University of California Transportation Center
UCTC Dissertation No. 164

New Vehicle Choice, Fuel Economy and Vehicle Incentives:
An Analysis of Hybrid Tax Credits and the Gasoline Tax

Elliot William Martin
University of California, Berkeley
2009

New Vehicle Choice, Fuel Economy and Vehicle Incentives:
An Analysis of Hybrid Tax Credits and the Gasoline Tax

by

Elliot William Martin

B.A. (Johns Hopkins University) 2001
M.S (University of California, Berkeley) 2006
M.C.P (University of California, Berkeley) 2006

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering-Civil and Environmental Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Adib Kanafani, Chair
Professor Joan Walker
Professor Robert Cervero

Fall, 2009

The dissertation of Elliot William Martin is approved:

Chair _____   Date _____

_____   Date _____

_____   Date _____

University of California, Berkeley

Copyright

New Vehicle Choice, Fuel Economy and Vehicle Incentives:
An Analysis of Hybrid Tax Credits and the Gasoline Tax

© 2009

by Elliot William Martin

Table of Contents

Dedication

To Heather,

Who endured each trial with me,

With unconditional and forgiving support,

I am forever grateful for finding such a compatible partner in life.

Index of Figures

Index of Tables

Acknowledgements

There are many people that I want to thank for the completion of this work and my graduate career. I feel it is most appropriate to begin with my parents. Dwight and Cheryl Martin supported me regardless of my successes and failures. They provided invaluable guidance in making decisions, in reflecting on ideas, and in giving support whenever it was needed. They never were judgmental of the path that I wanted to take and were always there to provide support. They raised me well and provided me with a good home and childhood. I appreciate these things more everyday and I could not have asked for better parents.

I extend a sincere gratitude to my adviser, Adib Kanafani. It goes without saying that this work could not have been completed without his guidance. He provided invaluable feedback in developing the research and executing its completion. I remained continually impressed by the force and acuity of his intelligence, which was repeatedly evident in his ability to provide sound and efficient advice on a number of issues related to the research and my future work. He exhibited a patience with my thinking and with the trials that I faced unique to my graduate experience that I very much appreciate. His contributions to this work, my career and my life will not be forgotten.

I am grateful to the faculty who offered a good education and helpful mentorship along the way. I am grateful to Joan Walker and Robert Cervero for their advice and support in the completing this dissertation. I greatly admire their capacities as researchers and instructors. Another mentor of great importance in my graduate career was Susan Shaheen. Susan is an outstanding researcher with a great dedication to the

profession. She has played a critical and indelible role in my professional development, and I am thankful that I was given the opportunity to work with her.

There are several other people who played important roles in my work at Berkeley. Sadly, some of them I can no longer thank personally. I am still struck by the tragic loss of Alexander Farrell, another great leader in this profession. I admired his life path transitioning from the US Navy to excellence in academic research. He was extremely thorough in his thinking of the many problems facing transportation energy and I can still hear the tone and style of his arguments in my mind as I work through similar problems.

Finally, I have had the pleasure and honor of working with a great number of bright and intelligent peers during my time at Berkeley. Many of these peers have become good friends. They include Mikhail Chester, Chris Cherry, Ipsita Banerjee, Josh Pilachowski, Stella So, Megan Smirti, Juan Carlos Herrera, C. Edward Huang, Offer Grembeck, Anthony Patire, Jing Xiong, Jeff Lidicker, and many others. I am grateful that they were a part of my experiences here and it was an honor to be a part of theirs.

Abstract

New Vehicle Choice, Fuel Economy and Vehicle Incentives:
An Analysis of Hybrid Tax Credits and the Gasoline Tax

by

Elliot William Martin

Doctor of Philosophy in Civil and Environmental Engineering

University of California, Berkeley

Professor Adib Kanafani, Chair

Automobiles impose considerable public costs in the form of emissions and foreign oil dependence. Public policy has thus taken a considerable interest in influencing the technology and fuel economy associated with new vehicles brought to market. In spite of this interest, there is very limited information on the effectiveness of these policies in reducing greenhouse gas emissions or shifting vehicle demands. This is in part due to the fact that modeling the demand for automobiles is wrought with many challenges. These include large choice sets that change frequently over time and significant data collection obstacles. This work proposes a methodology for data development that simplifies many of the challenges associated with data collection in automotive modeling. The methodology explores a technique to merge data on aggregate sales with disaggregate vehicle holdings data to synthesize a complete dataset that preserves the strengths of both. The merged dataset is used to estimate a logit choice model of automotive choice

that is applied in evaluating the effectiveness of hybrid tax credits and the gasoline tax in reducing greenhouse gas emissions. Policy simulations suggest that hybrid tax credits have saved an average 1.5 million metric tons of greenhouse gas emissions based on sales between 2006 and 2007. When considered in conjunction with the cost of the policies, the credits appear to have a cost effectiveness ranging between $1000 to $3000 per metric ton of greenhouse gas emissions reduced. Hybrid tax credits are also found to be more effective than a doubling of the gasoline tax in shifting the new vehicle stock towards more fuel efficient vehicles. Finally, the model evaluates the market willingness to pay for fuel cost reduction. The results suggest an average willingness to pay of $522 in purchase price per 1¢ reduction in fuel cost per mile. This means that reasonable circumstances exist in which some buyers will pay more for fuel economy than they save in fuel cost expenses over the life span of their automobiles.

Professor Adib Kanafani (Chair)   Date

1. Introduction

Individual consumer decisions of automotive choice determine the aggregate composition of the vehicle fleet. Although consumer choice is generally a private matter, the collective vehicle choices of individuals have broad implications on overall fuel consumption. Aggregate fuel consumption, in turn, has critical impacts on air quality, energy consumption and global climate change.

Because the composition of the automotive fleet influences the external costs of transportation, it is natural that policy makers would be interested in influencing the private decisions that affect the country's vehicle fleet. The motivation for exerting this influence stems from the fact that external impacts and public costs are rarely priced by the private market and hence are rarely considered in private decisions. Public agencies often find that extending such influence is inherently difficult due to limited policy channels and imperfect information on consumer behavior. While automotive companies are most directly influenced by regulations and technological mandates, consumers primarily respond to expected personal costs.

Federal and state automotive policies have been established to influence consumer decisions in favor of alternative fuels and hybrid vehicles. A primary policy instrument applied in the recent past has been tax incentives on new vehicles. The existing political structure of fuel economy regulation make this an interesting policy to study.   Federal policy has absolute jurisdiction over air quality and fuel economy regulation of the auto industry. The notable exception to this federal mandate is the special status of California, which pre-empts federal regulation of air quality (CARB, 2009). This has allowed the

state to act independently with respect to air quality policy and has permitted other states to adopt either federal regulations or those of California. But other states cannot invent their own air quality policy. Direct incentives for purchasing a vehicle or buying a fuel do not fall within this dichotomous regulatory structure. Any state, independent of the federal government, can choose to place its own tax incentives on specific vehicles or fuels. Therefore, research that improves our understanding of how consumers respond to vehicle costs, fuel economy and incentives can inform many government agencies on policy design. In this way, vehicle tax credits can be implemented independent of federal consensus, and in many respects this is a considerable advantage.

Understanding the effectiveness of vehicle incentives is critical because there are many proposed policies available to state and federal governments aimed at reducing petroleum consumption and greenhouse gas emissions. However, there are limited resources with which to evaluate these policies. The State of California's AB 32 climate change legislation requires that greenhouse gases be lowered to 1990 levels by 2020 (California Assembly, 2006). Meeting these goals will require long-term planning and commitment on many fronts, but encouraging greater efficiency in the automotive fleet is one policy capable of contributing to reductions. Therefore, the development of a tool that helps forecast and assess the potential benefits of vehicle incentives is an important contribution to the policy arsenal of the federal and state governments.

But the effort to answer critical policy questions in automotive choice is inhibited by some considerable methodological and data challenges that are specific to the automotive market. Since consumers are the ultimate respondents to incentive-driven

policy, a study of the impact of incentives requires some model for automotive choice. But modeling automotive choice has faced several unique challenges. The choice set is very large and is constantly evolving over time. This market dynamic implies that automotive choice models estimated within a given time period can become obsolete rather quickly.

Thus, automotive choice models would appear to be well served by an empirical foundation that can be updated to incorporate present information quickly and efficiently. But this is not a common attribute of datasets that are typically applied in choice analysis. Often, the data generated for discrete choice models are disaggregate in nature, containing the decision and choice sets of individual respondents. Such datasets take considerable time, effort, and expense to generate. In the case of automotive policy, the choice of interest is the new vehicle choice. But disaggregate datasets that describe the choice and attributes of new vehicle buyers are well protected by industry, and it is unlikely that this will change. As a result, public policy making has been hampered by this lack of information.

