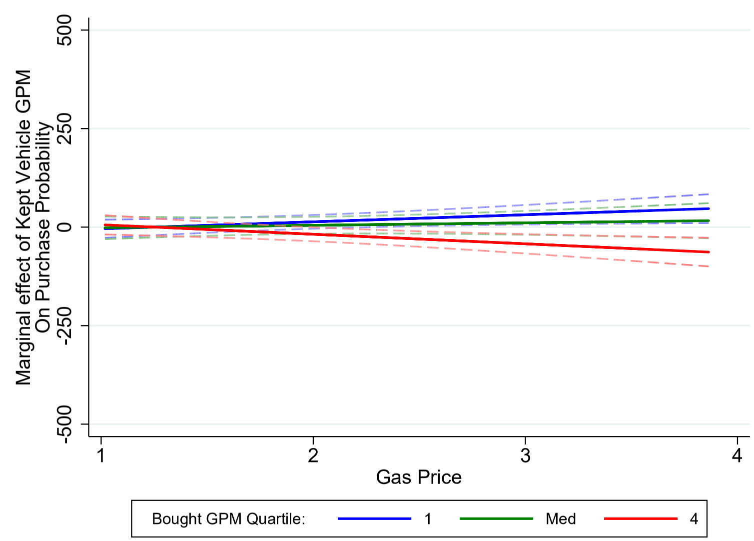

Pointwise 95% confidence intervals robust to heteroskedasticity shown in dashed lines.

Figure 7: Marginal Effect of Kept Vehicle GPM, Used Vehicles

(a) Price Difference IV

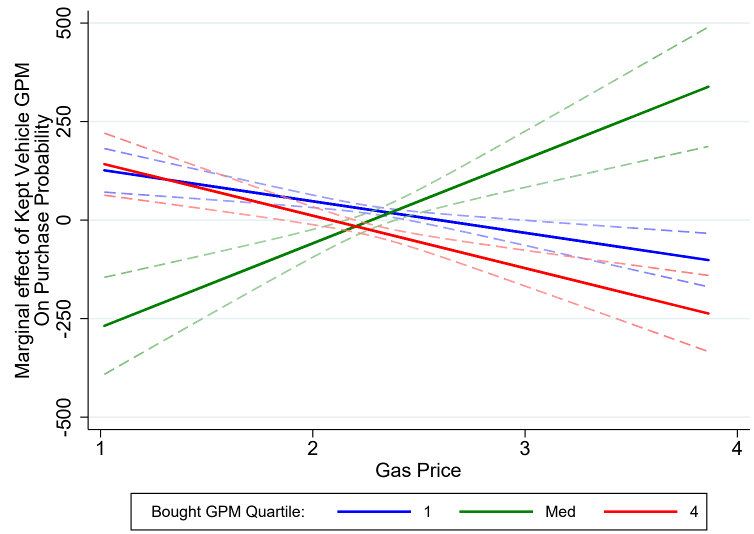

(b) Price DiD IV

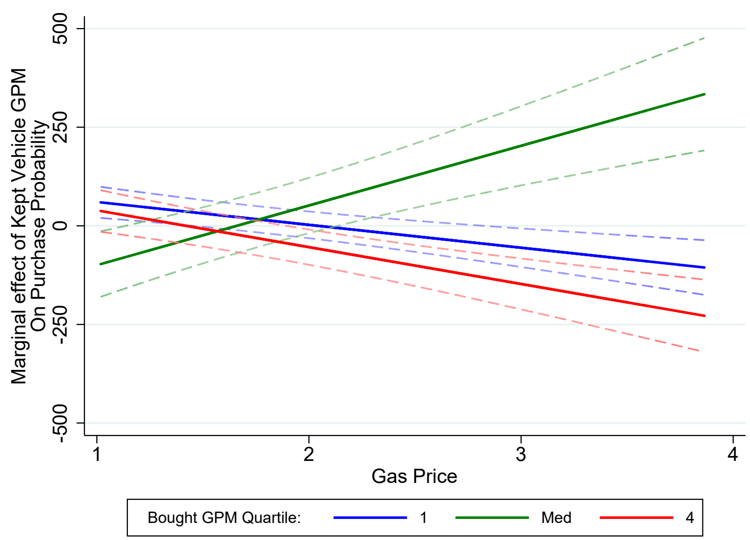

(c) Price Deviation DiD IV

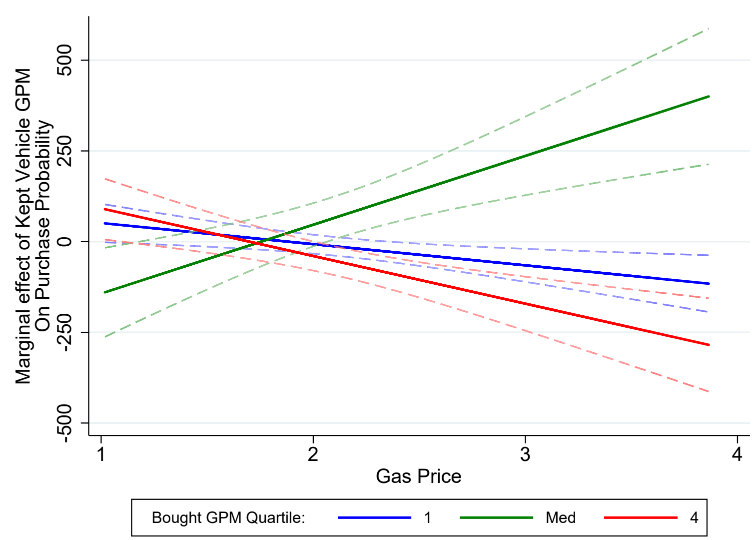

Pointwise 95% confidence intervals robust to heteroskedasticity shown in dashed lines.

The Welfare Impact of Second Best Uniform-Pigouvian Taxation: Evidence from Transportation

Christopher R. Knittel and Ryan Sandler*

September 20, 2015

Abstract

A basic tenet of economics posits that when consumers or firms don't face the true social cost of their actions, market outcomes are inefficient. In the case of negative externalities, Pigouvian taxes are one way to correct this market failure, where the optimal tax leads agents to internalize the true cost of their actions. A practical complication, however, is that the level of externality nearly always varies across economic agents and directly taxing the externality may be infeasible. In such cases, policy often taxes a product correlated with the externality. For example, instead of taxing vehicle emissions directly, policy makers may tax gasoline even though per-gallon emissions vary across vehicles. This paper estimates the implications of this approach within the personal transportation market. We have three general empirical results. First, we show that vehicle emissions are positively correlated with a vehicle's elasticity for miles traveled with respect to fuel prices (in absolute value)—i.e., dirtier vehicles respond more to fuel prices. This correlation substantially increases the optimal second-best *uniform* gasoline tax. Second, and perhaps more importantly, we show that the optimal second-best tax performs very poorly in eliminating deadweight loss associated with vehicle emissions; in many years in our sample over 75 percent of the deadweight loss remains under the optimal second-best gasoline tax. Substantial improvements to market efficiency require differentiating based on vehicle type, for example vintage. Finally, a uniform gasoline tax performs poorly on equity grounds as well. Such a tax would be highly regressive, and substantially more regressive than a direct tax on emissions.

*This paper has benefited from conference discussion by Jim Sallee and conversations with Severin Borenstein, Joseph Doyle, Michael Greenstone, Michael Grubb, Jonathan Hughes, Dave Rapson, Nicholas Sanders, and Catherine Wolfram. The paper has also benefited from participants at the NBER Energy and Environmental Economics Spring Meeting and seminar participants at Northeastern University, University of Chicago, MIT, and Yale University. We gratefully acknowledge financial support from the University of California Center for Energy & Environmental Economics. The research was also supported by a grant from the Sustainable Transportation Center at the University of California Davis, which receives funding from the U.S. Department of Transportation and Caltrans, the California Department of Transportation, through the University Transportation Centers program. The views expressed in this article are those of the authors and do not necessarily reflect those of the Federal Trade Commission. *Knittel:* William Barton Rogers Professor of Energy Economics, Sloan School of Management, MIT and NBER, *email:* knittel@mit.edu. *Sandler:* Federal Trade Commission, *email:* rsandler@ftc.gov.

1 Introduction

A basic tenet of economics posits that when consumers or firms do not face the true social cost of their actions, market outcomes are inefficient. In the case of externalities, Pigouvian taxes can correct this market failure, and the optimal tax leads agents to internalize the true cost of their actions. A practical complication is that directly taxing the externality may be infeasible. In such cases, policy makers might tax a product correlated with the externality. This introduces an additional complication: the level of externality generated can vary across agents. For example, instead of taxing vehicle emissions, policy makers tax gasoline even though emissions per gallon of gasoline consumed varies across vehicles. Similarly, a uniform alcohol tax may be imposed to reduce the negative externalities associated with use, even though externalities likely vary by the type of alcohol or who is consuming it. We refer to uniform taxes intended to address a heterogenous externality *second-best optimal (SBO) taxes*. When the level of externalities produced differs across consumers, a uniform tax will be second best and deadweight loss will remain.

In this paper, we study the size of the SBO gasoline tax, the amount of deadweight loss (DWL) from pollution that would remain if this tax is imposed, and the incidence of such a tax in the California personal transportation market between 1998 and 2008. Policy makers are often concerned about four externalities in the transportation sector: (1) local pollution from tailpipe emissions, known as criteria pollutants,[1] (2) climate change externalities resulting from carbon dioxide associated with the engine's combustion process, (3) road congestion, and (4) externalities associated with accidents. For all but the climate change externality, a gasoline tax is an imperfect instrument. While fuel consumption is positively correlated with criteria pollutant emissions, congestion, and accident externalities, it is not perfectly correlated. In contrast, burning a gallon of gasoline leads to the same amount of carbon dioxide emissions regardless of the vehicle, so a gasoline tax is the optimal instrument for climate change externalities.

In the case of the local pollution externalities of driving, the relationship between the

[1] Criteria air pollutants are the only air pollutants for which the Administrator of the U.S. Environmental Protection Agency has established national air quality standards defining allowable ambient air concentrations. Congress has focused regulatory attention on these pollutants (i.e., carbon monoxide, lead, nitrogen dioxide, ozone, particulate matter, and sulfur dioxide) because they endanger public health and are widespread throughout the United States

second-best optimal gasoline tax and the first-best Pigouvian emissions tax depends on three empirical relationships: the distribution of pollution externalities across vehicles; the extent to which gasoline prices affect the implicit demand for pollution; and the correlation between vehicle-specific demand responses and externality levels. If vehicles do not differ in their elasticity of vehicle miles traveled (VMT) with respect to gasoline prices—an elasticity that we hereafter call the VMT-elasticity—the SBO gasoline tax will simply be the average per-gallon externality across all vehicles. However, if price responsiveness and externalities are correlated, Diamond (1973) shows that the SBO gasoline tax will be a weighted average of vehicle per-gallon externalities, where the weights are the price derivatives of the vehicle-specific gasoline demand curves. In our empirical work, we allow for the VMT elasticity to vary depending on a vehicle's emissions per mile traveled, which we observe in our data.

An important empirical results in the paper is that we find that vehicle-level emissions are correlated with vehicle-specific VMT elasticities; dirtier vehicles are more price responsive.[2] Using detailed vehicle-specific data on miles driven, we show a positive correlation between criteria pollutant emissions and the VMT elasticity (in absolute value) holds for all three pollutants for which we have data: carbon monoxide (CO), hydrocarbons (HCs), and nitrogen oxides (NO_x). VMT elasticities are also positively correlated with greenhouse gas emissions and vehicle weight.

We find the average VMT elasticity is -0.13, but differences between cleaner and dirtier vehicles are substantial. When we allow VMT elasticity to vary by within-year quartiles of NO_x emissions, the elasticity for vehicles in the highest (i.e., dirtiest) quartile is -0.28. The VMT elasticity then falls monotonically with NO_x quartiles. The VMT elasticity is -0.15 in the third quartile; -0.05 in the second quartile, and 0.04 in the first quartile. Similar correlations between emissions and VMT elasticities hold for CO and HCs.

These correlations drive a wedge between the SBO gasoline tax associated with emissions and what we call the "naive" tax, which we define as the the tax based only on the *unweighted*-average externality across vehicles. We show the SBO gasoline tax is larger, on the order of 50 percent, than the naive gasoline tax in each of the years of our sample.

[2]While we tend to discuss the responsiveness of individuals to changes in prices, because drivers can shift miles from one car to another, the more relevant response is the response of the number of miles driven by a particular vehicle. Therefore, throughout we focus on the response of miles driven by a particular vehicle, not by a particular driver.

However, we show that even when instituting the SBO gasoline tax, the tax performs poorly in eliminating DWL, and only marginally better than the naive gasoline tax. Across our sample, we estimate the SBO gasoline tax eliminates only 30 percent of DWL associated with the pollutants studied. During the second half of our sample, the SBO gasoline tax eliminates only 25 percent of the DWL.

Given the shortfall of the SBO gasoline tax, we next investigate whether *any* indirect tax can eliminate a substantial portion of the DWL from emissions of criteria pollutants. We calculate efficiency improvements from county-specific gasoline taxes and from a policy that removes the dirtiest vehicles from the fleet. We find allowing gasoline taxes to vary by county would lead to a small improvement, eliminating DWL by less than an additional 5 percentage points. We find moderate benefits from "homogenizing" the fleet in terms of emissions per gallon, potentially through vehicle retirement (e.g., "Cash-for-Clunkers") programs; scrapping the dirtiest 10 percent of vehicles eliminates an additional 14 percentage points of DWL. Only a hypothetical tax linked to the weighted-average externality by vehicle age offers a substantial improvement over a simple uniform tax. Of course, such a tax is almost certainly politically and practically infeasible.

In addition to failing on efficiency grounds, the SBO gasoline tax also fails on equity grounds. Gasoline taxes are generally regressive (Chernick and Reschovski, 1997), because gasoline demand is income inelastic. Gasoline expenditures and gasoline taxes paid are relatively constant across the income distribution and necessarily make up an larger share of income at the lower end of the distribution. One might expect that an emissions tax might be more regressive still, especially if poorer households are more likely to own higher polluting vehicles. Our results show the opposite, however. Our empirical model predicts that the average household in *every* income decile would pay a higher percentage of income under the SBO gasoline tax than under an optimal emissions tax. This is more pronounced for households at the bottom of the income distribution. The intuition for this result is that the SBO tax is a weighted average of per-gallon emissions where the weights are the derivative of vehicle-specific VMT function with respect to gasoline prices. Because dirtier vehicles are more price responsive, the SBO tax is higher than unweighted average per-gallon emissions. Therefore, SBO tax revenues are higher than under a the optimal emissions tax. Furthermore, although on average lower income households are more likely to drive higher polluting vehicles, the correlation is weak. Vehicles that would have a higher tax burden under an optimal emissions tax are owned by households throughout the income distribution, and make up a small minority of vehicles in every income bracket.

We are not the first to analyze the optimal level of gasoline taxes. Parry and Small (2005) calculate the optimal gasoline tax for the US and UK accounting for local and global pollution, accidents, congestion, and the inefficiencies associated with income taxes. Our analysis differs in four key respects. First, Parry and Small (2005) implicitly assume vehicle externalities are uncorrelated with the sensitivity of each vehicle's demand for gasoline and gasoline prices, whereas we allow for, and find, such correlation. Second, we account for the possibility that marginal damage of vehicle emissions may vary geographically. Third, Parry and Small (2005) do not estimate the DWL that remains from instituting a gasoline tax, as opposed to the first best set of optimal emissions taxes, which is one of the main focuses of our paper. Fourth, our focus is more narrow in some respects, as we focus on externalities associated with local and global pollution and abstract from external costs related to accidents and congestion.

The closest paper to ours, in terms of our DWL results, is Fullerton and West (2010), who also investigate the amount of DWL eliminated by a uniform-gasoline tax. They do so by calibrating a numerical model with approximate miles and emissions obtained by matching inspection data from a small pilot study in California to quarterly gasoline expenditures in the Consumer Expenditure Survey. In contrast, our estimates are based on actual emissions, miles traveled, and gasoline prices for the universe of California vehicles. We find a uniform tax removes much less of the DWL of pollution compared to their calculations.[3]

The paper proceeds as follows. Section 2 draws on Diamond (1973) to derive the SBO gasoline tax and the amount of remaining DWL. Section 3 discusses the empirical setting and data. Section 4 provides descriptive support for the empirical results through graphical analysis. Section 5 presents the main empirical model and results on miles driven. Section 6 estimates empirically the optimal uniform tax and welfare effects, and Section 7 presents results on the incidence of gasoline and emissions taxes. Section 8 concludes the paper.

[3] In addition, there is a broad literature aimed at estimating how vehicle owners' driving and scrappage decisions respond to gasoline prices and vehicles policies, typically using either aggregate data or NHTS survey data. See for instance Li et al. (2009), Gillingham (2010) and Jacobsen and van Benthem (2013).

2 Optimal Uniform Taxes

In this section, we derive the second-best optimal uniform tax to internalize an externality, in the presence of heterogeneity in the externality. We closely follow the model of Diamond (1973) in deriving the optimal tax. We then add more structure to the problem to analytically solve for the amount of remaining DWL.

To simplify, we consider a setting where consumers own one vehicle, allowing us to treat consumers and vehicles interchangeably. We then discuss the implications of multiple vehicle households.

Consumer h derives utility (indirectly, of course) from her gasoline purchases, α_h, and a numeraire, μ_h, but is also affected by the gasoline consumption of others, α_{-h} (the externality). Assuming quasi-linear preferences, consumer h's utility can be written as:

$$U^h(\alpha_1, \alpha_2, ..., \alpha_h, ..., \alpha_n) + \mu_h. \tag{1}$$

We assume utility is monotone in own consumption, i.e., $\frac{\partial U^h}{\partial \alpha_h} \geq 0$. This yields demand curves α_h^*, given by:

$$\alpha_h^* = \alpha_h \left(p_g + \tau\right), \tag{2}$$

where p_g denotes the per-gallon price of gasoline, and τ a per-gallon tax on purchases of gasoline.

These assumptions, along with assuming an interior solution for each consumer, lead to:

Proposition 1. *The second-best-optimal uniform per-gallon gasoline tax, τ^*, is* (from Diamond (1973))*:*

$$\tau^* = \frac{-\sum_i \sum_{h\neq i} \frac{\partial U^h}{\partial \alpha_i} \alpha_i'}{\sum_i \alpha_i'}. \tag{3}$$

where $\sum_{h\neq i} \frac{\partial U^h}{\partial \alpha_i}$ is the external costs associated with a gallon of gasoline consumed by vehicle i and α_i' is the derivative of consumer i's demand for gasoline with respect to the price of gasoline.

Proof. See Appendix A. □

The SBO gasoline tax becomes a weighted average of vehicles' externalities where the weights are the derivative of the demand for the gasoline with respect to the tax. If there is a positive correlation between price responsiveness and emissions (i.e., dirtier cars are more responsive), this correlation will increase the SBO gasoline tax.[4]

As Diamond explicitly discusses, there is no requirement that all of the α_i's must be negative.[5] Indeed, if households hold multiple vehicles, they may shift miles from their least-fuel efficient, most polluting vehicle to a cleaner, more efficient vehicle.[6]

The presence of heterogeneity in the externality also implies that a uniform tax will not achieve the first-best outcome. The uniform tax will under-tax high externality agents and over-tax low externality agents. We extend Diamond (1973) to solve for the amount of DWL remaining in the presence of an SBO gasoline tax applied to a market with heterogenous externalities. Doing so requires a bit more structure. We first start with the case where VMT elasticities and emissions are uncorrelated and then relax this assumption.

Proposition 2. *Suppose drivers are homogenous in their demand for gasoline, but vehicles' per-gallon emissions differ. In particular, let α' denote the derivative of the demand for gasoline with respect to the price of gasoline.*

If the distribution of the per-gallon externality, E, is log normal, with probability density function:

$$\varphi(E_i) = \frac{1}{E_i\sqrt{2\sigma_E^2}} \exp\left(\frac{-(E_i - \mu_E)^2}{2\sigma_E^2}\right), \tag{4}$$

the DWL absent any market intervention will be given as:

$$D = \frac{1}{2\alpha'} e^{2\mu_E + 2\sigma_E^2}.$$

Proof. See Appendix A. □

This leads to the following calculation of remaining DWL under the SBO gasoline tax.

Proposition 3. *Under the assumptions in Proposition 2, the ratio of remaining DWL after*

[4] As an intuitive example, imagine the case where there are only two vehicle types. The first emits little pollution, while the second is dirtier. Also imagine the clean vehicles are completely price insensitive, while the dirty vehicles are price sensitive. The naive tax would be calculated based on the average emissions of the two vehicle types. However, the marginal emission is the emissions rate of the dirty vehicles; the clean cars are driven regardless of the tax level. In this case, we can achieve first best by setting the tax rate at the externality rate of the dirty vehicle. There is no distortion to owners of clean vehicles since their demand is completely inelastic, so we can completely internalize the externality to those driving the dirty vehicles.

[5] However, second-best optimal tax loses the interpretation as a weighted average if some α_i's are positive.

[6] Also note that the elasticity of gasoline consumption with respect to price accounts for households selling or scrapping their vehicles and buying different ones. That is, if the gasoline tax increases the scrappage rate of some vehicles, then the relevant derivative of the externality with respect to price is the expected change in gasoline consumption, not the change in gasoline demand, conditional on survival.

the tax is imposed to the DWL absent the tax is:

$$R = \frac{D - \frac{e^{2\mu_E + \sigma_E^2}}{2\alpha'}}{D} = 1 - \frac{e^{2\mu_E + \sigma_E^2}}{e^{2\mu_E + 2\sigma_E^2}} = 1 - e^{-\sigma_E^2}. \tag{5}$$

Proof. See Appendix A. □

With externalities uncorrelated with the demand for gasoline, the remaining DWL from a uniform tax depends only on the shape parameter of the externality distribution. The larger σ_E^2 is, the wider and more skewed will the distribution of the externality be, causing the uniform tax to "overshoot" the optimal reduction in gasoline consumption for more vehicles.

If the demand for gasoline is not homogeneous, and in fact is correlated with per-gallon externalities, the calculation changes. Let α_h' denote the derivative of the demand for gasoline associated with vehicle i with respect to the price of gasoline. For ease, define $B_i = \frac{1}{\alpha_h'}$, and assume that B_h is distributed lognormal with parameters μ_B and σ_B^2. Define ρ as the dependence parameter of the bivariate lognormal distribution (the correlation coefficient of $\ln E$ and $\ln B$). We then have:

Proposition 4. *When B_h and E_h are distributed lognormal with dependence parameter ρ, the optimal tax is:*

$$\tau^* = e^{\mu_E + \frac{\sigma_E^2}{2} + \rho\sigma_E\sigma_B}$$

Proof. See Appendix A. □

As we would expect, the optimal tax does not depend on the scale of the elasticity distribution, only on the extent to which externalities are correlated with elasticities. We can then calculate the amount of remaining DWL under both the naive and SBO gasoline tax.

Proposition 5. *When B_i and E_i are distributed lognormal with dependence parameter ρ, the ratios of the remaining DWL after the SBO gasoline tax to the original DWL will be:*

$$R(\tau^*) = 1 - e^{-\sigma_E^2}, \tag{6}$$

And, the ratios of the remaining DWL after the naive uniform tax to the original DWL will be:

$$R(\tau_{naive}) = 1 - e^{-\sigma_E^2}(2e^{-\rho\sigma_E\sigma_B} - e^{-2\rho\sigma_E\sigma_B}). \tag{7}$$

Proof. See Appendix A. □

As we would expect, the optimal tax correctly accounts for the correlation between the externality and demand responses, and thus the remaining DWL depends only on the variance and skewness of the externality distribution. However, in the presence of correlation the naive tax reduces less of the DWL from the externality, reducing it by a proportion related to the degree of correlation and the spread of the two distributions. The term in parentheses in Equation (7) is strictly less than 1, and strictly greater than zero if $\rho > 0$, but may be negative if $\rho < 0$ and the shape parameters are sufficiently large.

In Section 6, we will show that σ_E^2 is such that $R(\tau^*)$ is surprisingly large, and that while $R(\tau_{naive})$ is measurably larger, it is not much larger.

3 Empirical Setting

3.1 Data

Our empirical setting is the California personal transportation market. We bring together a number of large data sets. Our analysis is primarily based upon the universe of emissions inspections from 1996 to 2010 from California's vehicle emissions testing program, the Smog Check Program, which is administered by the California Bureau of Automotive Repair (BAR). An vehicle appears in the data for a number of reasons. First, vehicles more than four years old must pass a Smog Check within 90 days of any change in ownership. Second, in parts of the state (details below) an emissions inspection is required every other year as a pre-requisite for renewing the registration on a vehicle that is six years or older. Third, a test is required if a vehicle moves to California from out-of-state. Vehicles that fail an inspection must be repaired and receive another inspection before they can be registered and driven in the state. There is also a group of exempt vehicles. These are: vehicles of 1975 model-year or older, hybrid and electric vehicles, motorcycles, diesel-powered vehicles, and large natural-gas powered trucks.

These data report the location of the test, the unique vehicle identification number (VIN), odometer reading, the reason for the test, and test results. We decode the VIN to obtain the vehicles' make, model, engine, and transmission. Using this information, we match the vehicles to EPA data on fuel economy. Because the VIN decoding is only feasible for vehicles made after 1981, our data are restricted to these models. We also restrict our sample to 1998 and beyond, given large changes that occurred in the Smog Check Program in 1997. This yields roughly 120 million observations.

The Smog Check data report measurements for NO_x and HCs in terms of parts per million and CO levels as a percentage of the exhaust, taken under two engine speeds. As we are interested in the quantity of emissions, the more relevant metric is a vehicle's emissions per mile. We convert the Smog Check emissions readings into emissions per mile using conversion equations developed by Sierra Research for the California Air Resources Board in Morrow and Runkle (2005). The conversion equations are functions of both measurements of all three pollutants, vehicle weight, model year, and truck status. For more details on cleaning the Smog Check data, including the conversion equations for the three criteria pollutants, see online Appendix B.

As part of our simulation exercise, we also use data obtained from CARFAX Inc. to estimate scrappage decisions. These data contain the date and location of the last record of the vehicle reported to CARFAX for 32 million vehicles in the Smog Check data. This includes registrations, emissions inspections, repairs, import/export records, and accidents. Because the CARFAX data include import/export records, we are able to correctly classify the outcomes of vehicles which are exported to Mexico as censored, rather than scrapped, thus avoiding the issues identified in Davis and Kahn (2010).

For a subset of our Smog Check data, we are able to match vehicles to households using confidential data from the California Department of Motor Vehicles (DMV). These data track the registered address of the every vehicle in the state, with one address given for each year. We use the registration information to attach demographic information on income from U.S. Census data. Appendix C discusses the process of cleaning the registration data. The DMV data are only available for the years 2000 to 2008.

For a portion of our analysis, we use data from the 2009 National Highway Transportion Survey, which contains information on household vehicles, annual VMT, and household income for a sample of households.

Finally, we use gasoline prices from EIA's weekly California average price series to construct average prices between inspections.

Table 1 reports means and standard deviations of the main variables used in our analysis, for all observations and broken down by vehicle age and by year of Smog Check. The average fuel economy of vehicles in our sample is 23.5 MPG, with fuel economy falling over the period of the sample. The change in the average dollars per mile has been dramatic, almost tripling between 1998 and 2008. The dramatic decrease in vehicle emissions is also clear in the data, with average per-mile emissions of HCs, CO, and NO_x falling considerably from 1998 to 2008. The tightening of standards has also meant that more vehicles fail Smog Checks late in the sample, although some of this is driven by the aging of the vehicle fleet.

3.2 Automobiles, Criteria Pollutants, and Health

The vehicle inspection data report emissions of three criteria pollutants: NO_x, HCs, and CO. All three of these directly result from the combustion process within either gasoline or diesel engines. Both NO_x and HCs are precursors to ground-level ozone, but all three have been shown to have negative health effects on their own.[7]

While numerous studies have found links between exposure to the ozone or the three criteria pollutants and health outcomes, the mechanisms are still uncertain. These pollutants, as well as ozone, may directly impact vital organs or indirectly cause trauma. For example, CO can bind to hemoglobin, thereby decreasing the amount of oxygen in the bloodstream. High levels of CO have also been linked to heart and respiratory problems. NO_x reacts with other compounds to create nitrate aerosols, which are fine-particle particulate matter (PM). PM has been shown to irritate lung tissue, lower lung capacity, and hinder long-term lung development. Extremely small PM can be absorbed through the lung tissue and cause damage on the cellular level. On their own, HCs can interfere with oxygen intake and irritate lungs. Ground-level ozone is a known lung irritant, has been associated with lowered lung capacity, and can exacerbate existing heart problems and lung ailments such as asthma or allergies.

4 Preliminary Evidence

One of the main driving forces behind our empirical results is how gasoline demand elasticities for different vehicles vary systematically emissions levels. In this section, we present evidence that significant variation exists in terms of vehicle externalities within a year, across years,

[7] CO has also been shown to speed up the smog-formation process. For early work on this, see Westberg et al. (1971).

and even within the same vehicle type (make, model, model year, etc.) within a year. Further, simple statistics, such as the average miles traveled by vehicle type, suggest that elasticities may be correlated with externalities.[8]

Figure 1 plots the distributions of NO_x, HCs, and CO emissions in 1998, 2004, and 2010. The distribution of criteria pollutant emissions tends to be right-skewed in any given year, with a standard deviation equal to roughly one to three times the mean, depending on the pollutant. The skewness implies that some vehicles on the road are quite dirty relative to the mean vehicle. Over time, the distribution has shifted to the left, as vehicles have gotten cleaner, but the range remains.

The variation in emissions is not only driven by the fact that different types of vehicles are on the road in a given year, but also variation within the *same* vehicle type, defined as a make, model, model-year, engine, number of doors, and drivetrain combination. To see this, Figure 2 plots the distributions of emissions for the most popular vehicle/year in our sample, the 2001 four-door Toyota Corolla in 2009. The vertical red line is at the mean of the distribution. Here, again, we see that even within the same vehicle-type in the same year, the distribution is wide and right-skewed. The distribution of HCs is less skewed, but the standard deviation is 25 percent of the mean. CO is also less skewed and has a standard deviation that is 36 percent of the mean. Across all years and vehicles, the mean emissions rate of a given vehicle in a given year, on average, is roughly four times the standard deviation for all three pollutants (Table A.1).

To understand how the distribution within a given vehicle type changes over time, Figure 3 plots the distribution of the 1995 3.8L, front-wheel drive, Ford Windstar in 1999, 2001, 2004, and 2007.[9] These figures suggest that over time the distributions shift to the right, become more symmetric, and the standard deviation grows considerably, relative to the mean. Across all vehicles, the ratio of the mean emission rate of NO_x and the standard deviation of NO_x has increased from 3.16 in 1998 to 4.53 in 2010. For HCs, this increased from 3.59 to 5.51; and, for CO the ratio increased from 3.95 to 5.72.

These distributions demonstrate significant variation in emissions across vehicles and

[8] We are not the first to document the large variation across vehicles in emissions. See, for example, Kahn (1996). Instead, our contribution is in finding a link between elasticities and emissions.

[9] We chose this vehicle because the 1995 3.8L, front-wheel drive, Ford Windstar in 1999 is the second-most popular entry in our data and it is old enough that we can track it over four 2-year periods.

within vehicle type, and thus significant scope for meaningful emissions-correlated variation in elasticities along those lines.[10] We next present suggestive evidence that VMT elasticities may be correlated with emissions. For each criteria pollutant within each calendar year we rank vehicles by their observed emissions per mile and divide them into quartiles. We do the same for fuel economy. Next, for each quartile-year we compute the median annual VMT, and plot how this has changed over our sample, normalized by the 1998 level for each quartile.. Figure 4 foreshadows our results on VMT elasticities and externalities. For each pollutant, we see that the dirtiest quartile saw the largest decreases in miles driven during the run up gasoline prices from 1998 to 2008, when prices increased from roughly \$1.35 to \$3.20.[11] The ordering of the relative decreases suggests that dirtier vehicles may have been more responsive over this period.

5 Vehicle Miles Traveled Decisions

We now estimate how changes in gasoline prices affect decisions about vehicle miles traveled (VMT), and how this elasticity varies with vehicle characteristics. Our empirical approach mirrors Figure 4. For each vehicle receiving a biennial smog check, we calculate average daily miles driven and the average gasoline price during the roughly two years between Smog Checks. Obviously vehicle owners with more fuel efficient vehicles will respond less to changes in the per-gallon gasoline price, and to abstract from this we specify the elasticity with respect to the price in dollars per mile (DPM), by dividing the average per gallon price by fuel economy in gallons per mile. Thus, the price faced by each vehicle's owner will vary both with the exact period in between Smog Checks, and with the specific vehicles' fuel economy. We then allow the elasticity to vary based on the emissions of the vehicle. We begin by estimating:

$$\ln(VMT_{ijgt}) = \beta \ln(DPM_{ijgt}) + \gamma D_{truck} + \omega time + \mu_t + \mu_j + \mu_g + \mu_v + \varepsilon_{igt} \quad (8)$$

[10] Because of the way we handle multiple tests of a given vehicle with a year, our distributions likely understate the degree of on-road heterogeneity. In order for a vehicle to be registered, the vehicle must pass a Smog Check. In our data we see multiple tests of the same vehicle over a short time frame. We use the final test, which will necessarily have been passed, for our calculations. Furthermore, our calculations omit unregistered vehicles, many of which are likely to have high emissions.

[11] The levels also differ. Appendix Figure A.1 plots the median of daily miles traveled across our sample split up by the emissions quartile of the vehicle.

where i indexes vehicles, j vehicle-types, g geographic locations, t time, and v vehicle age, or vintage. DPM_{ijgt} is the average gasoline price per mile faced by vehicle i between time t and the date of the previous Smog Check, D_{truck} is an indicator variable for whether the vehicle is a truck, $time$ is a time trend, and ε_{igt} is a residual.[12] Our baseline specification assumes that gasoline prices are exogenous to individual driving decisions. Such an assumption is common in the literature, as gasoline prices are largely driven by movements in the world price of crude oil, which saw dramatic changes during the 2000s for reasons unrelated to driving choices in California.[13] However, we have also estimated our main analyses instrumenting for DPM with the Brent Crude oil price, and we obtain very similar results. In Online Appendix D we estimate equation 8 including a complete set of month-by-year fixed effects, thus relying on cross-sectional variation in gas prices, as opposed to time series variation. The results are qualitatively similar to our baseline specification.

We begin by including demographic characteristics by the zip code of Smog Checks, and year and vintage fixed effects. We then progressively include finer vehicle-type fixed effects by including make, then make/model/model-year/engine, and finally individual vehicle fixed effects. We also differentiate the influence of gasoline prices by vehicle attributes related to the magnitude of their negative externalities—criteria pollutants, CO_2 emissions, and weight.

We allow the VMT elasticity to vary with the magnitude of their externalities in two ways. For both approaches, we begin by ranking vehicles within each calendar year by their emissions per mile of NO_x, HCs, CO, fuel economy, or vehicle weight in pounds. In one set of specifications we split vehicles up by the quartile of these variables and allow each quartile to have a separate β. In another set, we include a linear interaction of centiles of these variables and the log of gasoline prices in dollars per mile.

Table 2 shows our results, focusing on NO_x. The changes from Models 1 to 4 illustrates the importance of controlling for vehicle-type fixed effects. Initially, the average elasticity falls from -0.265 to -0.117 when including make fixed effects, but then rises when including finer vehicle type fixed effects. Model 4 includes individual vehicle fixed effects yielding an

[12]The fuel economy in gallons per mile used to calculate our DPM variable uses the standard assumption that 45 percent of a vehicle's miles driven are in the city and 55 percent are on the highway. This is the standard approach used by the EPA for combined fuel economy ratings.

[13]See, for example, Busse et al. (forthcoming).

average elasticity of -0.134.[14] In Models 5 and 6 we examine heterogeneity with vehicle fixed effects. Model 5 includes interactions with quartiles of NO_x. The VMT elasticity for the cleanest vehicles, quartile one, is positive at 0.043, while the VMT elasticity for the dirtiest vehicles is twice the average elasticity at -0.280. To put these numbers in context, the average per-mile NO_x emissions of a quartile one vehicle is 0.163 grams, while the average per-mile NO_x emissions of a quartile four vehicle is 1.68 grams. Model 6 assumes the relationship is linear in centiles of NO_x and finds that each percentile increase in the per-mile NO_x emission rate is associated with an elasticity .001 larger in absolute value, from a base of essentially zero.

We find similar patterns across the other externalities. The range of the estimated VMT elasticities is somewhat larger when using quartiles of HCs and CO emissions compared to NO_x, with the dirtiest quartiles around -0.30 and the cleanest around 0.05. For CO_2 the cleanest vehicles are those with the highest fuel economy, and here we see the least fuel-efficient vehicles having a VMT elasticity of -0.183, compared to -0.108 for vehicles with fuel economy in the highest quartile. We observe some heterogeneity in the VMT elasticity across vehicle weights as well, although it is smaller than the other externalities. For the full set of results, see Appendix Table A.2.

6 Efficiency of the second best optimal gasoline tax

In this section, we consider the efficiency of using an SBO gasoline tax to abate the externalities caused by driving, specifically those resulting from emissions of NO_x, HCs, and CO. We begin by calculating both the naive and SBO gasoline tax, and then compare the remaining DWL left over from these second-best taxes to the optimal outcome obtained by a Pigouvian tax on emissions.

6.1 Second best optimal gasoline tax

We calculate the naive tax per gallon of gasoline as the simple average of the externality per gallon caused by all vehicles on the road in California in a particular year. We value the externalities imposed by NO_x and HCs using the marginal damages calculated by Muller

[14] Our average elasticity is larger than that found in Hughes et al. (2008) reflecting the longer run nature of our elasticity.

and Mendelsohn (2009), based on the county in which each vehicle has its Smog Check.[15] The damages calculated by Muller and Mendelsohn (2009) are ideal for this purpose, as they use an integrated assessment model to capture how a marginal unit of NO_x or HCs emitted in one location causes damages throughout the United States, both directly and through the formation and removal of ozone and particulate matter. For CO, we use the median marginal damage estimate from Matthews and Lave (2000).

Let the marginal damage per gram of pollutant p in county c be θ_c^p, with emissions rates in grams per mile by vehicle i of ϵ_i^p. Then the externality per mile of vehicle i, E_i is:

$$E_i = \theta_c^{HC} \cdot \epsilon_i^{HC} + \theta_c^{NO_x} \cdot \epsilon_i^{NO_x} + \theta_c^{CO} \cdot \epsilon_i^{CO} \tag{9}$$

The naive tax in year y will then be:

$$\tau_{naive}(y) = \frac{1}{N^y} \sum_{i=1}^{N^y} \frac{E_i}{MPG_i}, \tag{10}$$

where N^y denotes the number of vehicles on the road in year y, and MPG_i denotes the fuel economy rating of vehicle i. In practice, since the stock of vehicles represented in the Smog Check data in any given year will be less than the total stock of vehicles in the vehicles fleet, we weight each Smog Check observation by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole.

Following Proposition 1, we calculate the SBO gasoline tax, taking into account the heterogeneity in both levels of the externality and the responsiveness to gasoline prices. We estimate a regression similar to Equation (8), but allowing the elasticity of VMT with respect to DPM to vary over quartiles of all three criteria pollutants, fuel economy, vehicle weight, and three groups of vehicle age. For more details, see Appendix E. Let the group-specific elasticity for vehicle i be β_i^q, where q indexes cells by HC emissions, NO_x emissions, CO emissions, MPG, weight, and age, with the externalities again in quartiles by year. Further, let the average price per gallon and the quantity of gasoline consumed per year in gallons in

[15] Note that the values used in this paper differ from those used in the published version of Muller and Mendelsohn (2009). The published values were calculated using incorrect baseline mortality numbers that were too low for older age groups. Using corrected mortality data increases the marginal damages substantially. We are grateful to Nicholas Muller for providing updated values, and to Joel Wiles for bringing this to our attention.

year y be P_i^y and Q_i^y, respectively.[16] Then the optimal tax in year y will be

$$\tau^*(y) = \frac{-\sum_h \sum_{i \neq h} \frac{\partial U^h}{\partial \alpha_i} \alpha_i'}{\sum_h \alpha_h'}, \tag{11}$$

with

$$\alpha_i' = -\beta_i^q \cdot \frac{Q_i^y}{P_i^y}, \tag{12}$$

and

$$\frac{\partial U^h}{\partial \alpha_i} = \frac{E_i}{MPG_i}. \tag{13}$$

Table 3 shows the naive and SBO taxes for each year from 1998 to 2008. The naive tax would be 61.2 cents per gallon of gasoline consumed in 1998, while the SBO tax is 86 cents, 39 percent higher. The ratio of the naive and SBO gasoline tax increases even as the levels of the externalities decline over time. From 2002 on, the SBO gasoline tax is at least 50 percent larger than the naive tax in each year.

6.2 Welfare with Uniform Taxes

We have shown that because of the correlation between elasticities and externality rates, the SBO gasoline tax is much higher than the naive tax calculated as the average of per-gallon externalities. We now turn to the question of how much the SBO gasoline tax improves welfare beyond what is achieved by the naive tax. We note again that even the optimal uniform tax is still a second-best policy. Because of the heterogeneity in externality levels, the most polluting vehicles will be taxed by less than their external costs to society, leaving remaining DWL. Vehicles that are cleaner than the weighted average will be taxed too much, overshooting the optimal quantity of consumption and creating more DWL.[17]

In each of the following analyses, we compare the remaining DWL resulting from the local pollution externality with both the naive and SBO gasoline tax to the DWL without

[16] Again, we also weight vehicles based on the number of vehicles of that age and class that appear in the fleet as a whole; see Appendix E. We also account for vehicle owners' decisions to scrap their vehicles to the extent these are affected by gasoline prices. Appendix G discusses the details and results of this exercise. To summarize, we allow gasoline price to affect scrappage decisions, and allow this to vary over emissions profiles and vintages. We find that the main source of heterogeneity occurs across vintages; specifically, increases in gasoline prices increase the hazard rate of very old vehicles, but decrease the hazard rate of middle-aged vehicles. Because emissions of criteria pollutants are positively correlated with age, this has the effect of decreasing criteria pollutants, although the aggregate effect is small.

[17] We have also repeated the analysis under the assumption that policy makers adopt the second-best optimal VMT tax. The degree of DWL that remains is only slightly reduced.

any additional tax.

6.2.1 Simulation Results

We begin by approximating the ratios of DWL with and without the taxes using our data to simulate the change in miles driven and thus in gasoline consumption from a tax. Let $miles_i^y$ be the actual average miles per day traveled by vehicle i between its last Smog Check and the current one, observed in year y, and let $\hat{miles}_i^y(\tau)$ be the miles per day that a vehicle would travel if the average price of gasoline were raised by a tax of $\$\tau$ per gallon that is fully passed through to consumers. We approximate DWL as a triangle, such that the ratio of interest is:

$$r(\tau) = \frac{\sum_i \frac{1}{2} \cdot \frac{miles_i^y - \hat{miles}_i^y(\tau)}{MPG_i} \cdot \frac{E_i}{MPG_i} - \tau}{\sum_i \frac{1}{2} \cdot \frac{miles_i^y - \hat{miles}_i^y(\frac{E_i}{MPG_i})}{MPG_i} \cdot \frac{E_i}{MPG_i}}$$

The fully optimal tax would have a ratio of 0, while a tax that actually increased the DWL from gasoline consumption would be greater than 1. Table 4 shows these ratios for various taxes. The first two columns show ratios for a statewide tax based on the average and marginal externalities (i.e. the naive and SBO taxes), respectively, of all vehicles in California in each year. DWL with the naive tax averages 72.8 percent of the DWL with no additional tax over the sample period, and rises over time as the fleet becomes cleaner. The SBO gasoline tax is little better, averaging 69.8 percent of DWL with no tax during our sample period.

Is it even possible to effectively abate local pollution externalities using a tax on gasoline, or is there too much idiosyncratic variation in externality levels for this to be possible? That is, if hypothetically the tax were allowed to vary by groups observable to policymakers, would the SBO uniform tax perform better? Obviously, this may be politically infeasible depending on how the groups are defined and may be impractical to implement. The purpose of this analysis is to explore the nature of the failure of the uniform gasoline taxes.

The remaining columns of Table 4 show remaining deadweight loss from the naive and SBO taxes using the average or marginal externality for specific groups of vehicles, rather than the entire state. The marginal damages from Muller and Mendelsohn (2009) are designed to vary at the county level, and within California they vary substantially across counties, due to both baseline emissions levels and the extent to which population is exposed to harmful emissions.[18] As such, a county-specific tax on emissions might be expected to target externality levels more precisely. The third and fourth columns of Table 4 shows the DWL ratios for a naive and SBO gasoline tax computed this way, and it turns out there is relatively little improvement. In other words, county-by-county variation in emissions and elasticities does not explain the failure of a single, uniform tax to remove a substantial amount of deadweight loss. The average ratio over our sample is 0.684 for the naive tax and 0.653 for the optimal uniform tax.

Since emissions rates are highly correlated with vintage, another approach would be to allow taxes to vary by the age of the vehicle.[19] The fifth and sixth columns of table 4 show this, and here we see a substantial improvement: 0.342 for the naive tax and 0.34 for a SBO gasoline tax. Combining these and having the tax vary by both vintage and location, shown in the last two columns, reduces the ratios to 0.276 and 0.274, respectively.

This analysis shows two striking results. First, an SBO gasoline tax does a terrible job of addressing the market failure from local pollution externalities. The dirtiest vehicles are not taxed enough, and many clean vehicles are over-taxed. This is true even when the uniform tax is calculated taking the correlation between emissions and VMT elasticities into account. The roughly 50 percent increase in the tax level from a SBO gasoline tax correctly abates more emissions from the dirtiest vehicles, but also over-taxes the cleanest vehicles by a larger amount. The welfare benefits of the SBO gasoline tax are around 10% higher than those from a naive tax, but still fall far short of the benefits from a true Pigouvian tax linked to actual vehicle emissions. The number of vehicles for which the uniform tax overshoots is remarkable. Table 5 shows the proportion of vehicle-years over the 11 years of our sample for which each tax overshoots. Because the distribution of emissions is so strongly right skewed, the naive uniform tax overshoots for more than 72 percent of vehicle-years, and the SBO gasoline tax for even more. Second, there is enough heterogeneity in the distribution of the per-gallon externality that even a tax targeting broad groups leaves a substantial portion of DWL.

[18] We discuss this further in Online Appendix F.

[19] Such a system could perhaps be built within the Smog Check Program, with vehicle-specific taxes based on mileage since the previous test.

The variance and skewness in the distribution of externality per gallon causes a uniform tax to be less efficient than might otherwise be expected. Figure 5 shows this clearly, plotting the kernel density of the externality per gallon in 1998 and 2008, with vertical lines indicating the naive tax and the optimal tax, respectively. The long right tail of the distribution requires that either tax greatly exceed the median externality.

We next examine how the optimal uniform tax would compare to the optimal Pigouvian emissions tax if the distribution became less skewed. That is, how would a uniform tax perform if the right tail of the distribution—the oldest, dirtiest vehicles—were removed from the road? This could be achieved directly from a Cash for Clunkers-style program, or indirectly through tightening emissions standards in the Smog Check Program. Sandler (2012) shows that vehicle retirement programs are not cost-effective in reducing criteria emissions, and possibly grossly over pay for emissions. However, the overall welfare consequences of this sort of scheme may be more favorable if they improve the efficiency of a uniform gasoline tax. Table 6 shows the ratios of DWL with the SBO gasoline tax to DWL with no tax, removing increasing proportions of the top of the externality distribution. Removing the top 1 percent increases the DWL reduction from 30 percent to 38 percent of the total with no tax. Scrapping more of the top end of the distribution improves the outcome further. If the most polluting 25 percent of vehicles were removed from the road and the SBO gasoline tax was imposed based on the weighted externality of the remaining 75 percent, this would remove 58.3 percent of remaining DWL. Of course, the practical complications of scrapping this large a proportion of the vehicle fleet likely make this cost-prohibitive.

6.2.2 Analytical Results

We can also calculate the ratio of remaining DWL to original DWL by calibrating Equations (6) and (7) and with the moments in our data. The average value in our sample for the lognormal shape parameters σ_E^2 and σ_B^2 are 1.47 and 1.51, respectively. The average value of ρ, the correlation coefficient for the logs of externality and inverse elasticity, is 0.28.[20] These parameter values produce remaining DWL estimates in line with the simulation results in Table 4. With σ_E^2 around 1.47, the optimal uniform tax can only decrease DWL by 23

[20]This is the average of parameters calculated separately for each year from 1998 to 2008. The parameters do not vary much over time. For the year-by-year parameter estimates, see Table A.7 in the Online Appendix.

percent.

6.3 Treatment of Other Externalities

In the previous section we assumed that the difference between the socially optimal consumption of gasoline and actual consumption was entirely driven by externalities from local pollution. In practice, there are several other externalities from automobiles, as well as existing federal and state taxes on gasoline. Examples of additional externalities include congestion, accidents, infrastructure depreciation, and other forms of pollution. The combined state and federal gasoline tax in California was \$0.47 during our sample period.