This study addresses these challenges by developing a new approach for the design of datasets that can be used to model automotive choice. This approach uses automotive choice data to evaluate the impact and cost-effectiveness of the recent hybrid tax credits that have been in place since 2006. In addition, the model is also used for evaluating the degree to which changes in the gasoline tax can shift average fuel economy.

In conducting the policy analysis, this study, limited to the sedan market, addresses the following questions:

1) How cost effective are hybrid tax credits in reducing greenhouse gases in transportation?

2) What does current market behavior tell us about the extent of consumer willingness to pay for fuel economy?

This study demonstrates the synthesis of a new dataset from both aggregate and disaggregate data. In this context, aggregate data are defined as the total sales of vehicles broken down by month and by vehicle model, and disaggregate data are defined as observations of individuals that have made vehicle purchase decisions. Each data type has strengths lacked by the other. The aggregate data shows a complete picture of all consumer decisions over time. The disaggregate data, while constituting a snapshot of holdings that is not representative of market decisions, offers in-depth information on the decision-makers, which can help discern the factors that drive consumer decisions towards particular choices.

The contributions of this work are in the areas of data structure design and of vehicle policy. The choice model structures that are estimated and applied in this analysis are of standard formats developed for previous applications. This thesis demonstrates a unique data combination methodology using data as applied to new vehicle demand analysis. The development and exploration of this methodology may have applications in other fields within or closely related to transportation. As data becomes increasingly available through the internet and the automation of data collection

processes, opportunities for further development and application of this methodology are likely to become apparent. The availability of monthly vehicle sales data are a recent phenomenon, which permits an evaluation of new vehicle sales fluctuations with changes in gasoline prices, government policy and manufacturer incentives. In addition, the data explored in this study covers a unique period in the automotive industry in which fuel economy can be acquired for a premium. This has never been possible for consumers before. Thus, the maturity of hybrids during this period of study allows for a first look at the market willingness to pay for fuel economy from revealed preferences. Finally, the policy analysis evaluates the cost-effectiveness of the recent application of hybrid tax credits at the federal level to hybrid vehicles. This policy, begun in 2006, has a long enough history to offer empirical insights for its evaluation. The relevance of this evaluation will increase as new technologies and a broader array of hybrids come to market. Furthermore, the recent cash-for-clunkers policy illustrates the propensity of the federal government to engage in innovative policies similar to hybrid tax credits in order to impact overall fuel economy. While this study does not evaluate the cash-for-clunkers policy, insights gained here can offer a preliminary perspective on the likely ranges for cost-effectiveness of this new policy.

To proceed, Chapter 2 presents a review of the literature on choice studies that have been conducted with aggregate and disaggregate data with specific applications to automotive choice. Chapter 3 discusses the data sources and data collection procedures, the design of choice models, and the procedure of data synthesis. Chapter 4 presents the results of the model estimation within the sedan market and discusses their implications

for policy metrics such as consumer willingness to pay for fuel economy. Chapter 5 illustrates the application of the model to conduct policy analysis. This includes an analysis of the effectiveness of tax incentives to move the average fuel economy within the vehicle class, as well as a review of the cost-effectiveness of hybrid tax credits in lowering greenhouse gas emissions. Chapter 6 presents a summary and conclusions of the study. Appendix A provides a discussion of the American automotive market and offers insights on how automotive consumer choice sets are formed within the sedan market.

2. Literature Review of Automotive Choice Studies

Automotive demand and vehicle choice have been of interest to transportation researchers for several decades. Much of the research on this emerged during the 1970s and early 1980s with the onset of the energy crisis. The questions addressed then were similar to those of interest today, but they were asked within the context of a very different automotive market. The response of the automotive industry to the energy constraints of the 1970s involved a downsizing of vehicles as well as an emerging emphasis on diesel passenger cars. Models exploring automotive demand have been aggregate, predicting vehicle market shares with aggregate sales data, while others have been disaggregate, predicting the probability of individual choice estimated from datasets with disaggregate vehicle holdings and then reconstructing market choice.

2.1 Studies with Aggregate Data

Boyd and Mellman (1980) made an early and important contribution to aggregate automotive demand modeling by extending the theory of the multinomial logit model to produce what they call a hedonic demand model estimating the vehicle market shares of the entire 1977 vehicle market. Their nomenclature of "hedonic demand" was later qualified as being unrelated to the price regression models normally called hedonic demand models in the analysis of real estate economics (Train, 1986). The attributes incorporated into their aggregate model included the price of the vehicle, fuel economy, noise, acceleration, style, repair frequency and handling. Their model did not include any decision-maker parameters. The results of Boyd and Mellman (1980) generated a series of attribute valuation estimates. At that time, they found that consumers do make trade-

offs of lifetime fuel cost and purchase price in choosing among competing models. Using these results to generate policy simulations, they found that a doubling of gasoline prices (then $0.70/gallon) would raise average fuel economy of new vehicles by 6 percent, and that modest price changes favoring smaller vehicles could increase average fuel efficiency by about 3 percent. The work of Boyd and Mellman was ahead of its time far more than probably appreciated at its publication. It is the first known application of the mixed logit model, which would not start to receive significant attention until about twenty years later. The automotive market has changed significantly since then, with far more vehicle class options and greater diversity in fleet fuel economy.

Cardell and Dunbar (1980) also published a hedonic demand model (similarly misnamed), but for the purposes of comparing the consumer surplus impacts of CAFE standards versus fuel price changes. They also incorporate variation in consumer tastes by allowing coefficients to be random in their model, and they define weights by a predefined statistical density function with moments that are estimated via maximum likelihood. They concluded that policy which increased fuel prices would result in lower social costs than CAFE regulations for comparable oil import reductions.

One of most comprehensive models developed for aggregate automotive demand was produced by Berry, Levinsohn and Pakes (1995), commonly referred to as "BLP" in subsequent literature. Using only annual vehicle sales, they developed a general equilibrium model that simultaneously estimated parameters of consumer demand and producer prices. On the demand side, their model is very similar to that of Boyd and Mellman, in which the utility of a model depended on product attributes with random

8

coefficients. The incorporation of a production function for the purposes of estimating prices simultaneously with demand extended beyond the scope of many previous investigations. The aggregate sales data applied consisted of all vehicles marketed in the United States from 1971 to 1990, a period in which fuel economy increased significantly, and excluded "exotic" vehicles with extremely small market shares. The data set contained 2,217 total observations with 997 unique vehicle models. The time resolution was annual with each model/year as an observation, while the competing alternatives (i.e. the choice set) modeled for any given year was the entire automotive market, including sports cars, sedans, luxury sedans, minivans and pick-up trucks. A short list of parameters was considered, including Horsepower/Weight, Air Conditioning (a proxy for luxury at the time), miles/\$, size, and price. The results showed through demand elasticities that consumers of small fuel efficient vehicles are very responsive to changes in the fuel economy of competing models. Consumers of larger cars were actually found to have a disutility for fuel economy, with demand for large cars falling with an increase in fuel economy. The work Berry *et al.* was groundbreaking for its theoretical contribution and comprehensiveness in scope. But their model did not incorporate consumer attributes, and their results also illustrate the challenge of estimating theoretically consistent fuel economy parameters. Their final specification generated a negative mean coefficient for fuel efficiency, however the standard deviation coefficient of the mean fuel economy coefficient suggests a distribution of values that span both sides of zero. In theory, consumers of all vehicles should place some value on fuel economy in a vehicle when all other attributes are kept equal. Unlike attributes such as size and power, it is not possible for a vehicle to be too efficient. Because fuel efficiency

9

is an attribute and a function of operating cost, all consumers would consider themselves better off with a vehicle that is more efficient given no sacrifice of other attributes. Thus, a model that finds any disutility to fuel efficiency is one that does not capture the attributes that are offering the benefit for which efficiency is being sacrificed.