Many of these other externalities are similar to criteria pollution emissions in the sense that they also vary across vehicles. Congestion and accident externalities depend on when and where vehicles are driven. Accident and infrastructure depreciation depend to some degree on vehicle weight.[21] We lack vehicle-specific measures of these other externalities to measure how they impact our calculations of the amount of remaining DWL after imposing a SBO gasoline tax. However, because these other externalities also vary by vehicle, a gasoline tax will also be an imperfect policy instrument for these externalities. Therefore, the actual amount of deadweight loss will be the sum of the deadweight loss that we measure plus the deadweight loss arising from the externalities that we cannot measure. Insofar as additional variation exists, and they are not negatively correlated with the externalities that we do measure, we are understating the *level* of remaining DWL, although not necessarily the *share* of remaining DWL.[22] One way to interpret our results is that by ignoring the other externalities we are assuming that existing taxes exactly equal the SBO gasoline tax associated with these other externalities, and that we are also ignoring the remaining DWL due to the fact that these externalities are not uniform across vehicles.

One externality that does not vary across vehicles is the social cost of CO_2 emissions due to their contribution to climate change. Because CO_2 emissions are, to a first-order approximation, directly proportional to gasoline consumption, in this case a per-gallon gasoline tax

[21] For estimates on the degree of this heterogeneity, see Anderson and Auffhammer (2011) and Jacobsen (Forthcoming).

[22] In fact, Parry and Small (2005) find that the contribution of these other externalities to the second-best optimal gas tax may, in fact, be larger than the contribution arising from local pollutants. This would suggest that the degree for which we understate the remaining deadweight loss might be large.

is the optimal policy instrument. The larger the climate change externality, the greater the *share* of DWL eliminated by the SBO gasoline tax will be. To get a sense of how climate change externalities affect our calculations, we repeat the analysis for a range of social costs of carbon (SCC).

We calculate the remaining DWL after imposing the SBO gasoline tax based on local pollution externalities, varying the SCC from zero cents per gallon to $1.00 per gallon.[23] While our discussion focuses on the externalities associated with CO_2, we stress that these calculations are relevant for *any* externalities for which a per-gallon tax is the first-best instrument. They also represent the lower bound on the remaining DWL when we consider any other externality for which a per-gallon tax is a second-best instrument.

Figure 6 summarizes the results across all years in our sample. The values for an extra per-gallon externality of zero roughly correspond to the ratios reported in Table 4.[24] Not until the extra per-gallon externality exceeds $0.20 per gallon does a uniform gasoline tax eliminate the majority of DWL associated with both the criteria pollutants and a per-gallon externality. Even if the per-gallon externality is $1.00, nearly 20 percent of combined DWL remains under both the naive and SBO gasoline taxes.

7 Incidence of Gasoline and Pigouvian Taxes

Our results in section 6.1 demonstrate that a uniform gasoline tax is an ineffective policy tool on efficiency grounds. In this section we consider the implications of the SBO gasoline tax and the first-best optimal Pigouvian emissions tax for equity. Gasoline taxes are generally regressive (Chernick and Reschovski, 1997). However, it is possible that a uniform gasoline tax is less regressive than an emissions tax, particularly if poorer households tend to own dirtier vehicles. We begin by describing our methodology for assigning household income to the vehicles in our Smog Check data, and then present our results on the regressivity of the

[23]For comparison, Greenstone et al. (2011) estimate the SCC for a variety of assumptions about the discount rate, relationship between emissions and temperatures, and models of economic activity. For each of their sets of assumptions, they compute the *global* SCC; focusing only on the US impacts would reduce the number considerably. For 2010, using a 3 percent discount rate, they find an average SCC of $21.40 per ton of CO_2 or roughly 23.5 cents per gallon of gasoline, with a 95th percentile of $64.90 (71 cents per gallon).These calculations assume that the lifecycle emissions of gasoline are 22 pounds per gallon. Using a 2.5 percent discount rate, the average SCC is $35.10 (38.6 cents per gallon).

[24]The figure plots the weighted averages across the years, while the last row in Table 4 is a simple average of the annual weighted averages, hence a slight difference

SBO gasoline tax and the optimal Pigouvian emissions tax.

7.1 Assigning Income to Vehicles

For the analyses in this section, we link the Smog Check data to DMV registration information. We geocode the addresses from the DMV data and match them to Census block groups (CBGs) then link these data to CBG demographics from the 2000 Decennial Census. The DMV data allow us to match an address to each vehicle by calendar year for the period of 2000–2008. We then predict the annual average tax paid by the owners of each vehicle in the Smog Check data, using our estimates of optimal taxes and counterfactual VMT from section 6. All of the analyses in this section consider the average of predicted taxes over the whole period from 2000–2008. Although the levels of tax and externality are higher earlier in the period than later in the period, the distributional patterns are almost identical regardless of the time period used.

Figure 7 shows a map of California CBGs, shaded to show the average predicted annual tax paid as a fraction of year 2000 CBG median income. Panel (7a) shows the average tax burden of the SBO tax, while panel (7b) shows the average tax burden of an optimal Pigouvian emissions tax. Two patterns are evident from the maps. First, there is substantial geographic dispersion across the state of California, with vehicle owners in the urban cores of Los Angeles and San Diego paying a much higher proportion of CBG median income in taxes. Second, although the geographic dispersion differs slightly between the SBO gasoline tax and the optimal emissions tax, the *levels* differ substantially—the annual burden of the optimal tax is much lower than that of the SBO tax. We discuss the intuition behind this surprising result below.

The map analysis in Figure 7 is instructive, but is not enough to show the extent to which the SBO gasoline tax or an optimal Pigouvian emissions tax would be progressive or regressive. To estimate the progressivity of these taxes, median CBG income is insufficient. Borenstein (2012) shows that CBG median income masks substantial within-CBG variation in household income, which causes it to do a poor job of capturing effects on the top and bottom of the income distribution. We implement Borenstein's suggested correction for CBG income, which uses the full distribution of household incomes in each CBG from the Census data combined with a separate dataset that contains both annual VMT and household income. In brief, Borenstein's method requires the correlation between annual VMT and income within a CBG. In short, using the 2009 NHTS, we calculate the correlation between VMT and income. This allows us to assign vehicles in the Smog Check data to income brackets based on each vehicle's annual VMT and the proportion of households in each income bracket within the CBG the vehicle is registered in. For more details, see Appendix H and Borenstein (2012).

Using the Borenstein (2012) correction we assign vehicles to one of 10 income brackets, which aggregate the 16 income categories contained in the Census data into groups roughly approximating deciles of the California household income distribution.[25] For purposes of calculating the tax burden and progressivity, we use the midpoint of each income bracket. Because our data is at the vehicle level, not the household level, we account for multiple vehicle households by dividing the estimated household income by the average number of vehicles per household for that income bracket, taken from the 2009 NHTS.

7.2 Regressivity Results

Figure 8 plots for each decile of household income the average tax burden as a percentage of estimated income for the naive uniform tax, the SBO tax, and the optimal Pigouvian emissions tax. The figure also plots the average of pre-tax externality in dollars per year using the right axis. Aside from the 10th, highest, income decile, which has a much higher average annual VMT than the 9th decile, the average pre-tax externality is declining with income. In other words, we find that poorer households have dirtier cars, and pollute more in total even though their annual VMT is lower than richer households. As such, we expect an emissions tax to be regressive to some extent. Indeed, all three taxes are regressive, with the lowest income brackets predicted to pay the highest percentage of household income toward the tax. However, the curve is most steeply sloped for the SBO gasoline tax, indicating that it is the most regressive of the three. The optimal emissions tax imposes a lower average tax burden than the SBO gasoline tax in every income decile.

On its face, it seems surprising that an emissions tax would result in a lower tax burden for all parts of the income distribution. The explanation for this result is that under an

[25] Specifically, the break points for the groupings are at the 10.41, 19.87, 29.02, 41.68, 49.31, 59.30, 72.10, 81.29, and 93.5 percentiles of all households in California.

emissions tax, vehicles with the highest per-gallon tax rate have the fewest post-tax VMT, and vice versa. Even if all vehicles were equally price responsive, an emissions tax would raise less revenue on average than a uniform tax simply because the highest polluting vehicles pay the highest tax rate and thus reduce VMT and gasoline consumption the most. Of course, the core result of this paper is that vehicles are *not* equally responsive to gasoline prices, with the dirtiest vehicles having the greatest VMT elasticity. This further reduces the average burden of an emissions tax, while the variation in elasticities pushes the level and average burden of the SBO gasoline tax higher. Moreover, in practice older and dirtier vehicles have lower VMT pre-tax. Thus, the SBO gasoline tax would only impose a lower tax burden on an income bracket if high-polluting, low VMT vehicles were concentrated in one decile, which is not the case. Indeed, Table 5 shows that the SBO tax is higher than optimal on a per-mile basis for more than 80% of vehicles. Although households in the lower income brackets are more likely to have higher polluting vehicles, more than 75% of vehicles in every income group have emissions below the marginal externality that determines the SBO gasoline tax. As a result, switching from the SBO gasoline tax to an emissions tax lowers the tax burden for the vast majority of vehicles in every income bracket.

It is also important to note that while Figure 8 illustrates that gasoline taxes are regressive, the figure hides a tremendous amount of variation within income deciles. We find that the variance falls as income rises. Figure 9 shows this clearly, plotting several percentiles of SBO gasoline tax expenditures as a share of income within each income bracket. For instance, the interquartile range of the share of SBO gasoline tax as a fraction of income for households in the lowest income decile is between 0.5 percent of household income to over 1.5 percent of household income. In contrast, the interquartile range for the higher deciles is extremely small.

8 Conclusions

In this paper we present three empirical results, all stemming from the stylized fact that vehicle emissions are heterogeneous and highly right skewed. First, the sensitivity of a given vehicle's miles traveled to gasoline prices is correlated with the vehicle's emissions. Dirtier vehicles are more price responsive. This increases the size of the second-best optimal uniform gasoline tax by as much as 50 percent.

Second, gasoline taxes are an inefficient policy tool to reduce vehicle emissions. The optimal policy would differentially tax vehicles based on their emissions, not on consumption of gasoline. While gasoline consumption and emissions are positively correlated, we show that gasoline taxes are a poor substitute for a true Pigouvian emissions tax. The remaining DWL under the second-best optimal gasoline tax exceeds 75 percent in the second half of our sample. Although it comes as no surprise that an indirect tax fails to achieve the optimal result, the magnitude of that failure is striking.

Finally, we find that gasoline taxes are not only regressive, but are more regressive than a Pigouvian tax on emissions. Because the distribution of emissions is so strongly right skewed, with a small number of very high polluting vehicles contributing the bulk of total emissions, a uniform gasoline tax will tend to overtax relative to the social optimum, leaving the vast majority of vehicle owners paying more, and with the poorest households paying substantially more as a fraction of their income.

References

Anderson, M. and M. Auffhammer (2011): "Pounds that Kill: The External Costs of Vehicle Weight," Working Paper 17170, National Bureau of Economic Research.

Borenstein, S. (2012): "The Redistributional Impact of Nonlinear Electricity Pricing," *American Economic Journal: Economic Policy*, 4, 56–90.

Busse, M., C. R. Knittel, and F. Zettelmeyer (forthcoming): "Are Consumers Myopic? Evidence from New and Used Car Purchases," *The American Economic Review*.

Chernick, H. and A. Reschovski (1997): "Who Pays the Gasoline Tax," *Nation Tax Journal*, 50, 233–259.

Davis, L. W. and M. E. Kahn (2010): "International Trade in Used Vehicles: The Environmental Consequences of NAFTA," *American Economic Journal: Economic Policy*, 2, 58–82.

Diamond, P. A. (1973): "Consumption Externalities and Imperfect Corrective Pricing," *The Bell Journal of Economics and Management Science*, 4, pp. 526–538.

Fullerton, D. and S. E. West (2010): "Tax and Subsidy Combinations for the Control of Car Pollution," *The B.E. Journal of Economic Analysis & Policy*, 10.

Gillingham, K. (2010): "Identifying the Elasticity of Driving: Evidence from California," Working Paper.

Greenstone, M., E. Kopits, and A. Wolverton (2011): "Estimating the Social Cost of Carbon for Use in U.S. Federal Rulemakings: A Summary and Interpretation," Working Paper 16913, National Bureau of Economic Research.

Hughes, J. E., C. R. Knittel, and D. Sperling (2008): "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand," *Energy Journal*, 29.

Jacobsen, M. (Forthcoming): "Fuel Economy and Safety: The Influences of Vehicle Class and Driver Behavior," *American Journal of Economics: Applied Economics*, 39, pp. 1–24.

Jacobsen, M. R. and A. A. van Benthem (2013): "Vehicle Scrappage and Gasoline Policy," Working Paper 19055, National Bureau of Economic Research.

Kahn, M. E. (1996): "New Evidence on Trends in Vehicle Emissions," *The RAND Journal of Economics*, 27, 183–196.

Li, S., C. Timmins, and R. H. von Haefen (2009): "How Do Gasoline Prices Affect Fleet Fuel Economy?" *American Economic Journal: Economic Policy*, 1.

Matthews, H. S. and L. B. Lave (2000): "Applications of Environmental Valuation for Determining Externality Costs," *Environmental Science & Technology*, 34, 1390–1395.

Morrow, S. and K. Runkle (2005): "April 2004 Evaluation of the California Enhanced Vehicle Inspection and Maintenance (Smog Check) Program," Report to the legislature, Air Resources Board.

Muller, N. Z. and R. Mendelsohn (2009): "Efficient Pollution Regulation: Getting the Prices Right," *American Economic Review*, 99, 1714–39.

Parry, I. W. H. and K. A. Small (2005): "Does Britain or the United States Have the Right Gasoline Tax?" *American Economic Review*, 95, 1276–1289.

Sandler, R. (2012): "Clunkers or Junkers? Adverse Selection in a Vehicle Retirement Program," *American Economics Journal: Economic Policy*, 4, 253–281.

Westberg, K., N. Cohen, and K. W. Wilson (1971): "Carbon Monoxide: Its Role in Photochemical Smog Formation," *Science*, 171, 1013–1015.

Figures and Tables

Figures

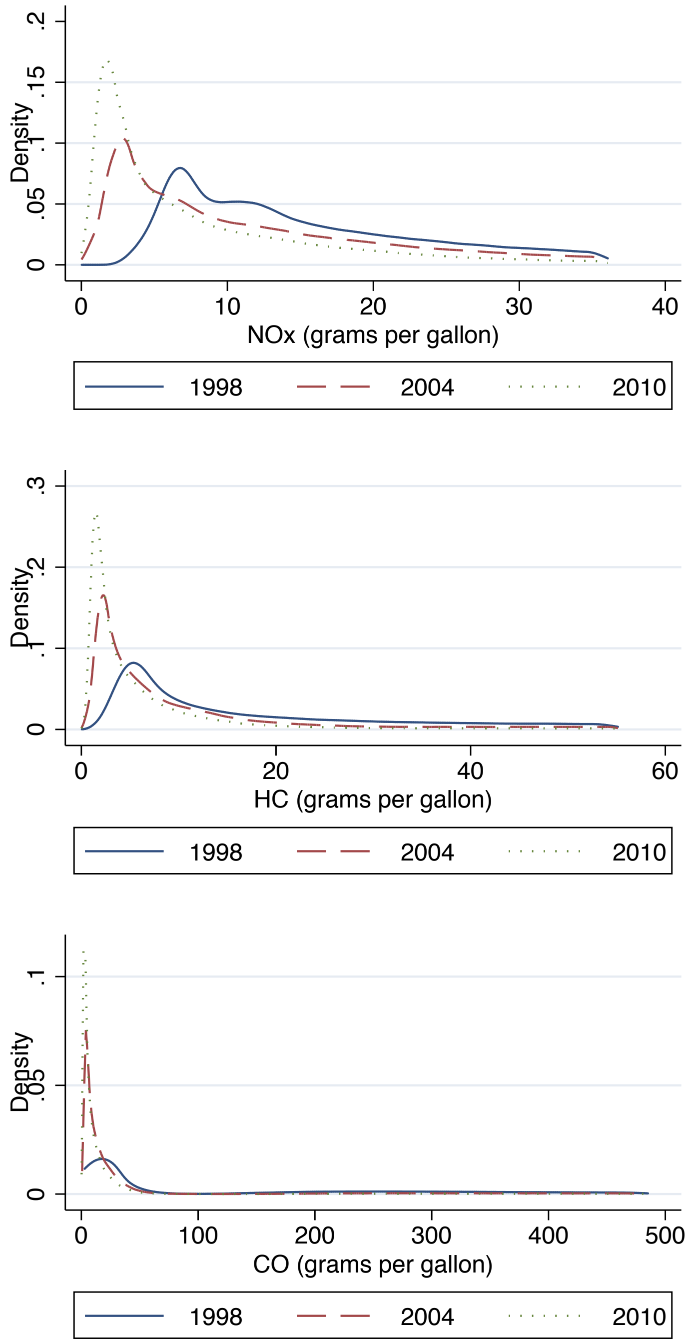

Figure 1: Distribution of three criteria pollutant emissions across all vehicles in 1998, 2004, and 2010 (observations above the 90th percentile are omitted)

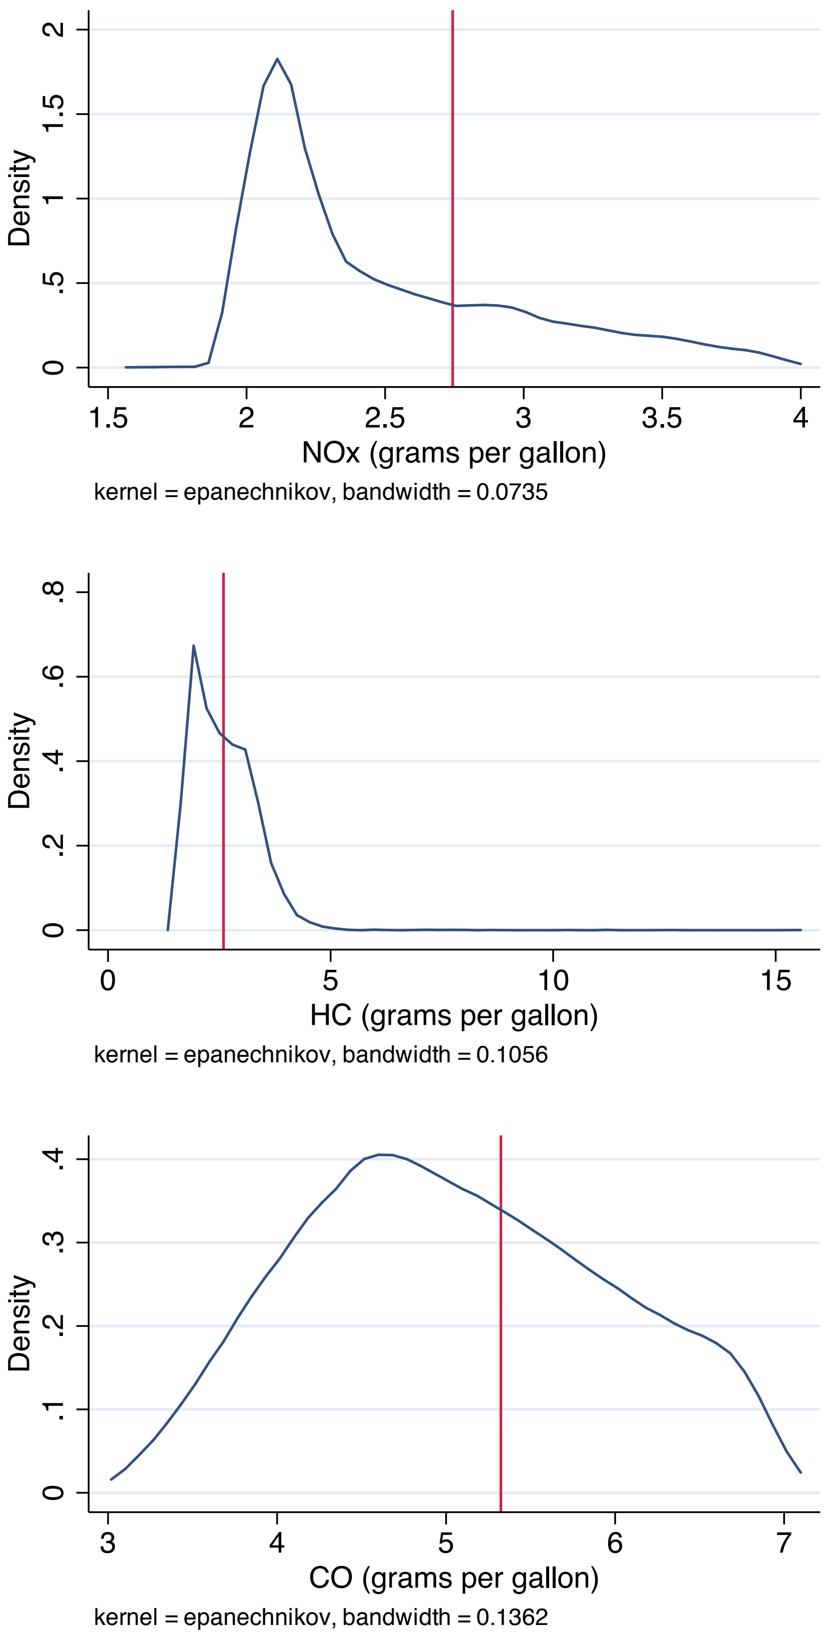

Figure 2: Distribution of three criteria pollutant emissions of a 2001 4-door, 1.8L, Toyota Corolla in 2009 (observations above the 90th percentile are omitted)

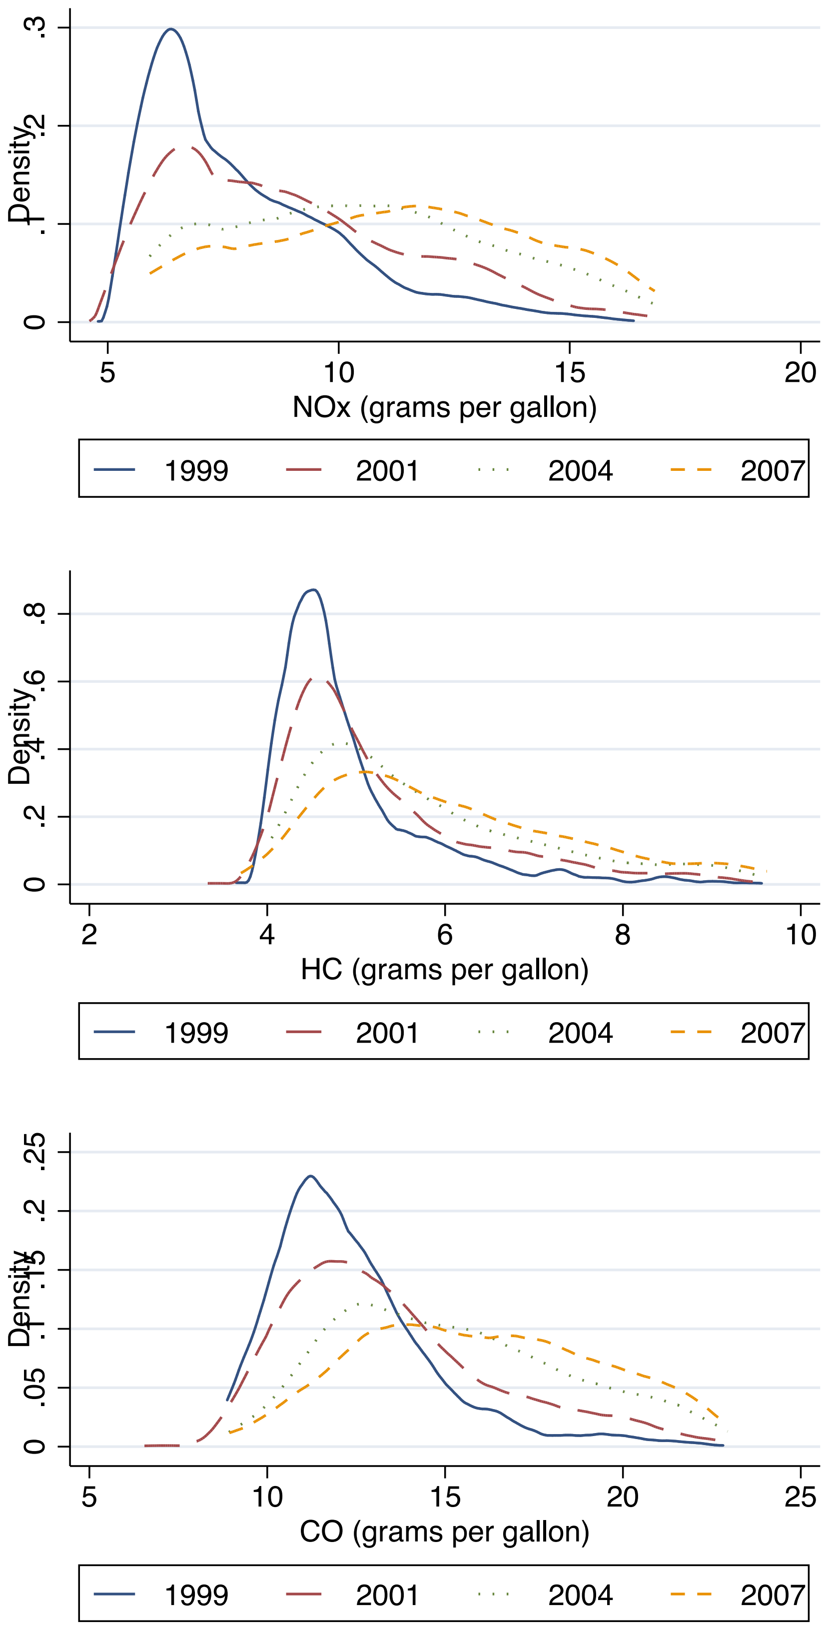

Figure 3: Distribution of three criteria pollutant emissions of a 1995 3.8L, FWD, Ford Windstar in 1999, 2001, 2005, and 2009 (observations above the 90th percentile are omitted)

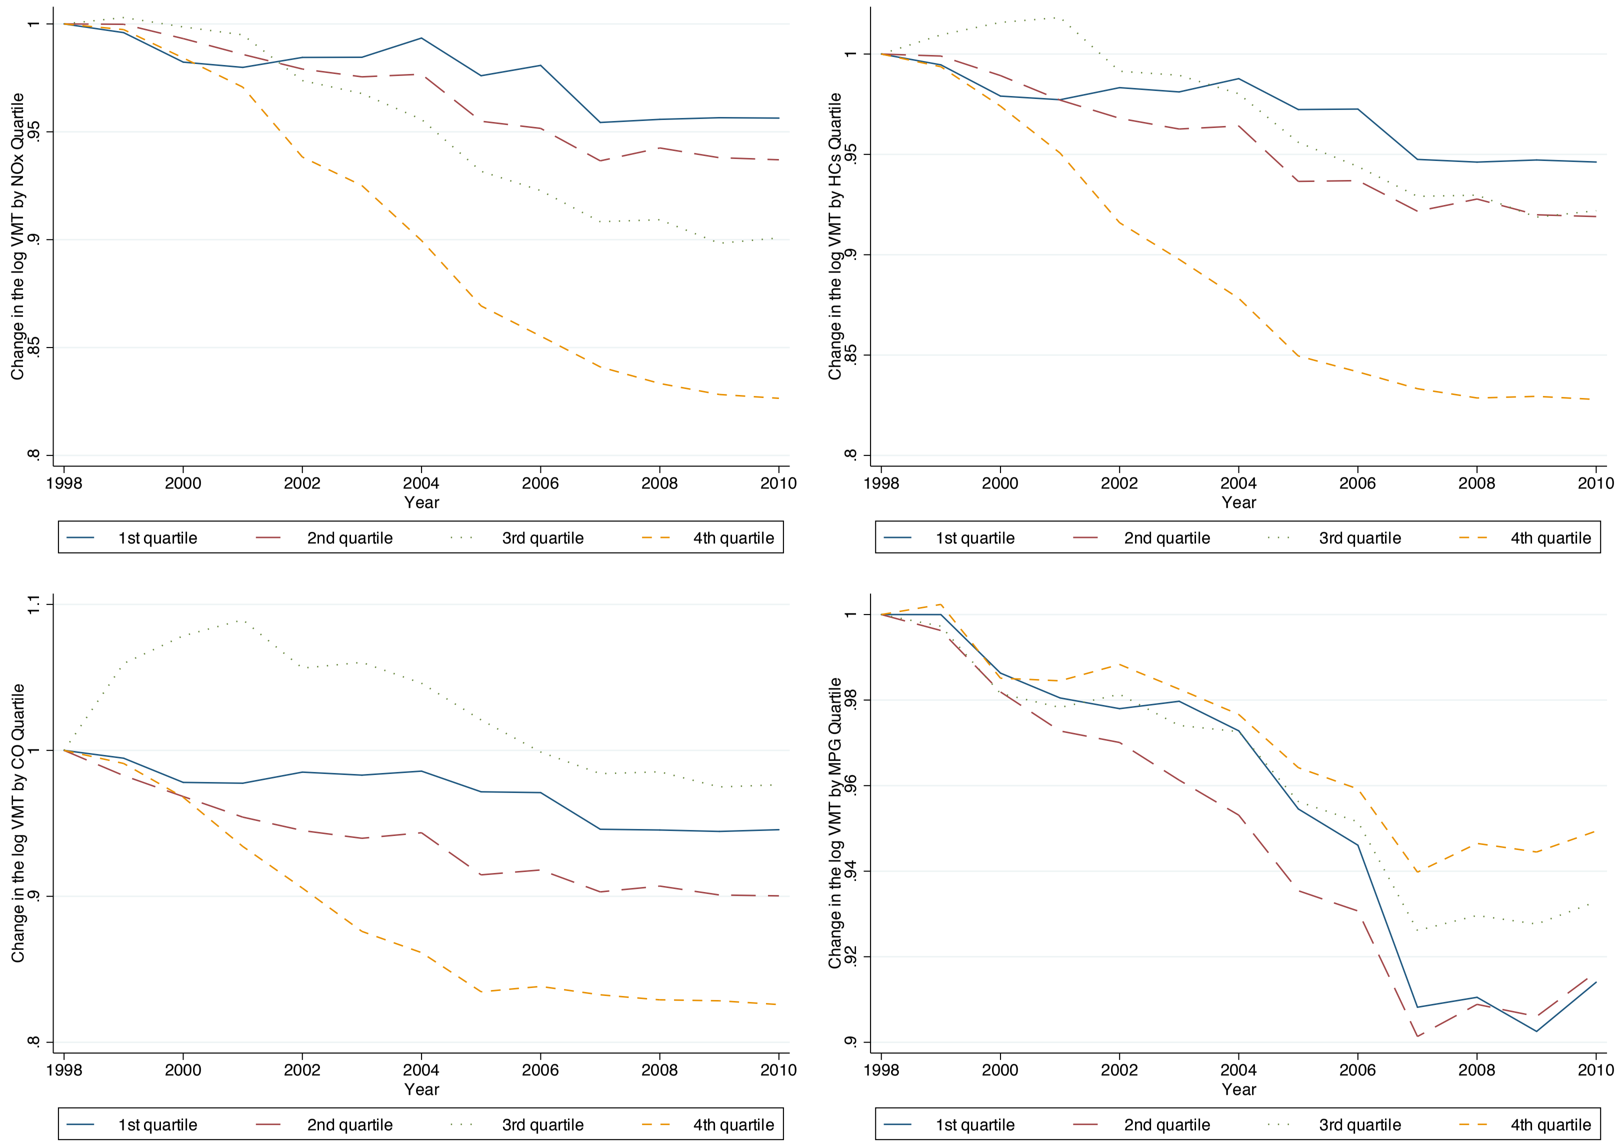

Figure 4: Change in the log of VMT over sample by pollutant quartile

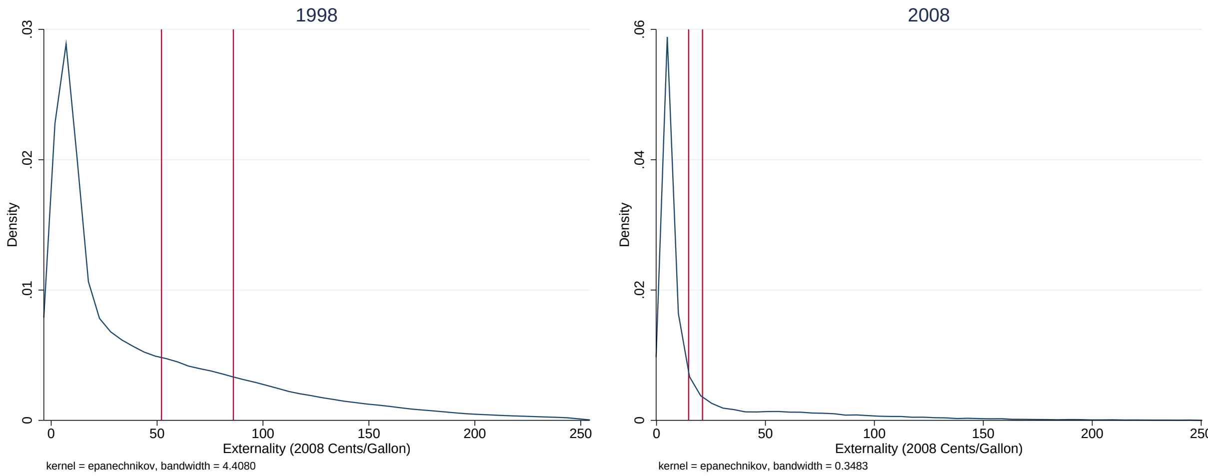

Figure 5: Distribution of externality per gallon—vertical lines indicate naive and marginal uniform tax

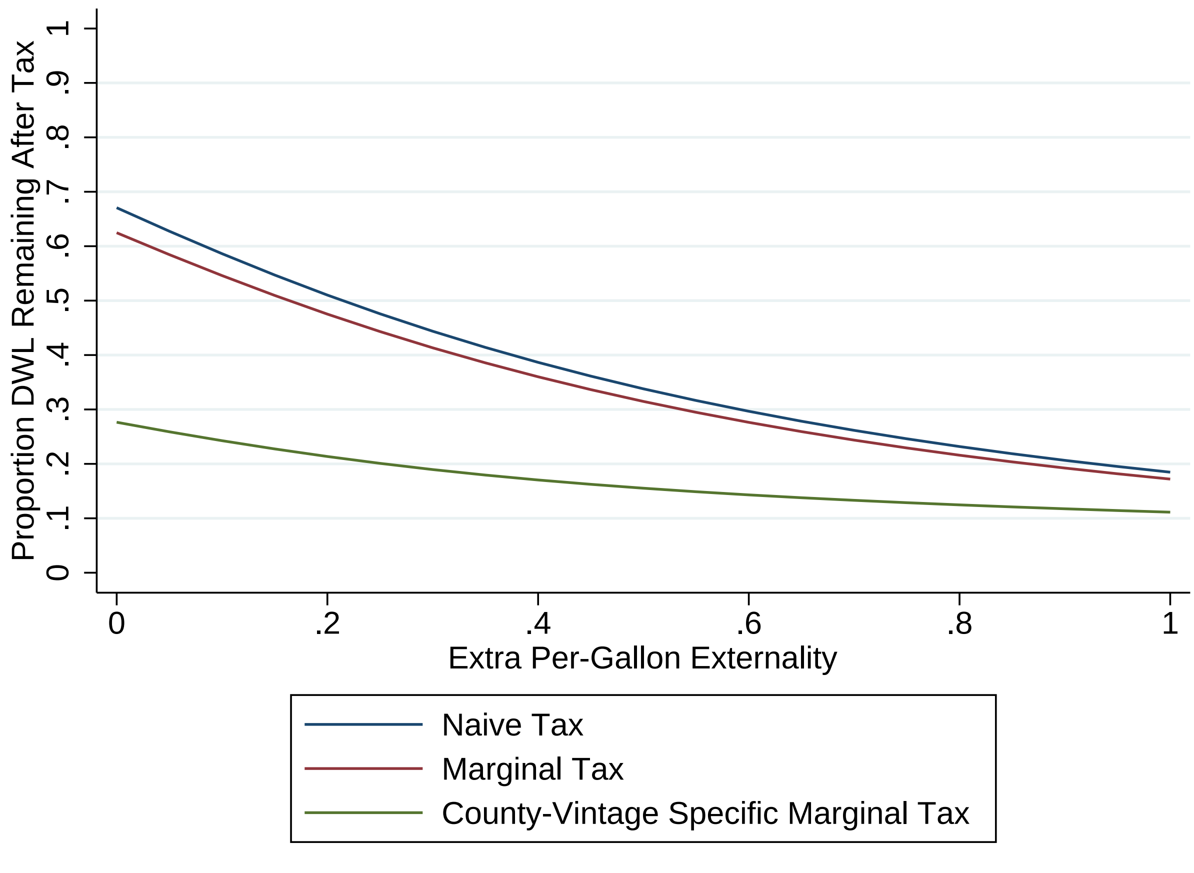

Figure 6: Remaining deadweight loss under alternatives gasoline-specific externalities

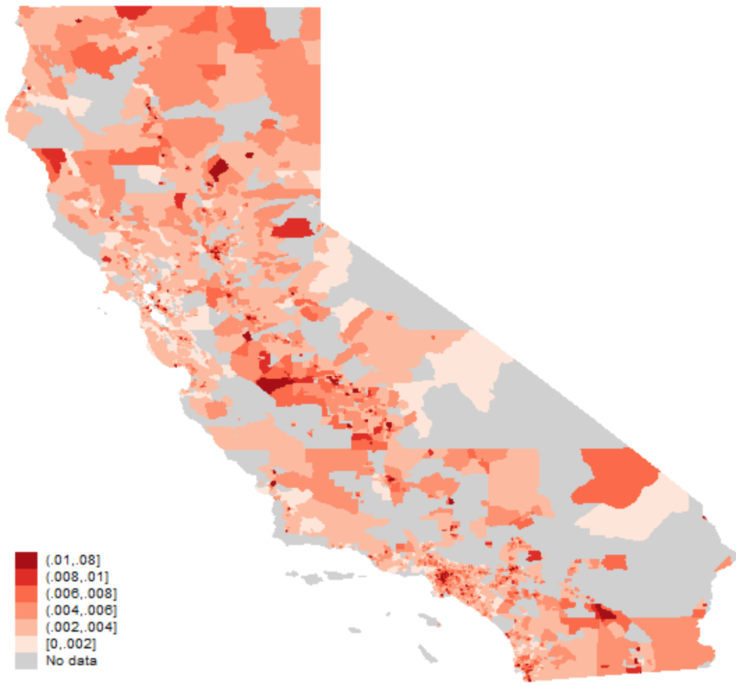

(a) SBO Tax

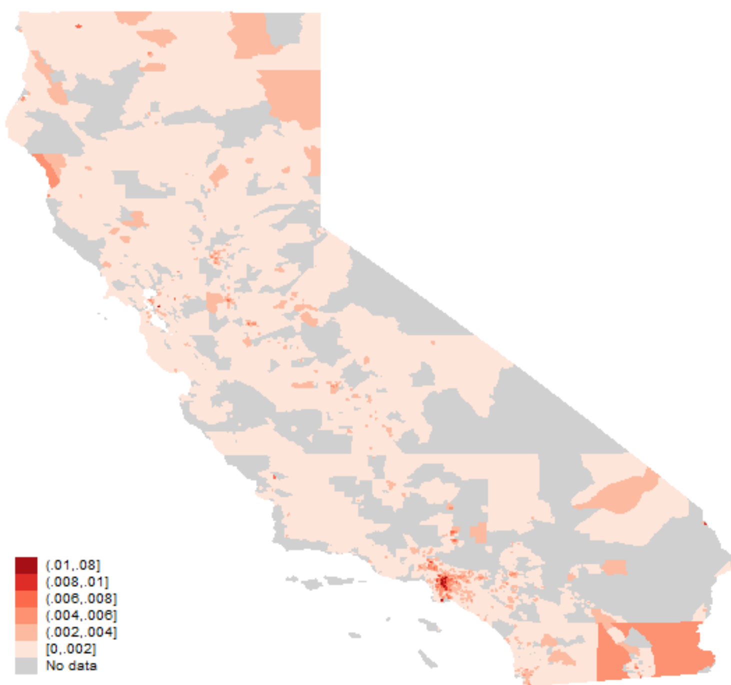

(b) Pigouvian Emissions Tax

Figure 7: Simulated Annual Tax Paid As a Proportion of CBG Median Income

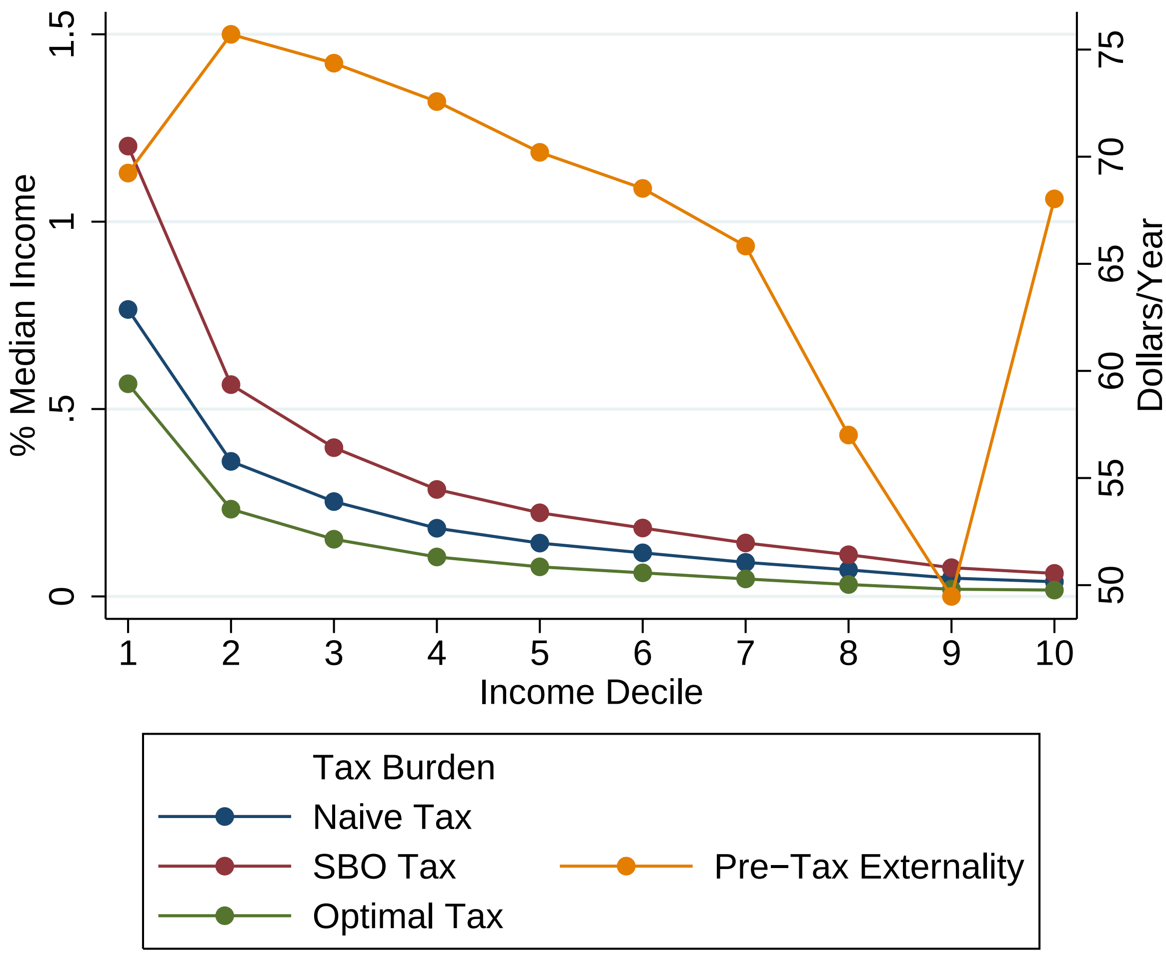

Figure 8: Annual Tax Burden as a Percent of Income, and Annual Pre-Tax Externality, by Household Income "Deciles"

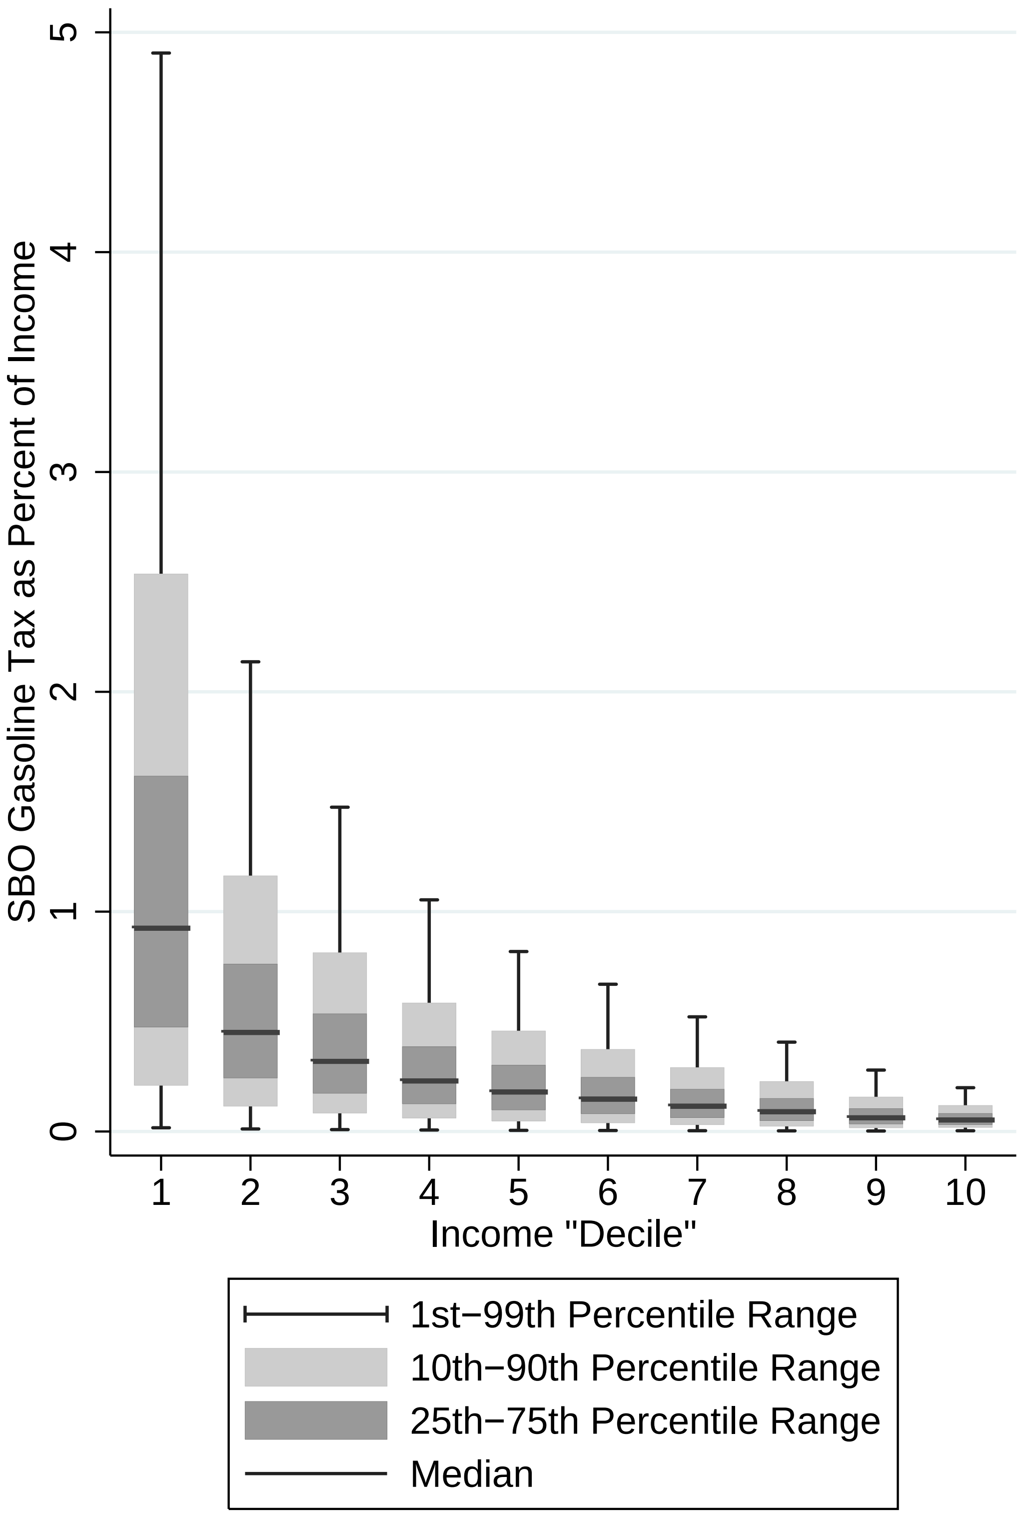

Figure 9: Distribution of Tax Burden, by Household Income “Deciles”

Tables

Table 1: Summary Statistics

	All	Vehicle Age			Year	
		4-9	10-15	16-28	1998	2008
Weighted Fuel Economy	23.53	23.30	23.70	23.80	24.27	23.07
	(5.320)	(5.235)	(5.331)	(5.507)	(5.478)	(5.169)
Average $/mile	0.0973	0.0928	0.0977	0.109	0.0581	0.143
	(0.0416)	(0.0403)	(0.0417)	(0.0426)	(0.0134)	(0.0361)
Odometer (00000s)	1.214	0.932	1.376	1.626	1.023	1.323
	(0.605)	(0.454)	(0.567)	(0.688)	(0.528)	(0.622)
Average VMT/Day	24.50	30.23	23.09	15.64	29.48	22.62
	(95.96)	(43.29)	(101.2)	(148.9)	(65.64)	(28.24)
Grams/mile HC	0.762	0.219	0.739	2.017	1.403	0.542
	(1.177)	(0.270)	(1.019)	(1.670)	(1.506)	(1.022)
Grams/mile CO	5.525	0.510	4.814	18.15	12.44	3.488
	(12.91)	(1.646)	(10.90)	(20.35)	(18.97)	(10.49)
Grams/mile NOx	0.664	0.317	0.731	1.297	1.042	0.516
	(0.638)	(0.303)	(0.599)	(0.728)	(0.904)	(0.547)
Failed Smog Check	0.0868	0.0435	0.106	0.165	0.0515	0.0992
	(0.282)	(0.204)	(0.307)	(0.371)	(0.221)	(0.299)
Average HH Income	48066.3	49955.1	47117.3	44970.8	49768.5	47778.3
	(17031.0)	(17685.0)	(16556.3)	(15555.7)	(17952.5)	(16791.9)
Truck	0.385	0.406	0.368	0.367	0.322	0.426
	(0.487)	(0.491)	(0.482)	(0.482)	(0.467)	(0.494)
Vehicle Age	10.68	6.694	12.14	18.54	9.244	11.77
	(4.587)	(1.615)	(1.686)	(2.478)	(3.552)	(4.854)
N	76510820	34713936	29775806	12008157	4172978	5849644

Statistics are means with standard deviations presented below in paranetheses. Weighted fuel economy is from EPA. Dollars per mile is the average gasoline price from EIA in between the vehicle's current and previous Smog Checks, divided by the vehicle's fuel economy. Average household income is taken from the 2000 Census ZCTA where the Smog Check occured. The dataset used for this table contains one observation per vehicle per year in which a Smog Check occured.