Berry, Levinsohn and Pakes produced a second study on automotive demand about ten years later that builds on their previous work (Berry et al., 2004). A key contribution of this study was their use of second choice data of consumer purchases of automobiles. The authors obtained a unique data set from General Motors called the CAMIP dataset, which was a large (N = 37500) proprietary survey of consumers making recent purchase decisions in 1993. Among other questions, the survey asked the purchaser which vehicle they would have bought had their first choice been unavailable. The objective of Berry et al. (2004) was to explore the impact that second choice data had on accurately modeling substitution patterns. The CAMIP data was a choice based sample, in which General Motors determined the number of households to sample from the registrations of each vehicle. The characteristics of these households were sampled alongside their second-choice. The authors link the CAMIP sample with the demographic profile of the U.S. Census Bureau's Current Population Survey. They use this model to predict both first and second choices of consumers (Berry et al., 2004). The results found that the logit model produced sensible signs for coefficients, and provided an adequate fit for the observed household and vehicle characteristics. But the logit was less effective in matching the characteristics of the first and second choice predictions. This led the authors to conclude that the logit is not as good at modeling the substitution

patterns within the market. They run predictions on several market scenarios, including the General Motors retirement of the Oldsmobile division and the introduction of SUVs by Mercedes and Toyota. They found that the introduction of second-choice data are far more helpful in describing these substitution patterns over a standard logit framework (Berry et al., 2004).

2.2 Studies with Disaggregate Data and Other Research

Disaggregate automotive choice models also emerged during the late 1970s and early 1980s. They were motivated by the myriad of obstacles facing automotive companies during this period including a rising challenge to American dominance of the industry. Early models encountered challenges similar to those found in aggregate models in obtaining theoretically consistent or significant signs in operating costs or fuel economy (Cowling and Cubbin, 1972, Hogarty, 1975; Manski and Sherman, 1980). Later models were more successful in developing estimates with the proper sign (Train and Lohrer, 1983; Berkovec, 1984; Berkovec and Rust, 1985; Mannering and Mahmassani, 1985). These and other disaggregate studies were constructed of survey data that provided socioeconomic information along with vehicle attributes for each vehicle-household pair. An important distinction of all these studies in comparison to aggregate models is that for practical reasons they were mostly based on observed vehicle holdings as opposed to actual new car purchases. Furthermore, the influence of fuel costs or fuel efficiency within these models is a minor point of these early studies and not the subject of policy. Many of the models were challenges in their own right to implement given limits in computing power and the burgeoning state of knowledge in choice theory

at the time. Early disaggregate models were either linear regression or standard logit. By the mid-eighties, nested multinomial logit applications to disaggregate automotive choice were just beginning to emerge as a means to handle the undesirable independence of irrelevant alternatives property (IIA).

Many of the disaggregate models generated throughout this period were also limited by the resolution of choice they sought to predict. For most, vehicle class, not the specific vehicle model, was the dependent variable (Mannering and Train, 1985). This generalization was motivated by data limitations, study scope, and computational challenges tied to estimating valid probabilities for every vehicle model given sample size limitations (Train and Lohrer, 1983). Under these circumstances, researchers would average vehicle attributes (such as horsepower, efficiency, space) within an entire market subclass (Mannering and Train, 1985; Beggs and Cardell, 1980). While the reasons for this treatment are understood given past data limitations, the averaging of physical attributes across vehicle models is problematic for many reasons. Among those reasons include problems that arise with models that have outlier attributes that can shift the average value of an attribute disproportionately in one direction. In addition, the averaging of attributes by class essentially washes out the information contained within class competition. The data reverts to distinguishing market preferences for a particular class. For example, in illustrating the value of fuel economy, most of the valuable information can be washed out in the average. Consumers are drawn to a particular class of vehicle because of a package of attributes, but distinctions in fuel economy within that class will increase their propensity to buy the more efficient vehicle all else equal.

However, if these attributes are averaged by class, then the distinctions in fuel economy, or other attributes, that make some models outperform others within a particular class are no longer apparent. If the class with the highest fuel economy is not popular overall, then it can appear that the market does not value fuel economy. Thus averaging of physical attributes across vehicle classes greatly reduces the potential for the choice model to describe how consumers react to those attributes that distinguish performance within the class.

Other research that has addressed American response to changes in gasoline price includes studies exploring the price and income elasticity of gasoline demand. The most recent study summarized the state of knowledge rather well, and estimated U.S. price and income elasticities for the period of 1975 to 1980 and 2001 to 2006 for comparative purposes (Hughes *et al.*, 2007). The authors used regression analysis on the log of gasoline consumption to estimate coefficients for price and income over the two periods to compute elasticity. Interestingly, the authors found that short-run gasoline price elasticities are near zero, far lower than they had been three decades earlier. In other words, gasoline prices would rise, but consumption (and driving) would not fall commensurately. They suggested that fewer transit options and greater suburbanization enforce today's driving distances regardless of fuel price. Indeed, American responsiveness to gasoline prices has historically been quite sluggish at the pump due to the near exclusive reliance on automotive transportation in many parts of the country. The conclusions of Hughes *et al.* are important for policy. They suggest that gasoline taxes reduce consumer surplus more than inducing behavioral change once a household's

vehicle stock has been established. However, this does not mean that Americans do not

react to contemporaneous and expected fuel prices when the opportunity arises for them

to readjust their vehicle holdings. Since the publishing of Hughes *et al*., remarkable

circumstances in the gasoline market have pushed prices to levels in which Americans

began to reduce their driving (Krauss, 2008).

Turrentine and Kurani (2007) took an entirely non-quantitative approach in

studying consumers' consideration of fuel efficiency in vehicle choice. Through the

course of fifty-seven household interviews, the researchers probed the subjects' thinking

behind all of their automotive purchases as a household. They asked directly about

respondent willingness to pay for a 1.5 times increase in fuel efficiency, and about

household methodologies in assessing the importance of fuel efficiency in their purchase.

This included probing desired pay back periods from gasoline savings. In their

discussion, Turrentine and Kurani encountered a range of responses from those who were

willing to give a confident answer, which was less than half of interviewed households.

Though they interviewed households with a wide range of intellectual skills, they found

no one making economically rationale calculations remotely close to those that vehicle

adoption models typically apply. Turrentine and Kurani ultimately suggest that

attempting to pin down a common mechanistic numerical treatment of fuel economy by

the consumer in the vehicle purchase process may be futile. In this, they do not mean to

suggest that car buyers do not consider fuel economy, but that cost functions are at best a

very general proxy for the decision process conducted by consumers. Although the style

of Turrentine and Kurani is unique to studies addressing these issues, the type of data

they produce is well-defined as stated preference. That is, the decisions with respect to willingness to pay are not tied to real decisions that must be made with actual resources. Stated preference data are still informative and important, especially when the object of measurement is difficult to measure, as in the case of willingness to pay for fuel economy. However, stated preference has natural limitations that must be acknowledged, and absent these limitations on measurement, revealed preference data are generally preferred.

The literature exploring automotive choice and demand reviewed here emphasizes the work that has been done with implications for valuing fuel economy. But fuel economy is just one perspective that researchers have explored in judging automotive demand. While fuel efficiency has been discussed in the past, it has rarely been the motivation for many of these studies. The value that consumers place on fuel efficiency has been challenging to ascertain in studies that have explored the issue from a variety of perspectives. Beyond the few aggregate and more numerous disaggregate choice studies, studies in hedonic demand have similarly encountered coefficients describing fuel efficiency to be insignificant or of counterintuitive sign (Espey and Nair, 2005, Arguea et al., 1994).

Surprisingly few studies of the automotive market have been done expressly to discern the consumer's valuation of fuel economy in car buying decisions (Greene et al., 2005). This study explores automotive demand from several new angles that have been rarely taken in the literature. To start, few studies of the automotive market have simultaneously combined the complete market share insights of aggregate data with the

consumer characteristics of disaggregate data to model new vehicle choice. Furthermore, this study is among the first to use hybrid sales to capture information on the consumer valuation of fuel economy. Finally, this study incorporates data on manufacturer incentives into price variables, which are excluded from most studies because of the difficultly in collection. Market data on hybrid purchases is now mature enough to permit better insights on the consumer value of fuel economy. Table 1 presents a summary of the literature reviewed in this chapter addressing automotive choice using choice models. The literature is vast and dispersed throughout many sources. This table includes only a sample of studies that use hedonic price models.