Table 2: Vehicle Miles Traveled, Dollars Per Mile, and Nitrogen Oxides (Quartiles by year)

	(1) Model 1	(2) Model 2	(3) Model 3	(4) Model 4	(5) Model 5	(6) Model 6
ln(DPM)	-0.269**	-0.123**	-0.183**	-0.134**		-0.038
	(0.044)	(0.038)	(0.027)	(0.022)		(0.028)
ln(DPM) * NO_x Q1					0.043*	
					(0.021)	
ln(DPM) * NO_x Q2					-0.054*	
					(0.022)	
ln(DPM) * NO_x Q3					-0.152**	
					(0.025)	
ln(DPM) * NO_x Q4					-0.280**	
					(0.028)	
ln(DPM)*NO_x Centile						-0.001**
						(0.000)
NO_x Q2					0.216	
					(0.663)	
NO_x Q3					-1.742	
					(0.881)	
NO_x Q4					-2.417*	
					(1.003)	
NO_x Centile						-0.001
						(0.001)
Truck	0.054	0.057	0.005			
	(0.033)	(0.045)	(0.055)			
Time Trend	-0.244**	-0.314**	-0.278**	-0.035	-0.057	-0.062*
	(0.037)	(0.024)	(0.015)	(0.028)	(0.032)	(0.025)
Time Trend-Squared	0.002**	0.002**	0.002**	0.000	0.000	0.001*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Vintage Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes	Yes	Yes
Make Fixed Effects	No	Yes	No	No	No	No
Vin Prefix Fixed Effects	No	No	Yes	No	No	No
Vehicle Fixed Effects	No	No	No	Yes	Yes	Yes
Observations	36387455	36387455	36387455	36387455	29779909	29779909
R-squared	0.210	0.218	0.143	0.121	0.117	0.118

* $p < 0.05$, ** $p < 0.01$

Notes: Each observation is a vehicle's Smog Check inspection. Dependent variable is the log of average daily vehicle miles travelled since the previous inspection. DPM represents the average gasoline price over the period since the previous inspection, converted to dollars per mile by dividing by vehicle fuel economy Quartiles and centiles of NO_x are based on rankings of emissions per mile within the calendar year in which the Smog Check occurs. Standard errors clustered by vehicle make reported in parentheses.

Table 3: Average and Marginal Pollution Externality

	Average Externality (¢/gal)	Marginal Externality (¢/gal)
1998	61.48	91.27
1999	54.78	81.62
2000	48.55	74.31
2001	40.96	64.29
2002	34.18	54.09
2003	28.77	46.89
2004	24.31	39.26
2005	21.25	33.95
2006	18.61	29.52
2007	16.23	25.81
2008	14.36	22.84

Notes: Average Externality is the simple average of damages from emissions of criteria pollutants produced by each car in each year, divided by fuel usage. We refer to a tax on the average externality as the "naive tax". The marginal externality is computed as the weighted average of externality per gallon, using the negative slope of the vehicle's demand curve as the weight.. A tax on the marginal externality is the SBO gasoline tax. Both calculations also weight vehicles by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole. Dollar figures inflation adjusted to year 2008.

Table 4: Ratios of DWL with Tax to DWL With No Tax

	Statewide Tax		County-Level Taxes		Vintage Tax		County/Vintage Tax		Total DWL
	Naive	SBO	Naive	SBO	Naive	SBO	Naive	SBO	$
1998	0.616	0.568	0.573	0.523	0.348	0.341	0.296	0.290	196466745.4
1999	0.636	0.577	0.592	0.529	0.330	0.325	0.269	0.265	158104024.8
2000	0.635	0.583	0.587	0.532	0.320	0.317	0.253	0.251	131221907.6
2001	0.690	0.627	0.649	0.582	0.348	0.345	0.281	0.279	100426398.8
2002	0.700	0.675	0.652	0.625	0.348	0.346	0.284	0.283	76704235.2
2003	0.716	0.699	0.661	0.643	0.316	0.314	0.248	0.247	58869860.6
2004	0.746	0.740	0.699	0.693	0.313	0.312	0.246	0.245	42633365.5
2005	0.766	0.762	0.723	0.718	0.319	0.318	0.250	0.250	27431776.9
2006	0.801	0.796	0.762	0.757	0.338	0.337	0.272	0.271	20756466.1
2007	0.817	0.817	0.780	0.780	0.328	0.327	0.259	0.357	15589665.8
2008	0.838	0.836	0.805	0.802	0.331	0.331	0.264	0.264	12340287.7
Average	0.724	0.698	0.680	0.653	0.331	0.329	0.266	0.273	76413157.7

Notes: DWL with no tax calculated based on the difference in emissions from imposing a tax equal to the actual externality per gallon consumed by a particular car. SBO tax computed as the weighted average of externality per gallon, using the negative slope of the vehicle's demand curve as the weight. All taxes also weight vehicles by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole.

Table 5: Proportion of Vehicles for which a Uniform Tax Overshoots the Optimal Tax

	Mean
Naive Tax on Fleet Average Externality	0.724
SBO Tax on Fleet Marginal Externality	0.803
Naive Tax on County Average Externality	0.714
SBO Tax on County Marginal Externality	0.793
Naive Tax on Vintage Average Externality	0.708
SBO Tax on Vintage Marginal Externality	0.733
Naive Tax on County/Vintage Average Externality	0.673
SBO Tax on County/Vintage Marginal Externality	0.718
N	36023471

Proportion of vehicles over the period 1998-2008 whose VMT would be lower than optimal under the indicated tax. We assume that the tax is adjusted each calendar year to reflect changes in the average or marginal externality

Table 6: Ratios of DWL with Tax to DWL With No Tax, Scrapping Most Polluting Vehicles

	Percentile Scrapped					
	None	1%	2%	5%	10%	25%
1998	0.568	0.476	0.451	0.426	0.419	0.439
1999	0.577	0.484	0.467	0.453	0.452	0.469
2000	0.583	0.501	0.486	0.472	0.471	0.478
2001	0.627	0.531	0.515	0.501	0.501	0.456
2002	0.675	0.590	0.578	0.566	0.564	0.509
2003	0.699	0.634	0.625	0.615	0.613	0.488
2004	0.740	0.681	0.672	0.662	0.656	0.458
2005	0.762	0.704	0.693	0.678	0.657	0.390
2006	0.796	0.735	0.725	0.707	0.668	0.363
2007	0.817	0.763	0.753	0.733	0.672	0.378
2008	0.836	0.783	0.772	0.745	0.628	0.373
Average	0.698	0.626	0.612	0.596	0.573	0.436

Notes: DWL with no tax calculated based on the difference in emissions from imposing a tax equal to the actual externality per gallon consumed by a particular car. SBO tax computed as the weighted average of externality per gallon, using the negative slope of the vehicle's demand curve as the weight. All taxes also weight vehicles by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole.

Table 7: Ratios of DWL with Tax to DWL With No Tax

	Statewide Tax		County-Level Taxes		Vintage Tax		County/Vintage Tax	
	Average	Marginal	Average	Marginal	Average	Marginal	Average	Marginal
1998	0.571	0.434	0.536	0.397	0.318	0.295	0.280	0.254
1999	0.590	0.426	0.558	0.390	0.301	0.279	0.256	0.235
2000	0.591	0.433	0.556	0.397	0.299	0.281	0.250	0.235
2001	0.648	0.472	0.619	0.440	0.329	0.312	0.280	0.266
2002	0.619	0.490	0.586	0.459	0.324	0.310	0.281	0.268
2003	0.625	0.503	0.589	0.469	0.314	0.303	0.268	0.259
2004	0.647	0.544	0.619	0.516	0.351	0.341	0.309	0.301
2005	0.644	0.548	0.617	0.522	0.347	0.336	0.306	0.296
2006	0.692	0.595	0.669	0.573	0.397	0.390	0.360	0.353
2007	0.674	0.585	0.653	0.564	0.368	0.362	0.329	0.325
2008	0.701	0.605	0.682	0.586	0.388	0.383	0.349	0.345
Average	0.636	0.512	0.608	0.483	0.340	0.327	0.297	0.285

Notes: DWL with no tax calculated based on the difference in emissions from imposing a tax equal to the actual externality per gallon consumed by a particular car. SBO tax computed as the weighted average of externality per gallon, using the negative slope of the vehicle's demand curve as the weight. All taxes also weight vehicles by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole.

For Online Publication

A Proofs of Propositions

Proposition 1. *The second-best-optimal uniform per-gallon gasoline tax, τ^*, is* (from Diamond (1973))*:*

$$\tau^* = \frac{-\sum_i \sum_{h\neq i} \frac{\partial U^h}{\partial \alpha_i}\alpha_i'}{\sum_i \alpha_i'}. \tag{14}$$

where α_i' is the derivative of consumer i's demand for gasoline with respect to the price of gasoline.

Proof. Consumers have quasi-linear utility functions, given as:

$$\max_{\alpha_h} \quad U^h(\alpha_1, \alpha_2, ..., \alpha_h, ..., \alpha_n) + \mu_h, \tag{15}$$

$$s.t. \quad (p_g + \tau)\alpha_h + \mu_h = m_h, \tag{16}$$

where p_g is the price, τ the tax per gallon, α_h the consumption of the polluting good by consumer h, μ_h consumption of a numeraire, and m_h consumer h's income. Assuming an interior solution, we have:

$$\frac{\partial U^h}{\partial \alpha_h} = (p + \tau). \tag{17}$$

This yields demand curves, which we represent by α_h^*, given by:

$$\alpha_h^* = \alpha_h\left(p_g + \tau\right). \tag{18}$$

The SBO gasoline tax maximizes social welfare, or the sum of utilities:

$$W(\tau) = \sum_h U^h[\alpha_1^*, ..., \alpha_h^*, ...\alpha_n^*] - p_g \sum_h \alpha_h^* + \sum_h m_h. \tag{19}$$

The first-order condition for the SBO gasoline tax is given as:

$$W'(\tau) = \sum_i \sum_h \frac{\partial U^h}{\partial \alpha_i}\alpha_i' - p_g \sum_h \alpha_h' = 0. \tag{20}$$

Rewriting this and plugging in the result from the consumers' problem, $\frac{\partial U^h}{\partial \alpha_h} - p_g = \tau$, we have:

$$W'(\tau) = \sum_i \sum_{h\neq i} \frac{\partial U^h}{\partial \alpha_i}\alpha_i' + \tau \sum_i \alpha_i' = 0. \tag{21}$$

Solving for the second-best tax yields:

$$\tau^* = \frac{-\sum_i \sum_{h \neq i} \frac{\partial U^h}{\partial \alpha_i} \alpha_i'}{\sum_i \alpha_i'}. \tag{22}$$

□

Proposition 2. *Suppose drivers are homogenous in their demand for gasoline, but vehicles' per-gallon emissions differ. In particular, let β denote the derivative of the demand for gasoline with respect to the price of gasoline.*

If the distribution of the per-gallon externality, E, is log normal, with probability density function:

$$\varphi(E_i) = \frac{1}{E_i\sqrt{2\sigma_E^2}} \exp\left(\frac{-(E_i - \mu_E)^2}{2\sigma_E^2}\right), \tag{23}$$

the DWL absent any market intervention will be given as:

$$D = \frac{1}{2\beta} e^{2\mu_E + 2\sigma_E^2}.$$

Proof. Given these assumptions, the deadweight loss absent any market intervention will be given as:

$$\begin{aligned} D &= \int_0^\infty \frac{(E_i)^2}{2\beta} \varphi(E_i) dE_i \\ &= \frac{1}{2\beta} \mathbb{E}[E_i^2] \\ &= \frac{1}{2\beta} e^{2\mu_E + 2\sigma_E^2}. \end{aligned} \tag{24}$$

□

Proposition 3. *Under the assumptions in Proposition 2, the ratio of the remaining DWL with the deadweight loss after the tax is:*

$$R = \frac{D - \frac{e^{2\mu_E + \sigma_E^2}}{2\beta}}{D} = 1 - \frac{e^{2\mu_E + \sigma_E^2}}{e^{2\mu_E + 2\sigma_E^2}} = 1 - e^{-\sigma_E^2}. \tag{25}$$

Proof. The level of the externality is given as:

$$\overline{E} = \tau = e^{\mu_E + \sigma_E^2/2}. \tag{26}$$

The deadweight loss associated with all vehicles is given as:

$$\begin{aligned} D(\tau) &= \int_0^\infty \frac{(\tau - E_i)^2}{2\beta}\varphi(E_i)dE_i \\ &= \frac{1}{2\beta}\mathbb{E}[\tau^2 - 2\tau E_i + E_i^2] \\ &= \frac{1}{2\beta}(\tau^2 - 2\tau\mathbb{E}[E_i] + \mathbb{E}[E_i^2]) \\ &= \frac{1}{2\beta}(\tau^2 - 2\tau e^{\mu_E + \frac{\sigma_E^2}{2}} + e^{2\mu_E + 2\sigma_E^2}) \\ &= \frac{1}{2\beta}(\tau^2 - 2\tau e^{\mu_E + \frac{\sigma_E^2}{2}}) + D \\ &= D - \frac{e^{2\mu_E + \sigma_E^2}}{2\beta}. \end{aligned} \tag{27}$$

The ratio of remaining DWL with the deadweight loss absent the tax is therefore:

$$R = \frac{D - \frac{e^{2\mu_E+\sigma_E^2}}{2\beta}}{D} = 1 - \frac{e^{2\mu_E+\sigma_E^2}}{e^{2\mu_E+2\sigma_E^2}} = 1 - e^{-\sigma_E^2}. \tag{28}$$

□

Proposition 4. *When $B_i = \frac{1}{\beta_i}$ and E_i are distributed lognormal with dependence parameter ρ, the optimal tax, represented by τ^*, is:*

$$\tau^* = e^{\mu_E + \frac{\sigma_E^2}{2} + \rho\sigma_E\sigma_B}$$

Proof. The slope of the demand curve with respect to the cost of driving, defined as $B_i = \frac{1}{\beta_i}$, where β_i is the VMT elasticity for the vehicle owned by consumer i is distributed lognormal with parameters μ_B and σ_B^2. ρ is the dependence parameter of the bivariate lognormal distribution (the correlation coefficient of $\ln E$ and $\ln B$). The optimal tax is:

$$\begin{aligned} \tau^* &= \frac{\sum E_i\beta_i}{\sum \beta_i} \\ &= \frac{\frac{1}{N}\sum E_i\beta_i}{\frac{1}{N}\sum \beta_i} \\ &= \frac{\mathbb{E}[E_i\beta_i]}{\mathbb{E}[\frac{1}{B_i}]} \\ &= \frac{e^{\mu_E + \frac{\sigma_E^2}{2} - \mu_B + \frac{\sigma_B^2}{2}} e^{\rho\sigma_E\sigma_B}}{e^{-\mu_B + \frac{\sigma_B^2}{2}}} \\ &= e^{\mu_E + \frac{\sigma_E^2}{2} + \rho\sigma_E\sigma_B}. \end{aligned} \tag{29}$$

□

Proposition 5. *When $B_i = \frac{1}{\beta_i}$ and E_i are distributed lognormal with dependence parameter*

ρ, the ratios of the remaining deadweight loss after the SBO gasoline tax to the original deadweight loss will be:

$$R(\tau^*) = 1 - e^{-\sigma_E^2}, \tag{30}$$

and, the ratios of the remaining deadweight loss after the naive uniform tax to the original deadweight loss will be:

$$R(\tau_{naive}) = 1 - e^{-\sigma_E^2}(2e^{-\rho\sigma_E\sigma_B} - e^{-2\rho\sigma_E\sigma_B}). \tag{31}$$

Proof. The deadweight loss with no gasoline tax is:

$$\begin{aligned} \mathcal{D} &= \int_0^\infty \left(\int_0^\infty \frac{(E_i)^2 B_i}{2} \varphi(E_i) dE_i \right) \varphi_B(B_i) dB_i \\ &= \frac{1}{2}\mathbb{E}[E_i^2 B_i] \\ &= \frac{1}{2} e^{2\mu_E + 2\sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2} + 2\rho\sigma_E\sigma_B}. \end{aligned} \tag{32}$$

The deadweight loss with the optimal uniform tax is:

$$\begin{aligned} \mathcal{D}(\tau^*) &= \int_0^\infty \left(\int_0^\infty \frac{(\tau - E_i)^2 B_i}{2} \varphi(E_i) dE_i \right) \varphi_B(B_i) dB_i \\ &= \frac{1}{2}\mathbb{E}[\tau^2 B_i - 2\tau E_i B_i + E_i^2 B_i] \\ &= \frac{1}{2}(\tau^2 \mathbb{E}[B_i] - 2\tau \mathbb{E}[E_i B_i] + \mathbb{E}[E_i^2 B_i]) \\ &= \frac{1}{2}(\tau^2 e^{\mu_B + \frac{\sigma_B^2}{2}} - 2\tau e^{\mu_E + \frac{\sigma_E^2}{2} + \mu_B + \frac{\sigma_B^2}{2} + \rho\sigma_E\sigma_B} + e^{2\mu_E + 2\sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2} + 2\rho\sigma_E\sigma_B}) \\ &= \frac{1}{2} e^{2\mu_E + \sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2} + 2\rho\sigma_E\sigma_B} - e^{2\mu_E + \sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2} + 2\rho\sigma_E\sigma_B} + \mathcal{D} \\ &= \mathcal{D} - \frac{1}{2} e^{2\mu_E + \sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2} + 2\rho\sigma_E\sigma_B}, \end{aligned} \tag{33}$$

while the deadweight loss with the naive tax, equal to the average externality level is:

$$\mathcal{D}(\tau_{naive}) = \mathcal{D} - \frac{1}{2}(2e^{2\mu_E + \sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2} + \rho\sigma_E\sigma_B} - e^{2\mu_E + \sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2}}). \tag{34}$$

Then the ratios of the remaining deadweight loss after a tax to the original deadweight loss will be:

$$\begin{aligned} R(\tau^*) &= 1 - \frac{e^{2\mu_E + \sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2} + 2\rho\sigma_E\sigma_B}}{e^{2\mu_E + 2\sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2} + 2\rho\sigma_E\sigma_B}} \\ &= 1 - e^{-\sigma_E^2}, \end{aligned} \tag{35}$$

$$
\begin{aligned}
R(\tau_{naive}) &= 1 - \frac{2e^{2\mu_E + \sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2}\rho\sigma_E\sigma_B} - e^{2\mu + \sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2}}}{e^{2\mu_E + 2\sigma_E^2 + \mu_B + \frac{\sigma_B^2}{2} + 2\rho\sigma_E\sigma_B}} \\
&= 1 - e^{-\sigma_E^2}(2e^{-\rho\sigma_E\sigma_B} - e^{-2\rho\sigma_E\sigma_B}). \qquad (36)
\end{aligned}
$$

□

B Steps to Clean Smog Check Data

California implemented its first inspection and maintenance program (the Smog Check Program) in 1984 in response to the 1977 Clean Air Act Amendments. The 1990 Clean Air Act Amendments required states to implement an enhanced inspection and maintenance program in areas with serious to extreme non-attainment of ozone limits. Several of California's urban areas fell into this category, and in 1994, California's legislature passed a redesigned inspection program was passed by California's legislature after reaching a compromise with the EPA. The program was updated in 1997 to address consumer complaints, and fully implemented by 1998. Among other improvements, California's new program introduced a system of centralized "Test-Only" stations and an electronic transmission system for inspection reports.[26] Today, more than a million smog checks take place each month.

Since 1998, the state has been divided into three inspection regimes (recently expanded to four), the boundaries of which roughly correspond to the jurisdiction of the regional Air Quality Management Districts. "Enhanced" regions, designated because they fail to meet state or federal standards for CO and ozone, fall under the most restrictive regime. All of the state's major urban centers are in Enhanced areas, including the greater Los Angeles, San Francisco, and San Diego metropolitan areas. Vehicles registered to an address in an Enhanced area must pass a biennial smog check in order to be registered, and they must take the more rigorous Acceleration Simulation Mode (ASM) test. The ASM test involves the use of a dynamometer, and allows for measurement of NO_x emissions. In addition, a randomly selected two percent sample of all vehicles in these areas is directed to have their smog checks at Test-Only stations, which are not allowed to make repairs.[27] Vehicles that match a "High Emitter Profile" are also directed to Test-Only stations, as are vehicles that

[26] For more detailed background see http://www.arb.ca.gov/msprog/smogcheck/july00/if.pdf.

[27] Other vehicles can be taken to Test-Only stations as well if the owner chooses, although they must get repairs elsewhere if they fail.

are flagged as "gross polluters" (those that fail an inspection with twice the legal limit of one or more pollutant in emissions). More recently some "Partial-Enhanced" areas that require a biennial ASM test have been added, but no vehicles are directed to Test-Only stations.

Areas with poor air quality not exceeding legal limits fall under the Basic regime. Cars in a Basic area must have biennial smog checks as part of registration, but they are allowed to take the simpler Two Speed Idle (TSI) test and are not directed to Test-Only stations. The least restrictive regime, consisting of rural mountain and desert counties in the east and north, is known as the Change of Ownership area. As the name suggests, inspections in these areas are only required upon change of ownership; no biennial smog check is required.

Our data from the Smog Check Program essentially comprise the universe of test records from January 1, 1996 to December 31, 2010. We were able to obtain test records only going back to 1996 because this was the year when the Smog Check Program introduced its electronic transmission system. Because the system seems to have been phased in during the first half of 1996, and major program changes took effect in 1998 we limit our sample to test records from January 1998 on. For our analyses, we use a 10 percent sample of VINs, selecting by the second to last digit of the VIN. We exclude tests that have no odometer reading, with a test result of "Tampered" or "Aborted" and vehicles that have more than 36 tests in the span of the data. Vehicles often have multiple smog check records in a year, whether due to changes of ownership or failed tests, but we argue that more than 36 in what is at most a 12 year-span indicates some problem with the data.[28]

A few adjustments must be made to accurately estimate VMT and emissions per mile.

First, we adjust odometer readings for "roll overs" and typos. Many of the vehicles in our analysis were manufactured with 5-digit odometers–that is, five places for whole numbers plus a decimal. As such, any time one of these vehicles crosses over 100,000 miles, the odometer "rolls over" back to 0. To complicate matters further, sometimes either the vehicle owner or smog check technician notices this problem and records the appropriate number in the 100,000s place, and sometimes they do not. To address this problem, we employ an algorithm that increases the hundred thousands place in the odometer reading whenever a

[28] For instance, there is one vehicle in particular, a 1986 Volvo station wagon, which has records for more than 600 smog checks between January 1996 and March 1998. The vehicle likely belonged to a smog check technician who used it to test the electronic transmission system.

rollover seems to have occurred. The hundred thousands are incremented if the previous test record shows higher mileage, or if the next test record is shows more than 100,000 additional miles on the odometer (indicating that the odometer had already rolled over, but the next check took this into account). The algorithm also attempts to correct for typos and entry errors. An odometer reading is flagged if it does not fit with surrounding readings for the same vehicle–either it is less than the previous reading or greater than the next–and cannot be explained by a rollover. The algorithm then tests whether fixing one of several common typos will make the flagged readings fit (e.g., moving the decimal over one place). If no correction will fit, the reading is replaced with the average of the surrounding readings. Finally, if after all our corrections any vehicle has an odometer reading above 800,000 or has implied VMT per day greater than 200 or less than zero, we exclude the vehicle from our analysis. All of our VMT analyses use this adjusted mileage.