Table 1: Selected Literature of Automotive Choice and Demand

Authors	Title	Year	Methodology/Model	Data Type
Boyd and Mellman	The Effect of Fuel Economy Standards on the US Automotive Market: An Hedonic Demand Analysis'	1980	Mixed Logit	Aggregate
Cardell and Dunbar	Measuring the Societal Impacts of Automobile Downsizing	1980	Mixed Logit	Aggregate
Hogarty	Price-quality relations for automobiles: a new approach	1975	Regression	Aggregate
Manski and Sherman	An Empirical Analysis oif Household Choice Among Motor Vehicles	1980	Logit	Disaggregate
Train and Lohrer	Vehicle Ownership and Usage: An Integrated System of Disaggregate Demand Models	1983	Logit	Disaggregate
Berkovec	Forecasting Automobile Demand using Disaggregate Choice Models	1984	Logit with Equilibrium	Disaggregate
Berkovec and Rust	A Nested Logit Model of Automobile Holdings for one Vehicle Households	1985	Nested Logit	Disaggregate
Berry, Levinsohn and Pakes	Automobile Prices in Market Equilibrium	1995	"BLP" and Logit	Aggregate
Berry, Levinsohn and Pakes	Differentiated Product Demand Systems from a Combination of Micro and Macro Data: The New Car Market	2004	"BLP" and Logit	Disaggregate/Aggregate Population
Mannering and Mahmassani	Consumer Valuation of Foreign and Domestic Vehicle Attributes: Econometric Analysis and Implications for Auto Demand	1985	Logit	Disaggregate
Mannering and Train	Recent Directions in Automobile Demand Modeling	1985	Literature Review	Disaggregate
Hughes, Knittel and Sperling	Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand	2007	Regression	Aggregate
Turrentine and Kurani	Car buyers and fuel economy?	2007	Home Interviews	Disaggregate
Molley-Espey	Automobile Fuel Economy: What is it worth?	2005	Hedonic Regression	Aggregate
McManus	The Link Between Gasoline Prices and Vehicle Sales: Economic Theory Trumps Conventional Detroit Wisdom	2007	Hedonic Regression	Aggregate & Disaggregate (Separate)
Arguea, Hsiao and Taylor	Estimating Consumer Preferences Using Market Data - An Application to US Automobile Demand'	1994	Hedonic Regression	Aggregate
Cowling and Cubbin	Hedonic Price Indexes for United Kingdom Cars	1972	Hedonic Regression	Aggregate

Building on these studies, the intended audience of this work consists of modelers interested in combining aggregate and disaggregate data within related or other applications in discrete choice. In addition, the study aims to inform governments on the existing effectiveness of hybrid incentives as well as illustrating techniques that can be applied to forecast the impact of proposed incentives, fees and fuel taxes on vehicle choice now and into the future.

3. Methodology

This chapter is broken up into six sections. The first section reviews the

generalized extreme value (GEV) models that are applied for this work, describing their

foundation and mathematical construction. The second section introduces other GEV

models and discusses their relationship to this work as well as the broader field of

discrete choice analysis. The third section describes the scope and limitations of the

model and the data in the context of this study. The fourth section describes the data in

both its aggregate and disaggregate formats. The fifth section outlines the procedure that

merges the two data types together into a single data set. Finally, the sixth section

presents the attributes in the linear utility function that will be estimated.

3.1 Background on Generalized Extreme Value Models (GEV)

Generalized extreme value (GEV) models have been the structural backbone of

the choice modeling field since early applications almost 40 years ago. The core of all

choice models is the utility function (U_{in}). Utility is a real number that describes the

overall value of alternative i to individual n. In the context of this work, the decision-

maker is the car buyer and the alternative is a specific vehicle model (e.g, Honda Civic).

The higher the utility, the greater the appeal of the vehicle to the decision-maker. But it

is important to note that the value of utility for an alternative has no intrinsic meaning. It

only has relevance as it is compared to the utilities of competing alternatives. In other

words, only differences in utility matter. In theory, a car buyer will always choose the

vehicle that offers him or her the highest utility (U_{in}) among all competing alternatives.

But the value of (U_{in}) is unknown and always subject to some uncertainty. Hence the

utility function itself is divided into two parts, the systematic utility (V_{in}) and the

unobserved error (ε_{in}). Together they comprise the total utility ($U_{ni} = V_{ni} + \varepsilon_{ni}$). The

systematic utility is the portion of utility that is knowable and predictable. This is the

component of the model in which the coefficients are estimated on the attributes of the

vehicle. It can be specified in a variety of ways, but a linear specification is standard for

most models. That is, each vehicle has an established systematic utility, and one of those

vehicles in the choice set has the highest value of V_{ni}. But the error, ε_{ni}, can change the

rank order of the absolute utility U_{ni}. The unobserved error is assumed to follow a

distribution, and it is the selection of this distribution that ultimately defines the type of

choice model that is being applied. The assumption that underlies a GEV model is that

the unobserved errors of the utility function are extreme value distributed.

As described in Train (2003), the probability that a decision maker chooses

vehicle i is stated as follows:

$$P_{ni} = Prob(U_{ni} > U_{nj}, \forall j \neq i) \tag{1}$$

$$P_{ni} = Prob(V_{ni} + \varepsilon_{ni} > V_{nj} + \varepsilon_{nj}, \forall j \neq i) \tag{2}$$

$$P_{ni} = Prob(V_{ni} - V_{nj} > \varepsilon_{nj} - \varepsilon_{ni}, \forall j \neq i) \tag{3}$$

The term $\varepsilon_{nj} - \varepsilon_{ni}$ is the combination of two random variables and is itself a

random term. The error term for each individual is specified as a joint density vector

$\varepsilon_n = \langle \varepsilon_{n1}, \dots, \varepsilon_{nJ} \rangle$ (Ben-Akiva and Lerman, 1985). Across all alternatives faced by car

buyer n, this density can be described as $f(\varepsilon_n)$. The probability of alternative i is the

probability that the difference in systematic utilities across all alternatives is not

overridden by the difference in unobserved errors across all alternatives. The integral of

the probability term above is used to calculate the actual probability. Consider a

function $I(*)$ that acts as a boolean indicator function of the relationship represented by

the "*". This function is equal to 1 when the relationship is true, and 0 when the

relationship is false. The following integral describes the probability of car buyer n

choosing vehicle i.

$$\int_{\varepsilon} I\big(V_{ni} - V_{nj} > \varepsilon_{nj} - \varepsilon_{ni}, \forall j \neq i\big) f(\varepsilon_n) d\varepsilon_n \qquad (4)$$

When the difference between the systematic utilities is larger than the unobserved

error, the indicator function $I(*)$ returns a 1, which is multiplied by the density of the

error draw (Train, 2003). Across the distribution of all errors, the result is a value

between 0 and 1 that reflects the probability that the utility of vehicle i truly is larger than

the utility of all other alternatives j. This is the fundamental mechanism that drives all

choice models. The distinctions between choice models arise from the implications of

specifying different assumptions of $f(\varepsilon_n)$ (Train, 2003). When $f(\varepsilon_n)$ is specified iid

extreme value, the integral is closed form. Other distributional specifications do not

guarantee a closed form solution. If $f(\varepsilon_n)$ is normally distributed, then the integral

specifies the probit model, which requires numerical simulation to solve. The closed

form solution offered by the extreme value specification is what makes the logit formula

more popular than the probit. The functional form of the P_{ni} computation of the logit

model is stated in Equation (5):

$$P_{ni} = \frac{e^{V_{ni}}}{\sum_j e^{V_{nj}}} \tag{5}$$

The multinomial logit formula has several features that are useful in terms of predicting probabilities. To start, the sum $\sum_i \frac{e^{V_{ni}}}{\sum_j e^{V_{nj}}} = 1$, which is a basic property that ensures that the share acquired by any vehicle does not exceed 100% of all vehicles. At the same time, no single vehicle can have a probability equal to 1, unless that vehicle is the only choice. Otherwise, all probabilities must be less than 1, even if only by the smallest of margins. Secondly, the functional form of Equation (5) is a sigmoid and follows the intuitive logistic shape. The logistic curve, shown in Figure 1, has some convenient interpretative properties with respect to predicting probabilities.

Figure 1: Hypothetical Logistic Curve
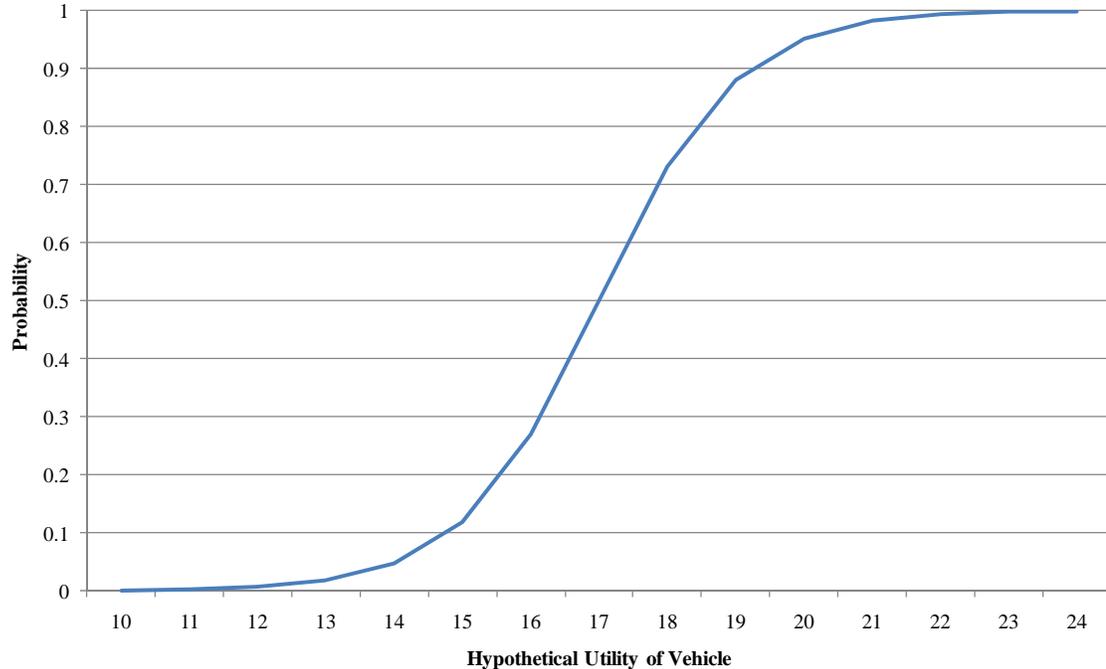
Figure 1 is a hypothetical diagram in which a vehicle with a utility of 17 has a 50% chance of selection. At this utility, a change of 1 utility in either direction makes a relatively large difference in the probability of selection. Whereas if the vehicle were relatively undesirable with a utility of 11, the same change in utility (by perhaps the same exact improvement) makes little difference in the probability of selection. The same phenomenon occurs when the utility of a vehicle is exceptionally high in comparison to its competition. A change of 1 utility does not significantly affect the degree to which it is likely to be selected.

The multinomial logit model has achieved its widespread application because this intuitive interpretation is unified with a closed-form mathematical foundation. In the next section, the application of the logit and more advanced GEV models will be discussed in the context of vehicle choice. In addition, this section will review challenges that were encountered in working with the more advanced GEV models in the context of vehicle choice.

3.2 The Application of Choice Models for Automotive Demand

The conventional multinomial logit model (MNL), with its closed form and intuitive interpretations has been the predominant approach to aggregate and disaggregate choice analysis (Berry et. al 1995; Train and Mannering, 1986; Choo and Mokhtarian, 2002). Multinomial logit is versatile, intuitive, and simple to estimate even for large datasets. For this reason, it is still used today for many applications. However, the MNL model is known for its inability to fully represent substitution patterns. The formulation of MNL results in the undesirable property of independence from irrelevant alternatives

(IIA). This property arises from the fact that the logit model assumes that the random components of utility are uncorrelated across alternatives. The implication is that unobserved attributes that impact the utility of one alternative will not impact the utility of any other alternative in a similar way. IIA also implies that the removal of one option from the choice set will increase the share or probability of all remaining alternatives by the same percent. In spite of its application to automotive choice, it is understood that this dynamic is not consistent with the true behavior of consumers in the automotive market.

Because of this shortcoming in substitution pattern modeling, researchers developed improved specifications that capture the substitution patterns with higher accuracy.   One such specification is the nested logit model, which has often been employed when the need to avoid IIA is significant and the structure and sequence of the choice set is well-defined. The functional form of the nested logit model as given in Train (2003) is shown as Equation (6):

$$P_{ni} = \frac{e^{V_{ni}/\lambda_k}\left(\sum_{j\in B_k} e^{V_{nj}/\lambda_k}\right)^{\lambda_k - 1}}{\sum_{l=1}^{K}\left(\sum_{j\in B_l} e^{V_{nj}/\lambda_l}\right)^{\lambda_l}} \tag{6}$$

The nested logit model divides the entire choice set into K nests that contain mutually exclusive subsets of all choices. Each nest is characterized by its own logsum coefficient λ_k which describes the degree to which the unobserved errors are correlated within the nest. The λ_k is estimated along with the coefficients of the utility function, but to be consistent with random utility theory, λ_k must be constrained to be between 0 and 1. The closer λ_k is to zero, the more highly correlated the unobserved errors of the

23

utilities within the nest. If all λ_k values across all nests are equal to 1, the nested logit

collapses to the logit. The IIA property holds within each nest, but does not hold for

alternatives in different nests. The extraction of one alternative from any nest raises the

probability of other alternatives in the same nest by a different proportion than

alternatives outside the nest. This is an improvement over the standard logit as attributes

such as vehicle size could be used as a good proxy for classifying nests. For example, the

unobserved correlation of errors pertaining to the utility of the Toyota Yaris and the

Nissan Versa, both small vehicles, is high. The nested logit would permit the extraction

of the Versa, and increase the probability of the Yaris proportionally more than larger

automobiles such as the Nissan Maxima.

While the nested logit is an improvement in modeling substitution patterns, its

rigid partition of choice sets introduces some undesirable properties as well (Berkovec

and Rust, 1985). For one, the choice of nesting structure is inherently subjective, as

researchers must try different specifications to discern the best fit. The degree of

correlation or independence of alternatives within the nest is constant across all the

within-nest alternatives. Different nests can have different degrees of within-nest

correlation, but all alternatives are governed by the single logsum parameter that

describes the degree of independence within the nest.   Furthermore, a change in utility of

an alternative in one nest, impacts all alternatives in all other nests by the same

proportion. The ratio of probabilities of two alternatives in different nests is independent

of choices in all other nests. In this way, the nested logit exhibits a property of

"independence from irrelevant nests" (Train, 2003).

Recent research in choice modeling has yielded structures that have improved the flexibility of alternatives within the choice structure while still avoiding IIA. Early structures included cross-nested logit (CNL), paired combinatorial logit (PCL), and the ordered generalized extreme value (OGEV) model (Vovsha, 1997; Small, 1987; Small 1994). The approach of these models was to permit overlapping nests, in which an alternative could be a member of more than one nest. Generally, these models allocate alternatives proportionally to specific nests. The estimation procedure estimates the coefficients of the utility function as well as the coefficients that indicate the proportion of membership for each alternative. Wen and Koppelman (2001) introduced the Generalized Nested Logit (GNL) model, which encompassed many of the above models as special cases. The functional form of the GNL is given as Equation (7).

$$
P_{ni} = \frac{\sum_k (\alpha_{ik} e^{V_{ni}})^{1/\lambda_k} \left(\sum_{j \in B_k} (\alpha_{jk} e^{V_{nj}})^{1/\lambda_k} \right)^{\lambda_k - 1}}{\sum_{l=1}^{K} \left(\sum_{j \in B_l} (\alpha_{jl} e^{V_{nj}})^{1/\lambda_l} \right)^{\lambda_l}}
\tag{7}
$$

The generalized nested logit model permits freedom of proportional allocation to established nests as well as flexibility in estimating different nesting coefficients for each nest. This proportional allocation is achieved with the coefficients α_{ik}, which correspond to each alternative i and nest k. Like the logsum coefficients, the α_{ik} terms are constrained to be between 0 and 1, and $\sum_k \alpha_{ik} = 1$. That is, the proportional allocation of any one alternative to all nests cannot exceed 100 percent.

The properties of the GNL are especially desirable in markets in which the divisions between products are not categorical but ordinal. Vehicles in the automotive

market are often classified by how they serve specific lifestyle needs. But the ability of vehicles within one class to satisfy lifestyle needs of another class is dependent on the general similarity of the vehicles within the two respective classes. The nested logit treats all vehicles outside the nest in the same manner. In this respect, the generalized nested logit model is a recent advance in choice modeling that has some very desirable attributes and has not yet been applied in analyzing the automotive market.

While the more advanced choice models are desirable from the standpoint of theoretical fit with product substitution patterns, the estimation of these models comes at a considerably higher cost. Because the proportional allocation of each alternative to each potential nest must be estimated, the number of additional parameters rises significantly with the structure of the CNL and even higher with the GNL. For a model with a small number of alternatives, the additional estimation of allocation parameters is not problematic, akin to adding a few more terms to the linear model. But the impact of the allocation parameters grows quickly as the number of alternatives increases. With just four nests and fifty alternatives, slightly smaller than the current sedan market, the allocation parameters number 200 in addition to the parameters of the utility function. These additional parameters are time consuming to estimate, but also can increase the propensity of encountering local optima or failing to achieve convergence. Therefore, in spite of improvements that aid the modeling of substitution patterns, these models are not always successful in generating parameter estimates that effectively characterize the behavior of decision-makers.