Emissions results from smog checks are given in either parts per million (for HC and NO_x) or percent (O_2, CO, and CO_2). Without knowing the volume of air involved, there is no straightforward way to convert this to total emissions. Fortunately, as part of an independent evaluation of the Smog Check Program conducted in 2002-2003, Sierra Research Inc. and Eastern Research Group estimated a set of conversion equations to convert the proportional measurements of the ASM test to emissions in grams per mile traveled. These equations are reported in Morrow and Runkle (2005) and are reproduced below. The equations are for HCs, NO_x, and CO, and estimate grams per mile for each pollutant as a non-linear function of all three pollutants, model year, and vehicle weight. The equations for vehicles of up to model year 1990 are

$$FTP\_HC = 1.2648 \cdot exp(-4.67052 \quad +0.46382 \cdot HC^* + 0.09452 \cdot CO^* + 0.03577 \cdot NO^*$$
$$+0.57829 \cdot \ln(weight) - 0.06326 \cdot MY^*$$
$$+0.20932 \cdot TRUCK)$$

$$FTP\_CO = 1.2281 \cdot exp(-2.65939 \quad +0.08030 \cdot HC^* + 0.32408 \cdot CO^* + 0.03324 \cdot CO^{*2}$$
$$+0.05589 \cdot NO^* + 0.61969 \cdot \ln(weight) - 0.05339 \cdot MY^*$$
$$+0.31869 \cdot TRUCK)$$

$$FTP\_NOX = 1.0810 \cdot exp(-5.73623 \quad +0.06145 \cdot HC^* - 0.02089 \cdot CO^{*2} + 0.44703 \cdot NO^*$$
$$+0.04710 \cdot NO^{*2} + 0.72928 \cdot \ln(weight) - 0.02559 \cdot MY^*$$
$$-0.00109 \cdot MY^{*2} + 0.10580 \cdot TRUCK)$$

Where

$$HC^* = \ln((Mode1_{HC} \cdot Mode2_{HC})^{.5}) - 3.72989$$
$$CO^* = \ln((Mode1_{CO} \cdot Mode2_{CO})^{.5}) + 2.07246$$
$$NO^* = \ln((Mode1_{NO} \cdot Mode2_{NO})^{.5}) - 5.83534$$
$$MY^* = modelyear - 1982.71$$
$$weight = \text{Vehicle weight in pounds}$$
$$TRUCK = 0 \text{ if a passenger car, } 1 \text{ otherwise}$$

And for model years after 1990 they are:

$$
\begin{aligned}
FTP\_HC = 1.1754 \cdot exp(-6.32723 \quad & +0.24549 \cdot HC^{*} + 0.09376 \cdot HC^{*2} + 0.06653 \cdot NO^{*} \\
& +0.01206 \cdot NO^{*2} + 0.56581 \cdot \ln(weight) - 0.10438 \cdot MY^{*} \\
& -0.00564 \cdot MY^{*2} + 0.24477 \cdot TRUCK)
\end{aligned}
$$

$$
\begin{aligned}
FTP\_CO = 1.2055 \cdot exp(-0.90704 \quad & +0.04418 \cdot HC^{*2} + 0.17796 \cdot CO^{*} + 0.08789 \cdot NO^{*} \\
& +0.01483 \cdot NO^{*2} - 0.12753 \cdot MY^{*} - 0.00681 \cdot MY^{*2} \\
& +0.37580 \cdot TRUCK)
\end{aligned}
$$

$$
\begin{aligned}
FTP\_NOX = 1.1056 \cdot exp(-6.51660 \quad & +0.25586 \cdot NO^{*} + 0.04326 \cdot NO^{*2} + 0.65599 \cdot \ln(weight) \\
& -0.09092 \cdot MY^{*} - 0.00998 \cdot MY^{*2} + 0.24958 \cdot TRUCK)
\end{aligned}
$$

Where:

$$
\begin{aligned}
HC^{*} &= \ln((Mode1_{HC} \cdot Mode2_{HC})^{.5}) - 2.32393 \\
CO^{*} &= \ln((Mode1_{CO} \cdot Mode2_{CO})^{.5}) + 3.45963 \\
NO^{*} &= \ln((Mode1_{NO} \cdot Mode2_{NO})^{.5}) - 3.71310 \\
MY^{*} &= modelyear - 1993.69 \\
weight &= \text{Vehicle weight in pounds} \\
TRUCK &= 0 \text{ if a passenger car, } 1 \text{ otherwise}
\end{aligned}
$$

C Steps to Clean DMV Data

We deal with two issues associated with DMV data. The main issue is that DMV entries for the same addresses will often have slightly different formats. For example, "12 East Hickory Street" may show up as "12 East Hickory St," "12 E. Hickory St.", etc. To homogenize the entries, we input each of the DMV entries into mapquest.com and then replace the entry with the address that mapquest.com gives.

Second, the apartment number is often missing in DMV data. Missing apartment numbers has the effect of yielding a large number of vehicles in the same "location." We omit observations that have over seven vehicles in a given address or more than three last names of registered owners.

D Robustness Checks

In this appendix, we report the results of several robustness checks to our main results on the intensive margin. Table A.2 reports elasticities by quartile for all five categories of externality.

Our base specification controls for the fixed effect of a each NO_x quartile on miles traveled. One might be concerned, however, that variation in dollars per mile (DPM) might be correlated with other characteristics such as age, odometer, and demographics, and that the DPM-quartile interactions may be picking up this correlation, rather than true heterogeneity. To test for this, in Table A.5 we present results with vehicle fixed effects and interactions between NO_x quartiles and various control variables. Adding these interaction terms actually makes the heterogeneity in the effect of DPM more pronounced. The final column in Table A.5 includes month-by-year fixed effects, therefore allowing for a completely flexible time trend. The degree of heterogeneity increases when we include these fixed effects.

Table A.6 repeats the same exercise, but uses levels rather than logs of DPM as the variable of interest. The results are qualitatively similar, with substantial heterogeneity in every specification. However, with a log-linear specification we do not observe the cleanest vehicles having a positive coefficient.

We also investigate the functional forms of these relationships in a semi-parametric way. For each externality, we define vehicles by their percentile of that externality. We then estimate Equation (8) with separate elasticities for vehicles falling in the zero to first percentile, first to second, etc. Appendix Figure A.2 plots a LOWESS smoothed line through these 100 separate elasticity estimates. For the three criteria pollutants, we find that the relationship is quite linear with the elasticity being positive for the cleanest 10 percent of vehicles. The dirtiest vehicles have elasticities that are roughly 0.4. For fuel economy, the relationship is fairly linear from the 60th percentile onwards, but begins steeply and flattens out from the 20th percentile to the 40th. The elasticity of the lowest fuel economy vehicles is nearly 0.6. To put these numbers into context across the different years, the average fuel economy of the 20th percentile is 18.7, while the average for the 40th percentile is 21.75. The variation in elasticities across weight is not monotonic. The relationship begins by increasing until roughly the 20th percentile, and then falls more or less linearly thereafter. The elasticity of the heaviest vehicles is roughly 0.3.

Note that the roughly linear relationship between criteria pollutant emissions and the elasticity is not due to "over smoothing." Appendix Figure A.3 plots the LOWESS smoothed lines for HCs under different bandwidths. The top left figure simply reports the 100 elasticities. There is some evidence that the relationship is not monotonic early on, but from the 5th percentile on, the relationship appears monotonic. Doing this exercise for the other criteria pollutants yields similar results.

E Details of the Gasoline Tax Policy Simulation

For the intensive margin, we estimate a regression as in column 6 of Tables 2 and A.2, except that we interact ln(DPM) with quartile of fuel economy, vehicle weight, and emissions of HC, NO_x, and CO, and dummies for vehicle age bins, again using bins of 4-9, 10-15, and 16-29 years, and control for the direct effects of quartiles of HC, NO_x, and CO emissions. We use quartiles calculated by year and age bin. The coefficients are difficult to interpret on their own, and too numerous to list. However, most are statistically different from zero, and the exceptions are due to small point estimates, not large standard errors.

As in Section G, we compress our dataset to have at most one observation per vehicle per year. Each vehicle is then assigned an elasticity based on its quartiles and age bin. Vehicle i's VMT in the counterfactual with an additional \$1 tax on gasoline is calculated by:

$$VMT^i_{counterfactual} = VMT^i_{BAU} \cdot \left(\frac{P_i + 1}{P_i} \cdot \beta_i \right),$$

where VMT^i_{BAU} is vehicle i's actual average VMT per day between its current and previous smog check, P_i is the average gasoline price over that time, and β_i is the elasticity for the fuel economy/weight/HC/NO/CO/age cell to which i belongs.

For the extensive margin, we estimate a Cox regression on the hazard of scrappage for vehicles 10 years and older, stratifying by VIN prefix and interacting DPM with all five type of quartiles and age bins 10-15 and 16-29. Similar to the intensive margin, we assign each vehicle a hazard coefficient based on its quartile-age cell. Cox coefficients can be transformed into hazard ratios, but to simulate the affect of an increase in gasoline prices on the composition of the vehicle fleet, we must convert these into changes in total hazard.

To do this, we first calculate the actual empirical hazard rate for prefix k in year t as:

$$OrigHazard_{kt} = \frac{D_{kt}}{R_{kt}},$$

where D_{kt} is the number of vehicles in group k, that are scrapped in year t, and R_{kt} is the number of vehicle at risk (that is, which have not previously been scrapped or censored). We then use the coefficients from our Cox regression to calculate the counterfactual hazard faced by vehicles of prefix k in quartile-age group q during year t as:[29]

$$NewHazard_{qkt} = OrigHazard_{kt} \cdot \exp\left\{\frac{1}{MPG_k} \cdot \gamma_q\right\},$$

where MPG_k is the average fuel economy of vehicle of prefix k and γ_q is the Cox coefficient associated with quartile group q. We then use the change in hazard to construct a weight H_{qkt} indicating the probability that a vehicle of prefix k in quartile group q in year t would be in the fleet if a \$1 gasoline tax were imposed. Weights greater than 1 are possible, which should be interpreted as a $H_{qkt} - 1$ probability that another vehicle of the same type would be on the road, but which was scrapped under "Business as Usual." Because the hazard is the probability of scrappage in year t, conditional on survival to year t, this weight must be calculated interactively, taking into account the weight the previous year. Specifically, we have:

$$H_{qkt} = \prod_{j=1998}^{t} (1 - (NewHazard_{qkj} - OrigHazard_{kt})).$$

We also assign each vehicle in each year a population weight. This is done both to scale our estimates up to the size of the full California fleet of personal vehicles, and to account for the ways in which the age composition of the smog check data differs from that of the fleet. We construct these weights using the vehicle population estimates contained in CARB's EMFAC07 software, which are given by year, vehicle age, and truck status. Our population weight is the number of vehicles of a given age and truck status in a each year given by

[29] Note that age group is determined by model-year and year.

EMFAC07, divided by the number of such vehicle appearing in our sample. For instance, if EMFAC07 gave the number of 10-year-old trucks in 2005 as 500, while our data contained 50, each 10-year-old truck in our data would have a population weight of 10. Denote the population weight by P_{tac}, where t is year, a is age, and c is truck status.

There is an additional extensive margin that we have not estimated in this paper: new car purchases. To ensure that the total vehicle population is accurate, we apply an *ad hoc* correction based on Busse et al. (forthcoming), who find that a \$1 increase in gasoline prices would decrease new car sales by 650,000 per year. Because California's vehicle fleet makes up about 13 percent of the national total, we decrease the population of model years 1998 and later by 84,500 when constructing the population weight for the counterfactual. We apply 40 percent of the decrease to trucks, and 60 percent to passenger cars. Denote the "new car effect" n_c.

We estimate the total annual emissions by passenger vehicle in California of NO_x, HC, CO, and CO_2 as actually occurred, and under a counterfactual where a \$1 gasoline tax was imposed in 1998. Let i denote a vehicle, a vehicle age, c truck status. Then the annual emissions of pollutant p in year t under "business as usual" are:

$$Emission^{pt}_{BAU} = \sum_i P_{tac} \cdot VMT^i_{BAU} \cdot r_i(p) \cdot 365,$$

and under the counterfactual they are:

$$Emission^{pt}_{counterfactual} = \sum_i (P_{tac} - 1(\text{model year} >= 1998) \cdot n_c) \cdot H_{qkt} \cdot VMT^i_{counterfactual} \cdot r_i(p) \cdot 365,$$

where $r_i(p)$ is the emissions rate per mile of pollutant p for vehicle i. For NO_x, HC, and CO, this is the last smog check reading in grams per mile, while for CO_2 this is the vehicle's gallons per mile multiplied by 19.2 pounds per gallon.

F California versus the Rest of the United States

Given that our empirical setting is California, it is natural to ask whether our results are representative of the country as a whole. At the broadest level, the local-pollution benefits from carbon pricing are a function of the per-capita number of miles driven, the emission characteristics of the fleet of vehicles, and the marginal damages of the emissions. We present evidence that the benefits may, in fact, be larger outside of California. The reason for this is that while the marginal damages are indeed larger in California, the vehicle stock in California is much cleaner than the rest of the country because California has traditionally led the rest of the U.S. in terms of vehicle-emission standards.

The results in Muller and Mendelsohn (2009) provide a convenient way to test whether California differs in terms of marginal damages. Table A.12 presents points on the distribution of marginal damages for NO_x, HCs, and the sum of the two, weighted by each county's annual VMT.[30] Figure A.6 plots the kernel density estimates of the distributions. We present the sum of because counties are typically either "NO_x constrained" or "VOC (HC) constrained," and the sum is perhaps more informative. As expected, the marginal damages are higher in California for HCs, but lower for NO_x, as California counties tend to be VOC-constrained. The sum of the two marginal damages is 78 percent higher in California. Higher points in the distribution show an even larger disparity.

Larger marginal damages are offset, however, by the cleaner vehicle stock within California—a result of California's stricter emission standards. To illustrate this, we collected county-level average per-mile emission rates for NO_x, HCs, and CO from the EPA Motor Vehicle Emission Simulator (MOVES). MOVES reports total emissions from transportation and annual mileage for each county. Table A.12 also presents points on the per-mile emissions, and Figure A.7 plots the distributions.[31] Mean county-level NO_x, HCs, and CO are 67, 36, and 31 percent lower in California, respectively. Other points in the distributions exhibit similar patterns.

Finally, we calculate the county-level average per-mile externality for each pollutant, as well as the sum of the three. Table A.12 and Figure A.8 illustrates these. As expected, the HC damages are higher, but the average county-level per-mile externality from the sum of the three pollutants is 30 percent lower in California than the rest of the country; the 25th percentile, median, and 75th percentile are 35, 30, and 9 percent lower, respectively. These calculations suggest that, provided the average VMT elasticities are not significantly

[30] All of the points on the distribution and densities discussed in this section weight each county by its total VMT.

[31] We note that the emissions reported in MOVES exceed the averages in our data. This may reflect the fact that smog checks are not required for vehicles with model years before 1975, and these vehicles likely have very high emissions because this pre-dates many of the emission standards within the U.S.

smaller outside of California and/or the heterogeneity across vehicle types is not significantly different (in the reverse way), our estimates are likely to apply to the rest of the country.

G Scrappage Decisions

Our next set of empirical models examines how vehicle owners' decisions to scrap their vehicles due to gasoline prices. Again we will also examine how this effect varies over emissions profiles.

We determine whether a vehicle has been scrapped using the data from CARFAX Inc. We begin by assuming that a vehicle has been scrapped if more than a year has passed between the last record reported to CARFAX and the date when CARFAX produced our data extract (October 1, 2010). However, we treat a vehicle as being censored if the last record reported to CARFAX was not in California, or if more than a year and a half passed between the last smog check in our data and that last record. As well, to avoid treating late registrations as scrappage, we treat all vehicles with smog checks after 2008 as censored. Finally, to be sure we are dealing with scrapping decisions and not accidents or other events, we only examine vehicles that are at least 10 years old.

Some modifications to our data are necessary. To focus on the long-term response to gasoline prices, our model is specified in discrete time, denominated in years. Where vehicles have more than one smog check per calendar year, we use the last smog check in that year. Also, because it is generally unlikely that a vehicle is scrapped at the same time as its last smog check, we create an additional observation for scrapped vehicles either one year after the last smog check, or six months after the last CARFAX record, whichever is later. For these created observations, odometer is imputed based on the average VMT between the last two smog checks, and all other variables take their values from the vehicle's last smog check. An exception is if a vehicle fails the last smog check in our data. In this case, we assume the vehicle was scrapped by the end of that year.

Because many scrapping decisions will not take place until after our data ends, a hazard model is needed to deal with right censoring. Let T_{jivg} be the year in which vehicle i, of vehicle type j, vintage v, and geography g, is scrapped. Assuming proportional hazards, our basic model is:

$$\Pr[t < T_{ijvg} < t+1 | T > t] = h^0_{jv}(t) \cdot \exp\{\beta x DPM_{igt} + \gamma D_{fail_{it}} + \psi G_{igt} + \alpha X_{it}\},$$

where DPM_{igt} is defined as before; $D_{fail_{it}}$ is a dummy equal to one if the vehicle failed a smog check any time during year t; G is a vector of demographic variables, determined by the location of the smog check; X is a vector of vehicle characteristics, including a dummy for truck and a sixth-order polynomial in odometer; and $h^0_{ijv}(t)$ is the baseline hazard rate, which varies by time but not the other covariates. In some specifications, we will allow each vehicle type and vintage to have its own baseline hazard rate.

We estimate this model using semi-parametric Cox proportional hazards regressions, leaving the baseline hazard unspecified. We report exponentiated coefficients, which may be interpreted as hazard ratios. For instance, a 1 unit increase in DPM will multiply the hazard rate by $\exp\{\beta\}$, or increase it by $(\exp\{\beta\} - 1)$ percent. In practice, we scale the coefficients on DPM for a 5-cent change, corresponding to a \$1.00 increase in gasoline prices for a vehicle with fuel economy of 20 miles per gallon.

Tables A.3 and A.4 show the results of our hazard analysis. Models 1 and 2 of Table A.3 assign all vehicles to the same baseline hazard function. Model 1 allows the effect of gasoline prices to vary by whether or not a vehicle failed a smog check. Model 2 also allows the effect of gasoline prices to vary by quartiles of NO_x.[32] Models 3 and 4 are similar, but stratify the baseline hazard function, allowing each VIN prefix to have its own baseline hazard function. Model 5 allows the effect of gasoline prices to vary both by externality quartile and age group, separating vehicles 10 to 15 years old from vehicles 16 years and older.

Models 1 and 2 indicate that increases in gasoline prices actually decrease scrapping on average, with the cleanest vehicles seeing the largest decreases. The effect is diminished once unobserved heterogeneity among vehicle types is controlled for, but is still statistically significant. However, the true heterogeneity in the effect of gasoline prices on hazard seems to be over age groups. Model 5 shows that when the cost of driving a mile increases by five cents, the hazard of scrappage decreases by about 23 percent for vehicles between 10 and 15 years old, while it increases by around 3 percent for vehicles age 16 and older, with little

[32] Quartiles in these models are calculated by year among only vehicles 10 years and older.

variation across NO_x quartiles within age groups. These results suggests that when gasoline prices rise, very old cars are scrapped, increasing demand for moderately old cars and thus reducing the chance that they are scrapped.

Table A.4 presents the quartile by age by DPM interactions for each of the 5 externality dimensions. Hydrocarbons and CO have the identical pattern to NO_x, with no heterogeneity within age-group. With fuel economy and vehicle weight, there is within-age heterogeneity, although the form is counter-intuitive. The heaviest and least fuel-efficient vehicles are relatively less likely thank the lightest and most fuel-efficient vehicles to be scrapped when gasoline prices increase. That is, while all 10- to 15-year-old vehicles are less likely to be scrapped, the decrease in hazard rate is larger for heavy, gas-guzzling vehicles. For vehicles 16 years and older, the heaviest quartile is less likely to be scrapped when gasoline prices increase, even though the lightest (and middle quartiles) are more likely. As the model stratifies by VIN prefix, this cannot be simply that more durable vehicles have lower fuel economy.

In summary, increases in the cost of driving a mile over the long term increase the chance that old vehicles are scrapped, while middle-aged vehicles are scrapped less, perhaps because of increased demand. Although vehicle age is highly correlated with emissions of criteria pollutants, there is little variation in the response to gasoline prices across emissions rates within age groups.

H Income Distribution Adjustment

In section 7, we assign income brackets to individual consumers using the method of Borenstein (2012). Here we briefly describe the details of that procedure; for more details see Borenstein (2012).

Borenstein shows that household consumption levels of some commodity (gasoline in our case) within income brackets can be bounded between the case where the ranking of household incomes is sorted by consumption levels (usage-ranking), and the case where the ranking of household incomes is random with respect to consumption levels (random ranking). If one can calculate the average consumption by income bracket, one can calculate a weighting between usage-ranking and random-ranking that correctly assigns households to income brackets based on their consumption. Borenstein proposes calculating these averages from a separate dataset that contains individual level income and consumption, if one can be found. We utilize the 2009 NHTS for this purpose.

Formally, let $\bar{g}_b$ denote the average gasoline consumption for consumers living in California in income group b in the 2009 NHTS. The N vehicles registered in each census block group (CBG) in California and appearing in the Smog Check data are to be assigned an integer rank from 1 to N, intended to correspond to the income ranking of the household those vehicles belong to. If s_b^c denotes the number of households falling into income bracket b in CBG c in the 2000 Census, and h_b denotes the number of vehicles per household in income bracket b, then, for instance, vehicles ranked from 1 to $fracNs_b * h_b$ will fall into bracket 1. The ranking for vehicle i will be $v_i(w) = (1-w) \cdot r_{rr} + w \cdot r_{ur}$, where r_{rr} is drawn randomly from a uniform distribution over $[1, N]$, producing random-ranking, while r_{ur} sorts vehicles by gasoline consumption, producing usage ranking. Any choice of w will produce a joint ranking within CBGs, leading to a statewide average within-bracket gasoline consumption level of $\tilde{g}_b$ within the Smog Check data.[33]

The income brackets given at the CBG level in the 2000 census can be pooled into groups roughly approximating deciles of the total income distribution in California. The NHTS gives income in brackets as well, which can be grouped into 8 groups corresponding to the first 7 "deciles" in the Census data, plus the top 3 deciles topcoded into one income bracket. We calculate w based on the 8 groupings in the NHTS data, but when using that w to assign vehicles to income brackets, we use the ranking implied by w to distribute vehicles across the top 3 deciles. We choose w to minimize the following goodness-of-fit measure:

$$G = \sum_{b=1}^{8} s_b(\tilde{g}_b - \bar{g}_b)$$

That is, we choose w such that when vehicles in the Smog Check Data are ranked into income brackets, the average gasoline usage in each income bracket matches the average gasoline usage for that income bracket in the 2009 NHTS.

[33] Ideally this calculation would use a CBG-specific w, however the NHTS does not provide geographic data at that level. Borenstein (2012) has the same limitation.

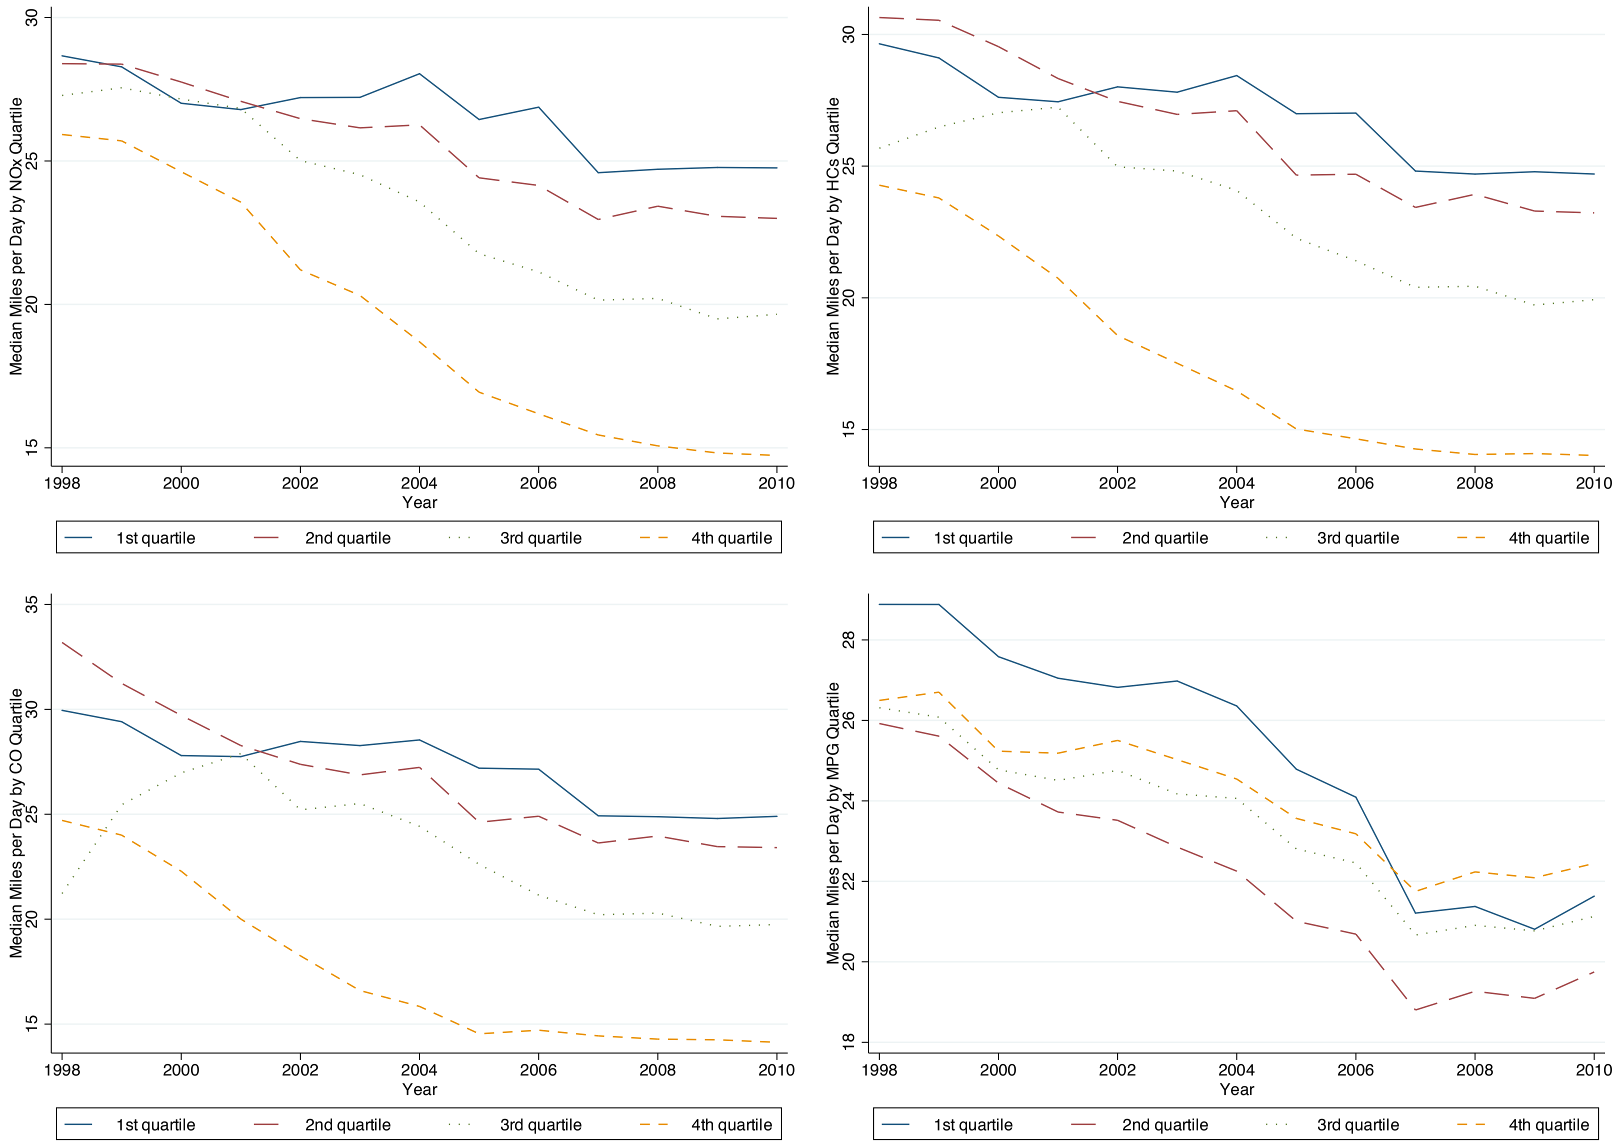

Figure A.1: Change in VMT over sample by pollutant quartile

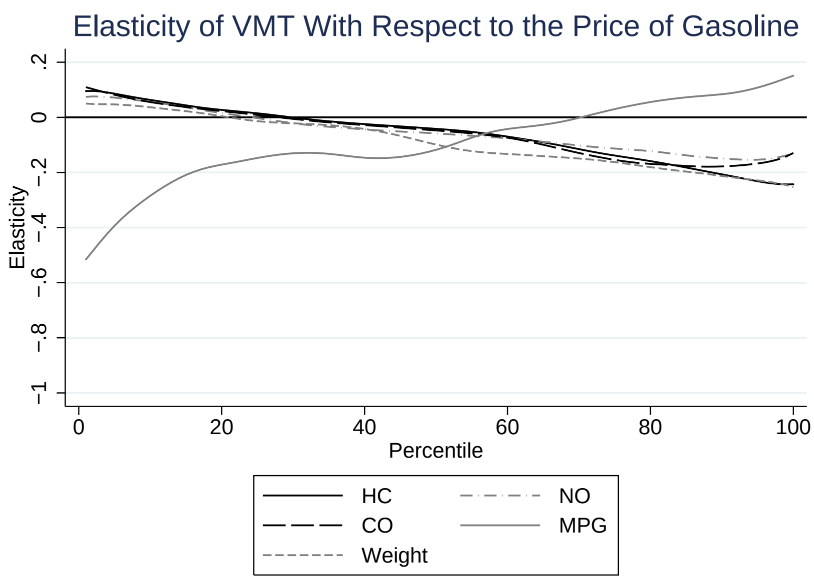

Figure A.2: Non-parametric relationships between elasticity and externality

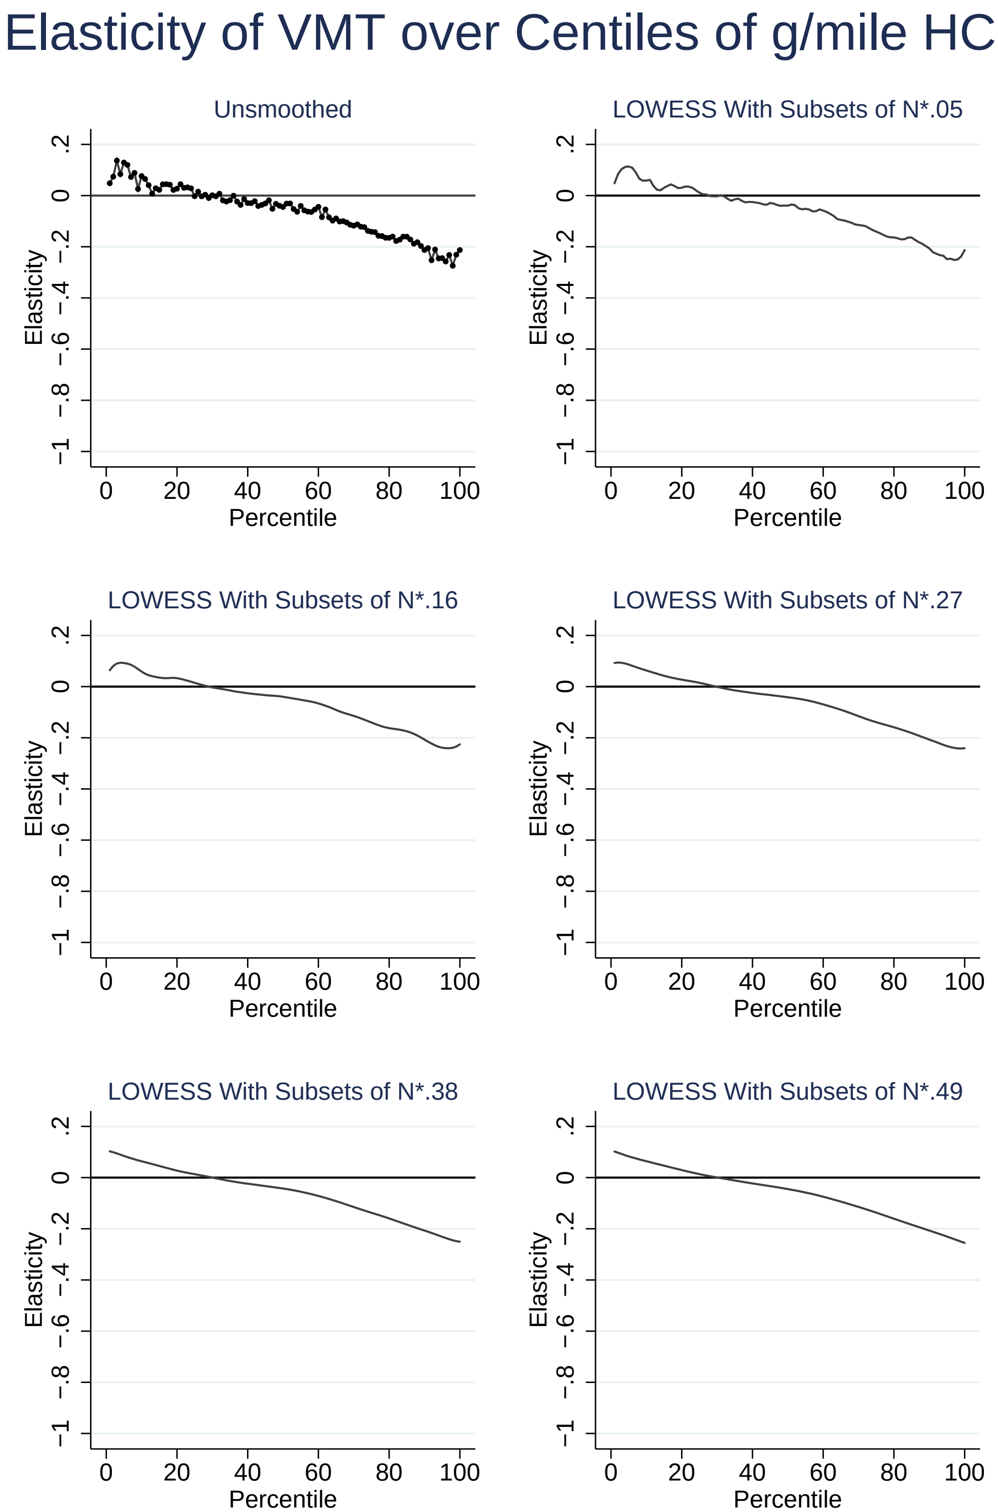

Figure A.3: The effect of bandwidth on the non-parametric function

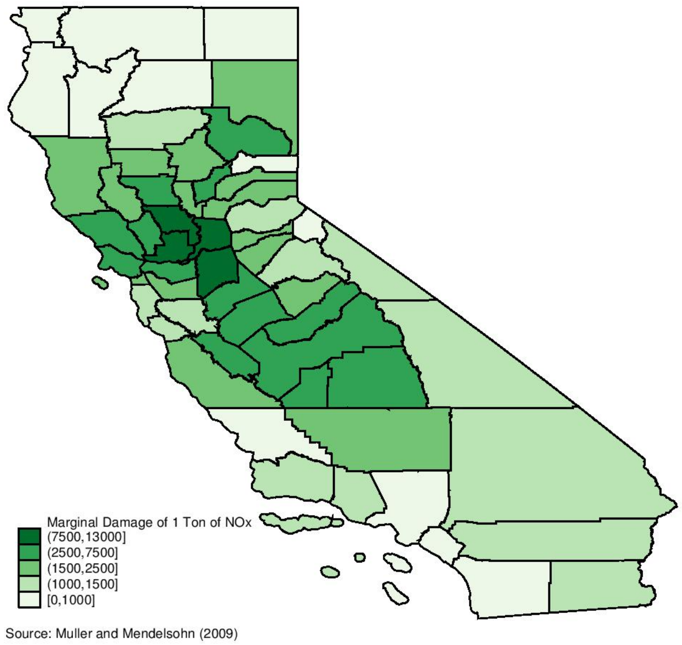

(a) NO_x Damage

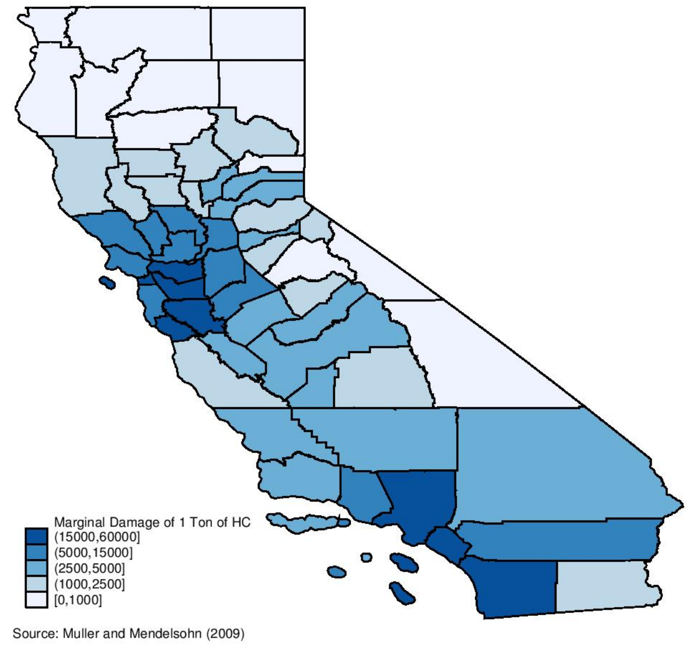

(b) HC Damage

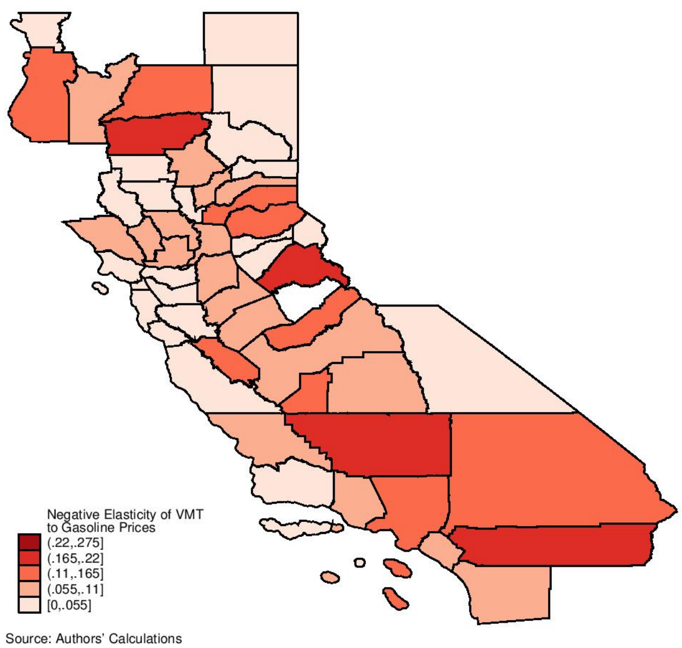

(c) VMT Elasticty

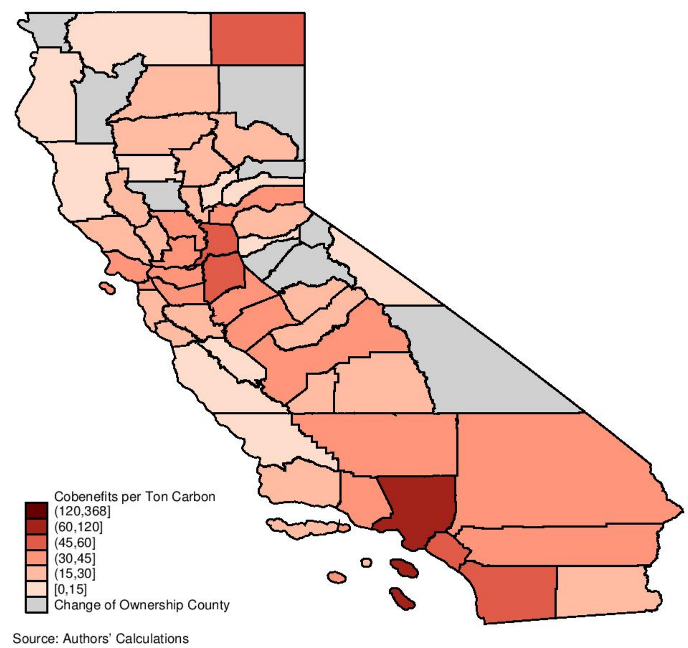

(d) Estimated Co-benefits

Figure A.4: Social damages of pollution, VMT elasticity, and local-pollution benefits of a gasoline tax, by county

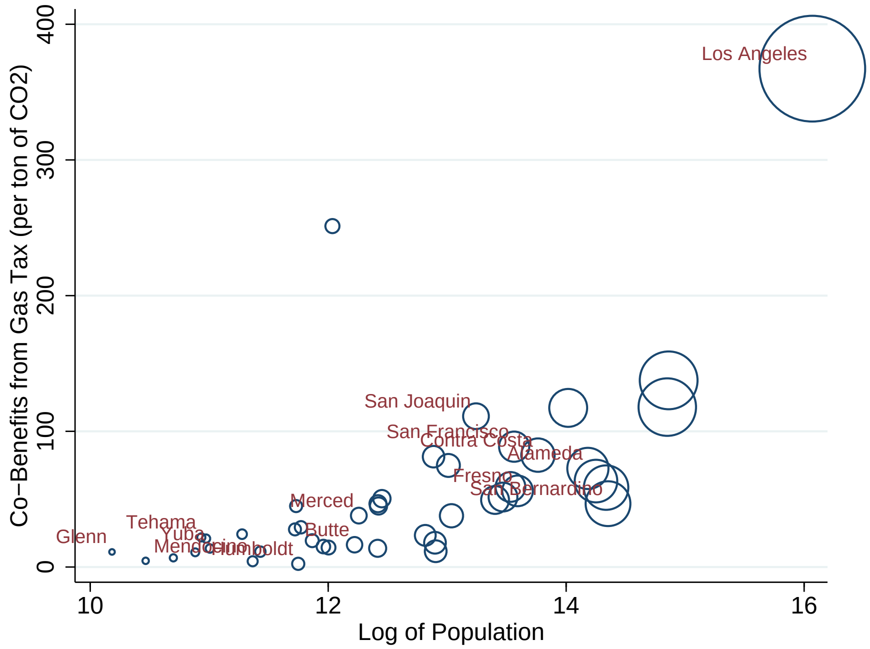

Figure A.5: Estimated county-level local-pollution benefits versus the log of county population

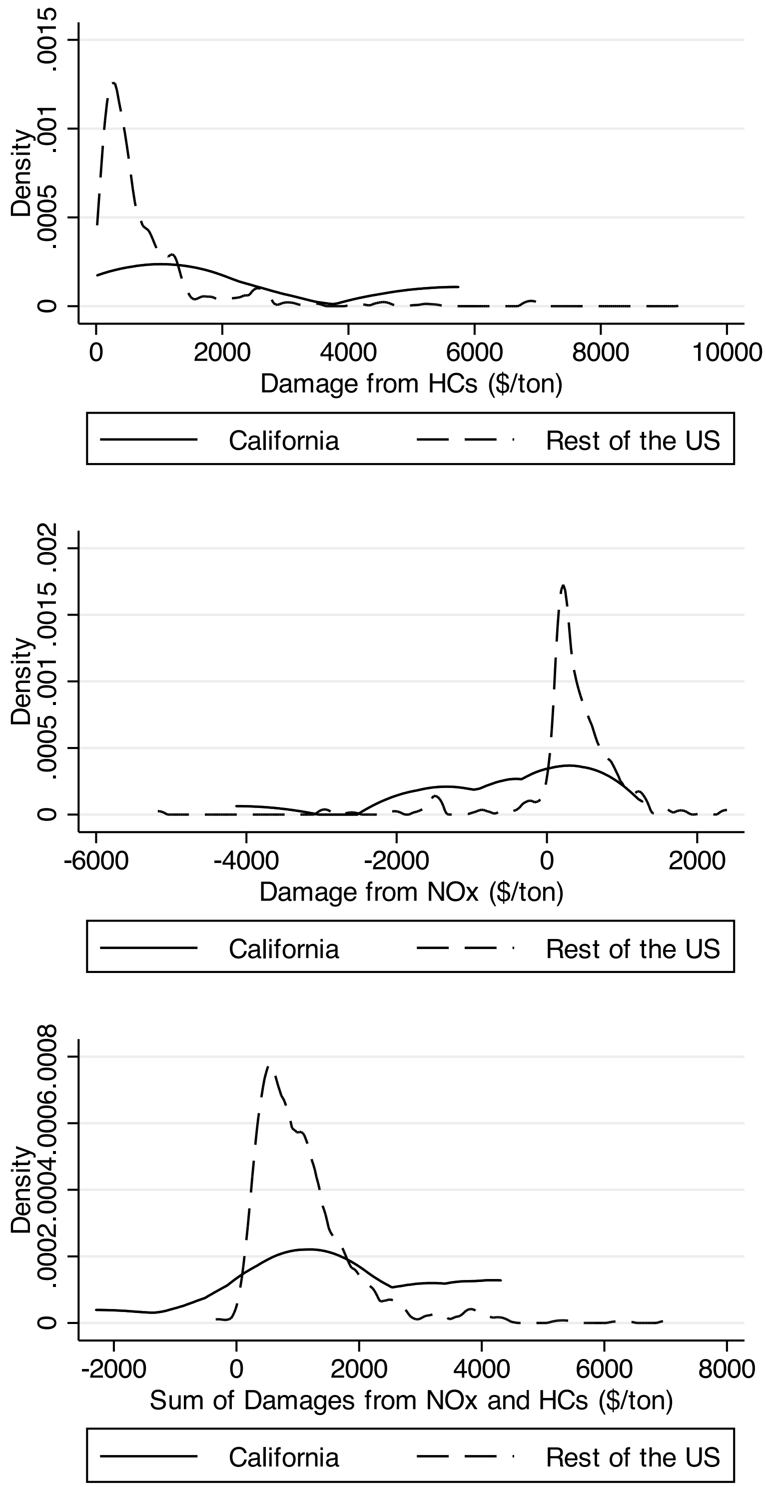

Figure A.6: Distributions of marginal damages from Muller and Mendelsohn (2009) for California and the rest of the U.S.

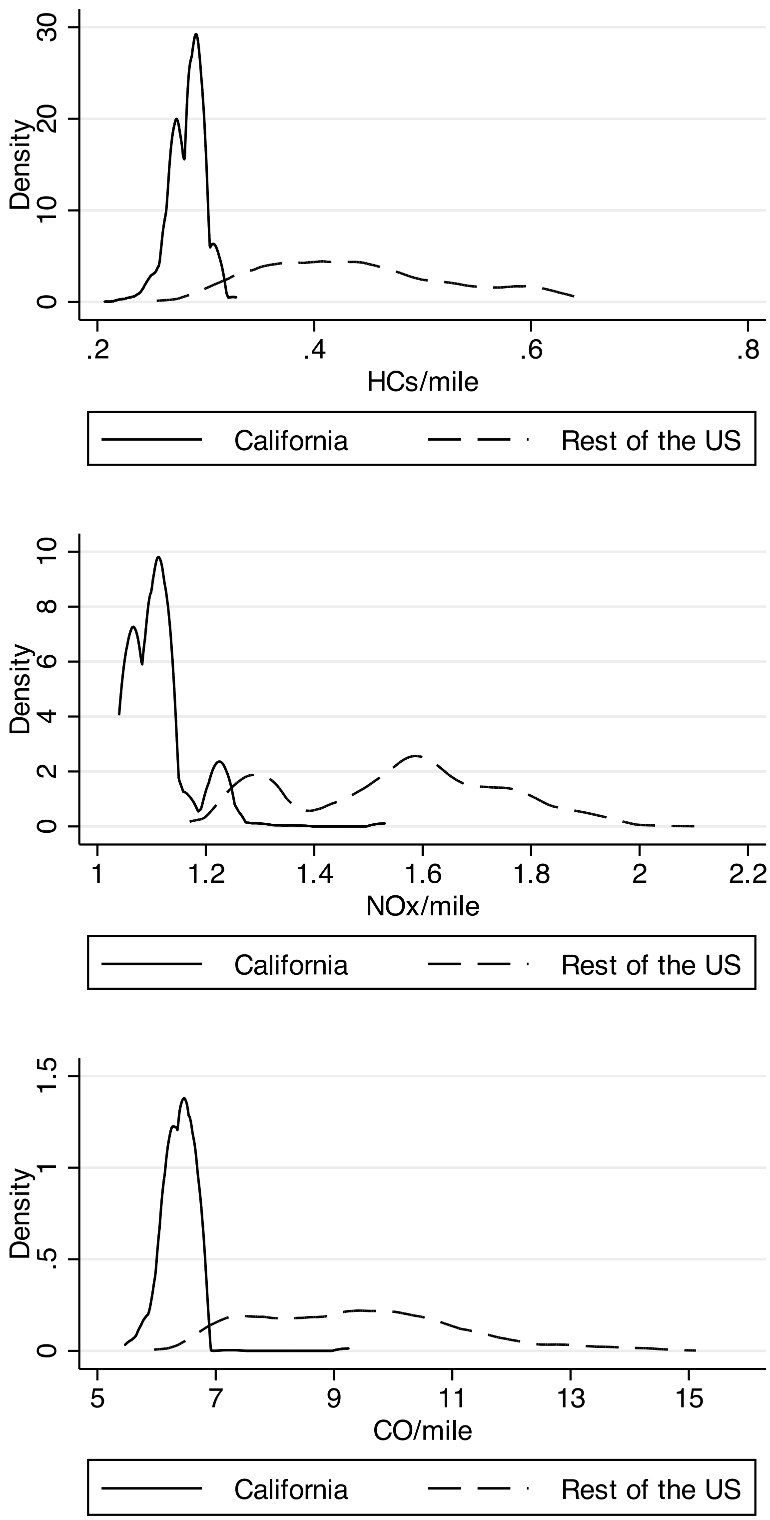

Figure A.7: Distributions of per-mile emissions for California and the rest of the U.S.