In this research, estimations of the NL, CNL, and GNL were attempted with

multiple specifications. However, all of these efforts failed to produce viable models. A

major obstacle was the inability of the final estimates to be consistent with random utility

theory. That is, many of the logsum terms if left unconstrained, would converge to

numbers outside the range (0, 1). The coefficient estimates were also unstable, and

dependent not only on starting values but also on the balance of allocation parameters. In

many instances of the higher order models, no convergence was achieved at all. Thus,

while these recently developed higher order GEV models show great promise for

capturing sophisticated substitution patterns, the research here could not find a reasonable

set of estimates within any structure except the logit, which exhibited considerable

stability in coefficient sign and magnitude across many specifications. Hence, the policy

analysis of this research is based on the logit.

3.3 Scope of the Model and Limitations

The scope of the analysis pertains strictly to new vehicle choice.   There is no

attempt to model whether or not a person tries to buy a new vehicle at all. This is

sometimes referred to as the "outside choice", in which the decision maker has the choice

to make none of the available choices. To effectively model this component of the

decision, the model would have to expand the sample to include all possible new car

buyers. This would involve additional model dynamics that would require inputs that

determine how large the population of new car buyers is at each time period, and how

many of them decide to purchase a new car. Such models are sometimes referred to as

vehicle transactions simulator (VTS) models, which attempt to simultaneously determine

the size of the market and the share of buyers making each choice. Such simulations are often extremely complicated, because they can involve a simulation of the vehicle stock and require assumptions about the pace of vehicle retirement over time. Such models can be useful when factors that bring specific cohorts of new buyers to the market are changing drastically over time. But most choice models take the market size to be exogenously determined and assume that the buyers entering the market over time do not vary in type significantly. For instance, an exogenous forecast of vehicle sales can be produced by a regression analysis on vehicle sales history, and then the choice model simulates how those forecasted sales are broken up based on forecasted inputs such as gasoline prices and policy variables. Speculating on how the results of standard choice models differ from the broader vehicle transactions models can be difficult as VTS models are not necessarily more accurate as a result of the increased complexity. The vehicle forecast is embedded in the choice model and this places more demands on the choice model structure. The VTS model could overestimate or underestimate the vehicle sales in comparison to a more simple regression approach. But the VTS model could have advantages in handling more abrupt and idiosyncratic changes in vehicle demand than simple regression. Regression analysis relies on stable patterns than are known in advance, such as seasonal adjustments and changes in income. But policies or changes in economic circumstances that have no precedent are difficult to model with regression. If such events occur often, then the VTS model may have more advantages in the long run. One example in which the VTS might be advantageous over an exogenous approach is the recent cash-for-clunkers policy, in which the incentives of a specific cohort to buy automobiles are changed, and this cohort is suddenly brought into the market. The VTS

model could be designed to account for this effect, whereas a regression model would likely miss the forecast. However, this scenario is only credible if the VTS model is effective in capturing these complex and random events. Otherwise, the simplicity of an exogenous forecast can be more appealing, as it can provide reasonable accuracy at a low cost and is guaranteed not to interfere with the more complex development of the choice model.

The model does not account for distinctions in state and local incentives. Regional policies such as state tax credits and HOV lane passes do offer new car buyers incentives for purchasing vehicles within specific regions. A minority of states offer hybrid tax credits, and those are given on top of the federal tax credits. Ideally, these policies should be incorporated into the model. Unfortunately, the resolution of the sales data as reported by manufacturers is strictly national. State based sales information can be obtained through new vehicle registrations by state. This data are collected by a company, R.L Polk & Co. Polk can even break such data down by zip code. Arrangements with Polk for data acquisition were explored, but obtaining just state-based data at a monthly resolution was prohibitively expensive, on the order of tens of thousands of dollars at academic prices. Such data resolution is possible, and it would improve the ability of the model to capture effects such as HOV stickers and state incentives. But obtaining this data with monthly resolution by state or zip code is an extensive undertaking. Such data would naturally be better, because it could account for additional incentives unseen just at the federal level. But this data resolution is also very

expensive, and it is not clear at this juncture whether the additional modeling benefits of this resolution merit the additional cost.

3.4 Data Sources and Data Collection

In spite of the important position that automobiles have in the nation's transportation system, there is a surprising dearth of organized public information regarding vehicle sales, attributes, and prices. Perhaps less surprising is the fact that much of the information available on automotive markets is geared towards the marketing and selling of cars. For this reason, data collection for this study had to build on a variety of separate sources that were ultimately combined together.

Today, the network of internet sites posting specialized information present vast opportunities to combine disparate sources of data together. To collect data for this project, programs were written in Visual Basic Applications (VBA) to traverse pages, pick up desired data points, and assemble a new macroscopic database where none existed before. Once collected, the data was organized and combined into a master dataset capable of conducting policy analysis. Because the datasets are large, the task of assembly and combination required the composition of additional VBA programs that work locally with data on the hard disk to align the attributes with the proper observations. The VBA programs were written for each organizational step working with both the aggregate, the disaggregate data, as well as the merging of the two datasets.

3.4.1 Aggregate Data Sources

The sources for the aggregate data include *Automotive News*, which publishes information regarding monthly vehicle sales to the resolution of the vehicle model. It also produces weekly summaries of the manufacturer and dealer sponsored customer incentives for purchasing a vehicle. These data describe the amount of "cash-back" given to customers when they purchase a vehicle model at a particular time. These data are an important component of price that is often omitted from aggregate automotive studies because they have not been readily available until recently. The collection of manufacturer summaries were imported from the PDF formats into Excel and a VBA procedure was written to extract, assemble, and align the incentive data with the appropriate vehicle.

The Environmental Protection Agency (EPA) also has a publicly available dataset containing information pertaining to the fuel economy of all vehicles. Unfortunately, it currently lacks in-depth information on the technical specifications of the models. To assemble a more comprehensive dataset containing vehicle specifications for the automotive fleet, information is collected from *Yahoo! Autos*, which makes the technical specifications of vehicles publicly available. This information is useful in defining the performance of the vehicle along metrics such as vehicle size, power-to-weight ratio, trunk space, fuel economy, and tire aspect ratio. The data collection also extended beyond the technical specifications of the vehicle to capture quality attributes that are defined by direct testing and review of the vehicle. Quality metrics are important because the technical specifications of a vehicle only describe what the manufacturer

states of a vehicle performance. This information cannot discern the quality of the

vehicle, whether the vehicle breaks down frequently, and cannot elaborate on more

qualitative measures that define the human interaction with the vehicle. For qualitative

data, *Consumer Reports* was used. *Consumer Reports* publishes a well-known and

widely consulted rating system of automobiles. The *Consumer Reports* rating systems

evaluates the quality and reliability of automotive systems. Furthermore, this rating

system is maintained historically, allowing the appropriate alignment of ratings for a

particular year to the vehicle sales of that year.

These data sources constitute the prime components of the aggregate dataset.

While each data source reported information differently, each presented enough

information to permit an accurate linking of the respective records together. The

observations across all datasets are linked by the vehicle model name as the key

identifier. This identifier was the most refined resolution available within the *Automotive*

News sales dataset. It is important to note that within the sphere of automotive

classification, there does exist an additional division of vehicle classification. This level

of classification within the vehicle model is called the vehicle "trim." Trims are usually

distinguished by a different set of standard interior options. These options can consist of

leather seats, a moon roof, among many other options that automakers put in the car to

distinguish their product from the competition. Most of these options do not affect the

drivability of the vehicle. The basic parameters of performance, including size,

horsepower, and fuel economy are the same across many trims. But some of the higher

end trims do come with a change in the engine size. For sedans, there are essentially two

engine sizes, a 4-cylinder and a 6-cylinder engine. The larger engine has higher

horsepower but lower fuel economy. This change in fuel economy is not substantial.

But ideally, the automotive sales data would distinguish sales by trim or engine size as

there is no level of classification that is lower than trim other than discrete combinations

of vehicle options. Unfortunately automakers do not report sales at the trim resolution.

The only way to construct such data would be to obtain the registration records of state

DMVs, which do contain information on engine size. Even the disaggregate survey data

collected for this study does not contain trim as not everyone knows the specific trim or

engine size of his/her vehicle. Thus, almost all studies that address automotive demand,

including those with access to proprietary disaggregate datasets use data with a resolution

no better than the vehicle model.

The scope of the aggregate data in this study is sedans sold from 2005 to 2007.

These are brands priced at a Manufacturer Suggested Retail Price (MSRP) of $30,000 or

less. The span of 2005 to 2007 offers several important attributes that are useful for

evaluating the consumer valuation of fuel economy and response to hybrid policy. This

time span covers a period in which hybrids were gaining respectable market share and

also the period in which the federal tax credits for hybrids were introduced. In addition,

the years of 2005 and 2007 saw a considerable rise and fluctuation in fuel prices

throughout each year. These three factors make this period exceptionally promising for

evaluating consumers' valuation of fuel economy through actual purchase decisions.

This study focuses only on the sedan submarket and there are several reasons for

this restriction. A primary reason is data availability with respect to key attributes of

other vehicle classes. The data sources covering other vehicle classes were more limited in their coverage of technical specifications. For example, the passenger volume and cargo volumes of SUVs and trucks, which is a key comparative feature of these vehicles, was not available. In addition, because the vehicle classes cater to different types of consumers, it was assumed that most choice sets constitute a collection of vehicles within same vehicle class. For instance, consumers who consider purchasing the compact sedans such as the Honda Civic are unlikely to simultaneously consider an SUV as a closely competing choice. Furthermore, most of the vehicle policy surrounding hybrid tax credits is targeted at the sedan market, where consumers respond more readily to changing costs. This is in contrast to the luxury sedan market, in which purchasing decisions are made on niche market attributes that more often are unobservable. Thus, these limitations led to a restriction of the model to the sedan market, which constitutes 56 vehicle models. Ideally, a model could handle the hundreds of choices available to consumers under one framework, but the construction of such a model must overcome considerable data and methodological challenges.

3.4.2   Disaggregate Data Sources

There are two primary sources for the disaggregate data employed in this study. These datasets are selected because of their availability and because they contain information that links the respondent's demographics to the vehicle that they held. One dataset is a sample of the population from the State of California. This dataset (hereafter the CEC dataset) is collected by a consulting firm under contract with the California Energy Commission (CEC). The CEC dataset is collected for the purposes of updating

the CALCARS model, which is the demand forecasting model that the CEC applies to forecast the impact of policies on vehicle demand. Specifically, the CALCARS model has been used to discern the degree of penetration of hybrids and alternative fuel vehicles within the state, given certain policies. The CALCARS model is a discrete choice model based on the CEC's Personal Vehicle Model, which was developed in 1983 (Page et al., 2007). But the CALCARS model was updated to use the WAVE datasets, which was collected in California during the early 1990s (Kavalec, 1996). This model builds on the work of Kenneth Train, David Bunch, David Brownstone, three California economists who utilized the WAVE survey datasets to merge stated preference (SP) and revealed preference data (RP) (Brownstone et al., 2000). The stated preference survey asked questions about preferences for hypothetical vehicles as well as preferences for common vehicle attributes. The CEC dataset was collected as an update to WAVE. It has both SP and RP components. The responses of interest in the CEC dataset are the RP responses because they revealed which vehicles the respondent actually owned. In addition, the CEC dataset includes a variety of important demographics, including income, family size, age of adults, education, among a few other household attributes. The total sample size of the CEC Dataset is 4,110 household observations. It was collected in 2007 and was generously provided by staff at the CEC.

The other disaggregate dataset is based in Chicago, Illinois. The Chicago Metropolitan Agency for Planning (CMAP) is a new planning agency formed in 2005 to unify the planning efforts of the ten surrounding counties including the City of Chicago. As part of this new effort, the agency has conducted a new survey of travelers within its

metropolitan region, known as the "Travel Tracker Survey." This survey was most recently completed for the year 2008, and the resulting dataset (CMAP dataset) was obtained with permission from CMAP. The CMAP dataset is larger than the CEC dataset with approximately 14,000 households. It was designed to be a metropolitan travel survey; hence, it is entirely a revealed preference dataset. The value of the CMAP is that similar to the CEC dataset; it collected the make, model and year of all household vehicles, allowing for the mapping of demographics to the specific vehicle holdings.

To build the disaggregate database, the CMAP and the CEC datasets were combined. All vehicles held by the households were listed as separate entries. That is, a household that held two vehicles was listed twice, as both vehicles are representative of the types of vehicles that the household chose to hold. While the aggregate data pertains specifically to new vehicle transactions, vehicle holding datasets that are collected by household surveys are not witness to the transaction through which the vehicle was acquired. Neither dataset states whether the household holding a particular vehicle chose to purchase it new or used. Fortunately, the time frame of the study covers recent vehicle transactions, and this permits a reasonable assumption with respect to vehicles held that are relatively new. Because both datasets are recent and because the sales data of interest is recent, any sedan within the CMAP or CEC dataset that had a model year of 2005, 2006 or 2007 was considered to have been acquired as a new vehicle. That is, although the disaggregate dataset contains many more vehicles of previous model years held by households, the only holdings that were considered were sedans of the model years that coincided with the sales data. In total, this constituted 2,114 vehicles.

3.5 The Combination of Aggregate and Disaggregate Data

3.5.1   Background

Aggregate data illustrates the complete market preference efficiently and continuously over time. In particular, aggregate data are well suited for capturing the consumer reaction to gasoline prices. In comparison, disaggregate datasets cannot easily capture such information from revealed preferences, and as such, they often rely on stated preference surveys to reveal consumer reactions to changes in operating costs (Brownstone et al, 2000). However, aggregate data has the significant disadvantage of lacking the individual characteristics of decision-makers.

Because each data type offers information lacked by the other, the two datasets are combined in a way that emphasizes the relative strengths of both. Merging the aggregate dataset's complete and continuous market perspective over time with the more detailed information of vehicle holdings contained within disaggregate datasets permits a richer understanding of which incentives are likely to impact specific consumer populations. Beyond what can be discerned solely through disaggregate vehicle holdings data, a combined dataset can yield a more manageable method for governments to apply in evaluating changes of consumer vehicle preferences over time.

For choice analysis, datasets can generally be classified as either exogenous or endogenous. The exogenous dataset is generated by a collection procedure that does not take choice as given, and the sequence of observations is random and independent of the choice of the respondent. This sampling procedure determines the market share of a particular choice. When the sample size is large and the sampling procedure of a survey

is simple and random, the market share of a choice within the disaggregate data is an unbiased estimator of the aggregate market share. On the other hand, an endogenous dataset is choice based and seeks out survey respondents that are making a particular choice to generate the dataset. As an example, an endogenous sample would be generated by a survey that is conducted within an auto dealership of a single manufacturer. If a researcher goes to a Toyota dealership, and only surveys car buyers within that dealership, the sample will be random within Toyota buyers, but not representative of the market. Hence, the endogenous sample collects rich data on the targeted choice, but does not produce an unbiased estimate of the aggregate share of decisions. The larger the sample of an exogenous dataset, the less necessary are the sacrifices imposed by the properties of the endogenous sample.

The exogenous sample is also more flexible as endogeniety is an undesirable restriction on data. From the perspective of an analyst, there is no advantage of an endogenous dataset over one that is exogenous. Both datasets are random and representative within the choice, but the exogenous dataset is also random and representative across choices. The advantage of an endogenous dataset is the reduction in collection costs (Manski and Lerman, 1977). These costs are considerably lower when decision-makers of a particular choice tend to cluster by location. Because an exogenous dataset is also random within a choice, it characterizes the within-choice distribution of respondents just as well as the endogenous sample when the sample size is large.

Given any choice within the disaggregate data, the joint distribution of attributes describe the type of decision-makers that are drawn to that choice. For example, the

owners of a given vehicle are characterized by a true distribution of income, education, family size, and age. A large sample that is simple and random within this choice should accurately characterize the distribution of any decision-maker characteristic. The observations of each choice should be randomly drawn from the population. The random draw can be conducted with an endogenous or exogenous sampling strategy. In either case, the distribution of decision maker attributes will be representative of the true distribution within each choice. What will be different based on the sampling strategy is the degree to which the share of each choice is representative of the market share in the population. If the sampling strategy is endogenous but random within choice, then the within-choice distribution of decision-maker attributes will be unbiased, but the market share of the decision will not. Alternatively, if the sampling strategy is exogenous, then both the within-choice distribution and the market share will be unbiased estimates of the population so long as the sample size is large.