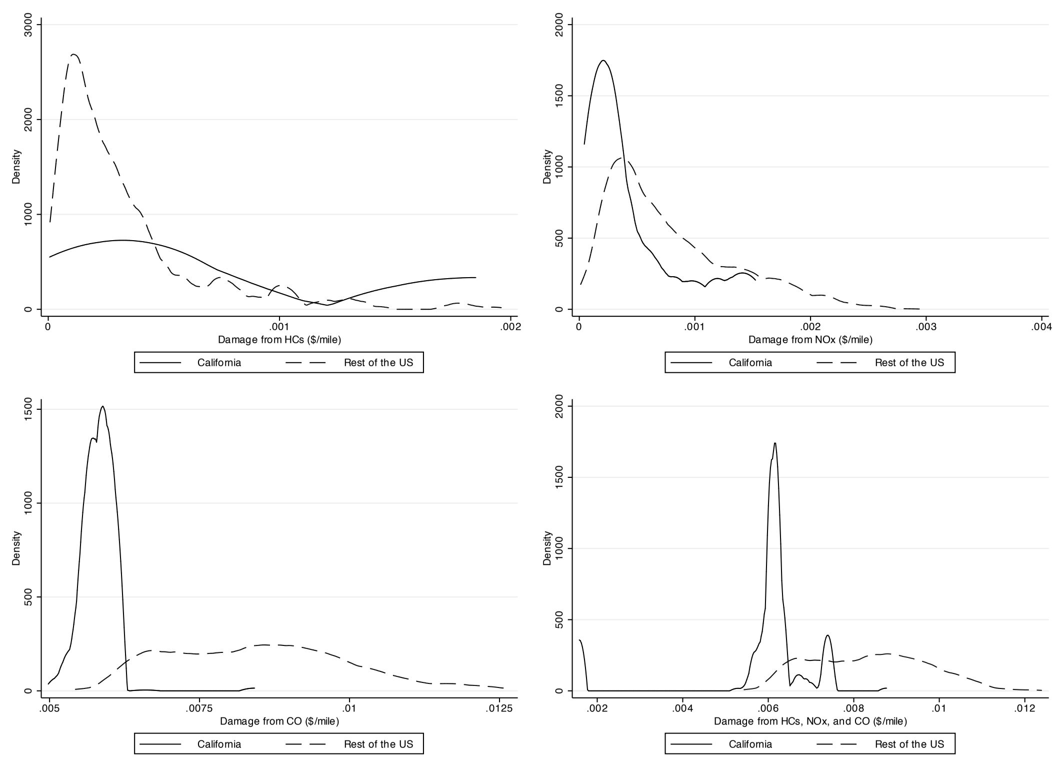

Figure A.8: Distributions of per-mile damages for California and the rest of the US

Table A.1: Average Pollutant Rates Per Mile Traveled by Year

	Nitrogen Oxides			Hydrocarbons			Carbon Monoxide			Gasoline	
Year	Mean	SD	Mean CV	Mean	SD	Mean CV	Mean	SD	Mean CV	Mean	SD
1998	1.161	1.051	0.536	1.662	1.866	0.504	15.400	24.739	0.516	0.043	0.010
1999	1.187	0.983	0.455	1.665	1.864	0.464	15.227	24.952	0.485	0.043	0.010
2000	1.094	0.915	0.441	1.535	1.826	0.460	13.539	23.849	0.476	0.044	0.010
2001	0.982	0.857	0.427	1.354	1.769	0.464	11.689	23.123	0.461	0.044	0.010
2002	0.876	0.816	0.418	1.145	1.679	0.446	9.694	21.381	0.430	0.044	0.010
2003	0.791	0.780	0.401	0.997	1.563	0.432	7.940	19.503	0.395	0.045	0.010
2004	0.715	0.742	0.380	0.855	1.469	0.421	6.561	17.581	0.363	0.045	0.010
2005	0.735	0.713	0.393	0.852	1.444	0.455	6.375	17.519	0.379	0.045	0.010
2006	0.638	0.667	0.382	0.718	1.351	0.430	5.157	15.887	0.350	0.045	0.010
2007	0.572	0.634	0.377	0.628	1.261	0.431	4.308	14.509	0.334	0.045	0.010
2008	0.512	0.602	0.373	0.545	1.185	0.400	3.556	13.064	0.317	0.046	0.010
2009	0.478	0.590	0.379	0.496	1.148	0.412	3.120	12.147	0.316	0.046	0.011
2010	0.462	0.566	0.402	0.460	1.002	0.427	2.741	10.901	0.323	0.046	0.010
N	10432374			10432374			10666348			13397795	

Note: Mean CV is the average VIN Prefix-level coefficient of variation (SD/Mean). Gasoline is measured in gallons per mile, while the remaining pollutant rates are measured in grams per mile.

Table A.2: Vehicle Miles Travelled, Dollars Per Mile, and Externality Quartiles

Quartile	Nitrogen Oxides	Hydrocarbons	Carbon Monoxide	Fuel Economy	Vehicle Weight
1	0.0425	0.0486	0.0466	-0.169	-0.111
2	-0.0540	-0.0550	-0.0527	-0.159	-0.114
3	-0.152	-0.149	-0.149	-0.104	-0.145
4	-0.280	-0.305	-0.307	-0.0986	-0.167

Coefficients are elasticities calcualted by regressing the log of average daily VMT between Smog Checks on the log of the gas price in dollars per mile, interacted with quartiles of the pollutants indicated. Quartiles are based on rankings of within the calendar year in which the Smog Check occurs. All regressions controll for direct effects of the quartiles, a quadratic time trend, demographics of the zip code where the Smog Check occurs, calendar-year fixed effects, vehicle age fixed effects, and vehicle fixed effects.

Table A.3: Hazard of Scrappage: Cox Proportional Hazard Model

	Model 1	Model 2	Model 3	Model 4	Model 5
Dollars per Mile	0.920*		0.965*		
	(0.039)		(0.018)		
DPM * Failed Smog Check	1.105**	1.074**	1.063**	1.043*	
	(0.029)	(0.026)	(0.021)	(0.020)	
Failed Last Smog Check	7.347**	7.800**	7.639**	8.155**	
	(0.242)	(0.246)	(0.161)	(0.173)	
DPM * NO Quartile 1		0.801**		0.893**	
		(0.044)		(0.035)	
DPM * NO Quartile 2		0.862**		0.923**	
		(0.038)		(0.026)	
DPM * NO Quartile 3		0.883**		0.956*	
		(0.034)		(0.017)	
DPM * NO Quartile 4		0.929*		0.983	
		(0.033)		(0.010)	
Vehicle Ages 10-15					
DPM * NO Quartile 1					1.285**
					(0.023)
DPM * NO Quartile 2					1.287**
					(0.018)
DPM * NO Quartile 3					1.291**
					(0.014)
DPM * NO Quartile 4					1.254**
					(0.012)
Failed Smog Check					8.732**
					(0.380)
DPM * Failed Smog Check					0.910**
					(0.014)
Vehicle Ages 16+					
DPM * NO Quartile 1					0.751**
					(0.016)
DPM * NO Quartile 2					0.745**
					(0.014)
DPM * NO Quartile 3					0.745**
					(0.011)
DPM * NO Quartile 4					0.737**
					(0.008)
Failed Smog Check					7.765**
					(0.286)
DPM * Failed Smog Check					1.185**
					(0.026)
Station ZIP Code Characteristics	Yes	Yes	Yes	Yes	Yes
Quadratic Time Trend in Days	Yes	Yes	Yes	Yes	Yes
Vehicle Characteristics	Yes	Yes	Yes	Yes	Yes
Quartiles of NO	No	Yes	No	Yes	Yes
Stratified on Vin Prefix	No	No	Yes	Yes	Yes
Observations	31567473	26720283	31567473	26720283	26720283

Note: Coefficients on dollars per mile scaled for a 5-cent change

Table A.4: Effect of a 5-cent/mile Increase in Driving Cost on the Hazard of Scrappage

	Nitrogen Oxides	Hydrocarbons	Carbon Monoxide	Fuel Economy	Vehicle Weight
Vehicle Ages 10-15					
Quartile 1	1.285	1.297	1.301	1.108	1.516
Quartile 2	1.287	1.288	1.302	1.237	1.349
Quartile 3	1.291	1.286	1.272	1.479	1.265
Quartile 4	1.254	1.236	1.240	1.707	1.174
Vehicle Ages 16+					
Quartile 1	0.751	0.754	0.749	0.674	0.866
Quartile 2	0.745	0.753	0.763	0.735	0.773
Quartile 3	0.745	0.757	0.745	0.797	0.715
Quartile 4	0.737	0.751	0.751	0.822	0.629

Statistics are exponentiated coeficients of a Cox proportional hazards model. Interpret as hazard ratios.

Table A.5: Robustness Check—Intensive Margin Interacting NOx Quartiles With Other Controls

	(1)	(2)	(3)	(4)	(5)	(6)
ln(DPM) * NO Q1	0.0406	0.0381	0.0678*	0.0605	0.0590	0.0666
	(0.0231)	(0.0250)	(0.0339)	(0.0335)	(0.0333)	(0.121)
ln(DPM) * NO Q2	-0.0617*	-0.0581*	-0.0453	-0.0478	-0.0484	-0.0410
	(0.0261)	(0.0269)	(0.0309)	(0.0310)	(0.0308)	(0.121)
ln(DPM) * NO Q3	-0.158***	-0.155***	-0.166***	-0.165***	-0.165***	-0.157
	(0.0271)	(0.0272)	(0.0282)	(0.0291)	(0.0294)	(0.120)
ln(DPM) * NO Q4	-0.288***	-0.298***	-0.355***	-0.353***	-0.351***	-0.344**
	(0.0300)	(0.0302)	(0.0325)	(0.0332)	(0.0331)	(0.120)
NO Q2	0.378	0.327	-2.622	-3.925*	-3.954*	-4.916**
	(0.800)	(0.735)	(1.622)	(1.693)	(1.673)	(1.732)
NO Q3	-1.246	-1.447	-5.233***	-6.846***	-6.793***	-7.987***
	(1.012)	(0.899)	(1.447)	(1.524)	(1.508)	(1.566)
NO Q4	-2.297*	-2.951**	-9.696***	-11.39***	-11.26***	-12.60***
	(1.116)	(1.084)	(2.257)	(2.253)	(2.271)	(2.301)
Quartile-Time Trend Interactions	Yes	Yes	Yes	Yes	Yes	Yes
Vintage-Quartile Interactions	No	Yes	Yes	Yes	Yes	Yes
Quartile-Year Interactions	No	No	Yes	Yes	Yes	Yes
Quartile-Lagged Odometer Interactions	No	No	No	Yes	Yes	Yes
Quartile-Demographics Interactions	No	No	No	No	Yes	Yes
Calendar Month Fixed-Effects	No	No	No	No	No	Yes
N	2979289	2979289	2979289	2979289	2979289	2979289

Note: All regressions include vehicle fixed-effects, year fixed effects, vintage/truck fixed effects, a quadratic time trend, a sixth order polynomial in the odometer reading at previous Smog Check, and ZIP code level demographic characteristics.

Table A.6: Robustness Check—Intensive Margin Interacting NOx Quartiles With Other Controls

	(1)	(2)	(3)	(4)	(5)	(6)
DPM * NO Q1	-2.676***	-2.807***	-2.294***	-2.412***	-2.421***	-5.089***
	(0.359)	(0.350)	(0.301)	(0.347)	(0.345)	(0.696)
DPM * NO Q2	-3.337***	-3.358***	-3.075***	-3.128***	-3.129***	-5.339***
	(0.359)	(0.357)	(0.334)	(0.355)	(0.354)	(0.631)
DPM * NO Q3	-3.925***	-3.941***	-3.858***	-3.881***	-3.875***	-5.728***
	(0.389)	(0.391)	(0.397)	(0.395)	(0.394)	(0.631)
DPM * NO Q4	-4.642***	-4.720***	-4.970***	-4.974***	-4.957***	-6.482***
	(0.425)	(0.433)	(0.444)	(0.440)	(0.442)	(0.653)
NO Q2	0.958	0.821	-5.997***	-7.404***	-7.433***	-4.917**
	(0.674)	(0.613)	(1.330)	(1.391)	(1.384)	(1.567)
NO Q3	0.242	-0.00702	-8.999***	-10.74***	-10.69***	-6.708***
	(0.889)	(0.798)	(1.563)	(1.619)	(1.605)	(1.527)
NO Q4	0.615	0.124	-12.63***	-14.43***	-14.34***	-9.222***
	(1.015)	(0.999)	(2.173)	(2.181)	(2.205)	(2.227)
Quartile-Time Trend Interactions	Yes	Yes	Yes	Yes	Yes	Yes
Vintage-Quartile Interactions	No	Yes	Yes	Yes	Yes	Yes
Quartile-Year Interactions	No	No	Yes	Yes	Yes	Yes
Quartile-Lagged Odometer Interactions	No	No	No	Yes	Yes	Yes
Quartile-Demographics Interactions	No	No	No	No	Yes	Yes
Calendar Month Fixed-Effects	No	No	No	No	No	Yes
N	2979289	2979289	2979289	2979289	2979289	2979289

Note: All regressions include vehicle fixed-effects, year fixed effects, vintage/truck fixed effects, a quadratic time trend, a sixth order polynomial in the odometer reading at previous Smog Check, and ZIP code level demographic characteristics.

Table A.7: Ratio of Remaining Deadweight Loss With Tax to Deadweight Loss with No Tax: Calibration

	σ^2	σ_B^2	ρ	$R(\tau_{naive})$	$R(\tau^*)$
1998	1.407	1.465	0.322	0.789	0.755
1999	1.408	1.471	0.299	0.785	0.755
2000	1.438	1.486	0.308	0.794	0.763
2001	1.457	1.496	0.311	0.799	0.767
2002	1.492	1.506	0.283	0.802	0.775
2003	1.517	1.535	0.283	0.807	0.781
2004	1.525	1.531	0.265	0.806	0.782
2005	1.474	1.539	0.265	0.796	0.771
2006	1.482	1.539	0.251	0.795	0.773
2007	1.487	1.547	0.247	0.796	0.774
2008	1.498	1.533	0.252	0.799	0.777
Average	1.471	1.513	0.281	0.797	0.770

Table A.8: Ratios of DWL with Tax to DWL With No Tax, Scrapping Most Polluting Vehicles

	Percentile Scrapped					
	None	1%	2%	5%	10%	25%
1998	0.434	0.338	0.316	0.293	0.286	0.323
1999	0.426	0.338	0.323	0.308	0.307	0.374
2000	0.433	0.350	0.336	0.323	0.323	0.405
2001	0.472	0.373	0.358	0.347	0.358	0.514
2002	0.490	0.407	0.396	0.388	0.398	0.546
2003	0.503	0.433	0.424	0.419	0.436	0.624
2004	0.544	0.464	0.456	0.455	0.485	0.686
2005	0.548	0.485	0.479	0.482	0.520	0.708
2006	0.595	0.511	0.506	0.518	0.577	0.757
2007	0.585	0.534	0.532	0.552	0.625	0.779
2008	0.605	0.556	0.558	0.590	0.681	0.806
Average	0.512	0.435	0.426	0.425	0.454	0.593

Notes: DWL with no tax calculated based on the difference in emissions from imposing a tax equal to the actual externality per gallon consumed by a particular car. SBO tax computed as the weighted average of externality per gallon, using the negative slope of the vehicle's demand curve as the weight. All taxes also weight vehicles by the frequency with which vehicles of the same vintage and class appear in the California fleet as a whole.

Table A.9: Cobenefits of a Gasoline Tax, Taking Heterogeneity into Account (No Extensive Margin Effect), Muller and Mendelsohn (2009) Baseline Values

	Δ Consumption	Δ CO2	DWL		Criteria Benefit						**Net Cost**
	(Gallons)	(Tons)	($)	(Per CO2)	(NOx $)	(HC $)	(CO $)	(Total $)	(% DWL)	(Per CO2)	(Per CO2)
1998	230.7	2.215	158.4	71.53	-0.473	93.88	118.3	210.6	134.4	96.13	-24.61
1999	230.0	2.208	153.4	69.49	-0.132	84.64	103.8	187.7	123.1	85.57	-16.09
2000	226.4	2.173	145.6	66.98	-0.0733	76.91	91.80	168.2	116.2	77.82	-10.83
2001	206.6	1.983	127.8	64.44	0.123	60.20	72.66	132.8	104.4	67.25	-2.807
2002	198.0	1.900	119.6	62.92	0.220	49.22	57.84	107.1	89.83	56.52	6.395
2003	191.8	1.841	113.2	61.50	0.359	41.86	47.85	89.91	79.68	49.00	12.50
2004	174.7	1.677	100.5	59.92	0.518	31.23	36.96	68.58	68.52	41.06	18.86
2005	134.9	1.295	74.84	57.79	0.447	21.20	24.51	46.05	61.74	35.68	22.11
2006	117.2	1.125	62.59	55.62	0.361	16.20	18.24	34.70	55.75	31.01	24.61
2007	104.7	1.006	54.15	53.85	0.311	12.79	14.09	27.12	50.38	27.13	26.72
2008	91.99	0.883	46.00	52.08	0.265	10.40	11.23	21.84	47.74	24.86	27.22
Average	173.3	1.664	105.1	61.47	0.175	45.32	54.30	99.50	84.70	53.82	7.645

Note: All units are in millions. Dollar figures are inflation adjusted to 2008 terms. NOx and HC are valued as in Muller and Mendelsohn's (2009) Baseline scenario by the county where each vehicle received its Smog Check. The values vary substantially across counties, but a population-weighted average for the state values HC at $3638.62 per ton per year, and NOx at $125.26 per ton per year. CO is valued at $802.88, as per the median value of Matthews and Lave (2000).

Table A.10: Cobenefits of a Gasoline Tax, No Heterogeneity, Muller and Mendelsohn (2009) USEPA Values

	Δ Consumption	Δ CO2	DWL		Criteria Benefit						**Net Cost**
	(Gallons)	(Tons)	($)	(Per CO2)	(NOx $)	(HC $)	(CO $)	(Total $)	(% DWL)	(Per CO2)	(Per CO2)
1998	426.6	4.095	292.9	71.53	23.94	687.5	122.2	832.3	287.3	205.5	-134.0
1999	434.6	4.172	289.9	69.49	25.38	629.1	107.3	761.2	264.4	183.7	-114.2
2000	406.5	3.902	261.4	66.98	20.80	507.5	81.80	609.7	234.6	157.1	-90.16
2001	346.1	3.323	214.1	64.44	15.52	335.9	51.89	403.2	189.3	122.0	-57.55
2002	343.3	3.295	207.3	62.92	13.31	238.9	34.09	286.2	138.6	87.19	-24.27
2003	337.4	3.239	199.2	61.50	11.69	173.6	22.24	207.5	104.6	64.32	-2.818
2004	336.8	3.234	193.8	59.92	9.785	103.0	11.06	123.8	64.22	38.48	21.44
2005	291.3	2.796	161.6	57.79	6.556	42.91	0.602	50.06	31.24	18.06	39.74
2006	276.6	2.655	147.7	55.62	3.928	-1.975	-7.260	-5.267	-3.585	-1.994	57.62
2007	268.8	2.581	139.0	53.85	2.190	-30.59	-12.31	-40.65	-29.43	-15.85	69.70
2008	284.1	2.728	142.1	52.08	0.394	-52.93	-16.11	-68.56	-48.51	-25.26	77.35
Average	341.1	3.275	204.5	61.47	12.14	239.4	35.95	287.2	112.1	75.75	-14.29

Note: All units are in millions. Dollar figures are inflation adjusted to 2008 terms. NOx and HC are valued as in Muller and Mendelsohn's (2009) USEPA scenario by the county where each vehicle received its Smog Check. The values vary substantially across counties, but a population-weighted average for the state values HC at $23897.10 per ton per year, and NOx at $1995.43 per ton per year. CO is valued at $802.88, as per the median value of Matthews and Lave (2000).

Table A.11: Cobenefits of a Gasoline Tax, Taking Heterogeneity into Account, Muller and Mendelsohn (2009) USEPA Values

	Δ Consumption	Δ CO2	DWL		Criteria Benefit						**Net Cost**
	(Gallons)	(Tons)	($)	(Per CO2)	(NOx $)	(HC $)	(CO $)	(Total $)	(% DWL)	(Per CO2)	(Per CO2)
1998	230.7	2.215	158.4	71.53	19.49	617.1	118.3	753.7	481.1	344.1	-272.6
1999	235.0	2.256	156.8	69.49	20.18	569.9	106.6	696.0	446.8	310.5	-241.0
2000	241.1	2.314	155.0	66.98	18.50	536.3	98.02	652.4	423.1	283.4	-216.4
2001	223.7	2.148	138.4	64.44	15.85	421.4	77.48	514.5	373.2	240.5	-176.1
2002	216.0	2.073	130.4	62.92	14.23	340.6	60.82	415.5	319.5	201.0	-138.1
2003	200.5	1.925	118.4	61.50	13.06	280.6	48.41	341.9	289.8	178.2	-116.7
2004	168.7	1.620	97.06	59.92	11.75	198.9	34.85	245.4	253.7	152.0	-92.11
2005	119.6	1.148	66.36	57.79	8.888	123.8	19.71	152.3	229.8	132.8	-75.02
2006	87.70	0.842	46.83	55.62	6.682	78.92	10.63	96.17	206.6	114.9	-59.32
2007	65.59	0.630	33.91	53.85	5.202	47.51	4.515	57.20	169.8	91.43	-37.58
2008	44.11	0.423	22.05	52.08	4.099	25.08	0.0752	29.25	133.5	69.52	-17.44
Average	166.6	1.599	102.1	61.47	12.54	294.6	52.67	359.5	302.4	192.6	-131.1

Note: All units are in millions. Dollar figures are inflation adjusted to 2008 terms. NOx and HC are valued as in Muller and Mendelsohn's (2009) USEPA scenario by the county where each vehicle received its Smog Check. The values vary substantially across counties, but a population-weighted average for the state values HC at $23897.10 per ton per year, and NOx at $1995.43 per ton per year. CO is valued at $802.88, as per the median value of Matthews and Lave (2000).

Table A.12: Percentage Difference Between California and the rest of the US

	25th Percentile	Median	75th Percentile	Mean
NOx g/mi	-0.230	-0.291	-0.338	-0.282
NOx Damage/ton (MM)	-0.439	-0.525	-0.558	-0.685
NOx Damage/mi	-0.595	-0.657	-0.712	-0.761
HC g/mi	-0.262	-0.321	-0.410	-0.354
HC Damage/ton	1.475	2.558	5.318	1.821
HC Damage/mi	0.602	1.134	3.358	1.035
CO g/mi	-0.226	-0.321	-0.366	-0.320
CO Damage/mi	-0.226	-0.321	-0.366	-0.320
NOx + HC Damage/ton (MM)	0.0191	0.994	2.337	0.787
NOx + HC + CO Damage/mi	-0.353	-0.299	-0.0883	-0.295

Notes: The table reports the coefficient on the California dummy divided by the constant. All differences are statistically significant at the 0.001 level, except for NOx g/mi and HC Damage/mi at the 25th percentile (significant at the 0.05 level), and NOx Damage/mi.