3.5.2 Methodology of Data Combination

A combined dataset is generated from the aggregate and disaggregate datasets. To generate the combined dataset, the households holding vehicles within the disaggregate dataset are organized by their respective vehicles. Each vehicle i is distinguished by a unique make, model, and year; and there is a collection of households associated with each vehicle i. This collection characterizes the distribution of household attributes that hold vehicle i. The distribution can be specified individually (e.g., along the lines of income, education, etc.), or taken together they characterize a joint distribution of attributes pertaining to households that hold the vehicle. This distribution

is merged with the market share that the vehicle represents in the aggregate data. When the data are merged, each individual household essentially represents a small proportion of car sales within the aggregate data.

For example, consider a single household n that represents a proportion p_{ni} of all households holding vehicle i within the disaggregate data. Let vehicle i be considered a unique make, model and year, and let the market share in the aggregate data of vehicle i in month m be k_{im}. The combination of aggregate and disaggregate data merges these proportions such that household n represents a share of $k_{im}p_{ni}$ of the merged dataset. That is, the merged data considers a share of $k_{im}p_{ni}$ decision makers that look like household n to have made the decision to purchase vehicle i during each month m. One assumption that must be made given the limitations of the data is that the demographic distribution of households purchasing vehicle i does not change across the months within any given year. An inherent limitation of the available vehicle holding datasets is the inability to determine the month of purchase. Based on this data limitation, this study assumes that it is unlikely that the demographic profile of the owners of any specific vehicle shifts significantly over the course of a year. Figure 2 presents a diagram that shows how the data are merged for a subsample of the data. The disaggregate categorical values are shown simply for illustration of the data flow into the merged dataset.

Figure 2: A Diagram of Sample Data Merged

Aggregate Data

Vehicle	Period	MSRP-Incentive Adjusted	Power to Weight	Cost per mile	Consumer Reports - Climate Rating	Sales
Ford Fusion	Oct-05	18,898	0.0511	9.92	4	4,078
Toyota Camry	Oct-05	21,717	0.0510	9.39	5	32,323
Toyota Prius	Oct-05	21,275	0.0381	4.94	5	9,939
Honda Accord	Oct-05	22,670	0.0535	8.94	5	26,977
Nissan Altima	Oct-05	22,175	0.0583	9.76	4	16,624

Merged Data

Respondent ID	Vehicle	Date	Income	Age	Education	MSRP-Incentive Adjusted	Power to Weight	Cost per mile	Consumer Reports - Climate Rating	Sales
1	Ford Fusion	Oct-05	4	45	2	18,898	0.0511	9.92	4	816
2	Ford Fusion	Oct-05	4	46	2	18,898	0.0511	9.92	4	816
3	Ford Fusion	Oct-05	4	49	6	18,898	0.0511	9.92	4	816
4	Ford Fusion	Oct-05	6	44	2	18,898	0.0511	9.92	4	816
5	Ford Fusion	Oct-05	5	41	6	18,898	0.0511	9.92	4	816
6	Toyota Camry	Oct-05	7	30	4	21,717	0.0510	9.39	5	6,465
7	Toyota Camry	Oct-05	6	33	6	21,717	0.0510	9.39	5	6,465
8	Toyota Camry	Oct-05	6	29	2	21,717	0.0510	9.39	5	6,465
9	Toyota Camry	Oct-05	6	43	2	21,717	0.0510	9.39	5	6,465
10	Toyota Camry	Oct-05	5	24	5	21,717	0.0510	9.39	5	6,465
11	Toyota Prius	Oct-05	6	47	6	21,275	0.0381	4.94	5	1,988
12	Toyota Prius	Oct-05	7	42	5	21,275	0.0381	4.94	5	1,988
13	Toyota Prius	Oct-05	6	27	6	21,275	0.0381	4.94	5	1,988
14	Toyota Prius	Oct-05	5	43	4	21,275	0.0381	4.94	5	1,988
15	Toyota Prius	Oct-05	7	33	6	21,275	0.0381	4.94	5	1,988
16	Honda Accord	Oct-05	7	43	2	22,670	0.0535	8.94	5	8,992
17	Honda Accord	Oct-05	4	48	6	22,670	0.0535	8.94	5	8,992
18	Honda Accord	Oct-05	7	50	5	22,670	0.0535	8.94	5	8,992
19	Nissan Altima	Oct-05	4	31	3	22,175	0.0583	9.76	4	3,325
20	Nissan Altima	Oct-05	5	42	2	22,175	0.0583	9.76	4	3,325
21	Nissan Altima	Oct-05	5	26	6	22,175	0.0583	9.76	4	3,325
22	Nissan Altima	Oct-05	4	50	4	22,175	0.0583	9.76	4	3,325
23	Nissan Altima	Oct-05	6	47	5	22,175	0.0583	9.76	4	3,325

Disaggregate Data

Respondent ID	Vehicle	Income	Age	Education
1	2005 Ford Fusion	4	45	2
2	2005 Ford Fusion	4	46	2
3	2005 Ford Fusion	4	49	6
4	2005 Ford Fusion	6	44	2
5	2005 Ford Fusion	5	41	6
6	2005 Toyota Camry	7	30	4
7	2005 Toyota Camry	6	33	6
8	2005 Toyota Camry	6	29	2
9	2005 Toyota Camry	6	43	2
10	2005 Toyota Camry	5	24	5
11	2005 Toyota Prius	6	47	6
12	2005 Toyota Prius	7	42	5
13	2005 Toyota Prius	6	27	6
14	2005 Toyota Prius	5	43	4
15	2005 Toyota Prius	7	33	6
16	2005 Honda Accord	7	43	2
17	2005 Honda Accord	4	48	6
18	2005 Honda Accord	7	50	5
19	2005 Nissan Sentra	4	31	3
20	2005 Nissan Sentra	5	42	2
21	2005 Nissan Sentra	5	26	6
22	2005 Nissan Sentra	4	50	4
23	2005 Nissan Sentra	6	47	5

There are three tables to consider in the diagram; the aggregate data table, the disaggregate data table, and the merged data table. The aggregate data table at the top shows sample data for five vehicles including their sales recorded for the month of October. The aggregate data shown includes selected attributes of these vehicles during the year 2005, as well as the purchase price and cost per mile during the month of October (in cents per mile). The disaggregate data table to the right illustrates a (hypothetical) collection of households that hold each of these vehicles alongside a selected set of household attributes. The arrows show the data flow into the merged data table in the center for the Ford Fusion and the Toyota Prius. The flows are similar for the vehicles with no arrows. The combined dataset translates the household attributes of each disaggregate observation to an equivalent number of observations in the merged dataset. The disaggregate observations are unique and the same as they were in the

41

disaggregate table, whereas the aggregate data are repeated (because each vehicle is the same). The only difference of the aggregate data in the merged data are the sales column, in which the sales is divided by the number of disaggregate observations holding that particular vehicle. Thus, the 4078 Ford Fusions are divided by 5 to roughly equal 816 vehicles each. Each Ford Fusion observation in the disaggregate data now represents 816 new vehicles in the merged data for October 2005. This adjusted sales number is the new weight placed on the observation in the aggregate data. For example, if the merged data table in Figure 2 is the entire dataset, then the total sales of the market is 89,941. The weight placed on a single Ford Fusion observation in the merged dataset is $\frac{816}{89941} = .009$. In the disaggregate dataset, the weight placed on a single Ford Fusion observation is $\frac{1}{23} = .04$. Effectively, the combined data set preserves the in-choice proportion of each of the decision makers, but the market share of the choice is adjusted to match that found within the aggregate data. That is, the total number of sales of Ford Fusions is preserved in the aggregate and merged data. Finally, recall that the disaggregate data requires the assumption that demographic distribution of car buyers through the year is constant, the same merge procedure applies for November 2005 with the same disaggregate data. The aggregate data changes by month, but the disaggregate distributions change annually.

The disaggregate data are treated as if it were endogenously collected, but the use of endogenous datasets for choice models presents a problem with the consistency of the coefficient estimates. The maximum likelihood procedures applied in the estimation of choice models using disaggregate data assume that the dataset is exogenously sampled. Manski and Lerman (1977) discuss this problem, acknowledging the economies of scale

42

obtained by sampling endogenously. They show that endogenous samples can be used for estimation as long as the choice share in the endogenous sample is scaled appropriately by the choice share in the population. They develop a modification to the weighting of the log-likelihood function, which when applied to an endogenous sample will produce a consistent estimator. The discussion that follows explains how the process of merging the data are equivalent to the weighting scheme established by Manski and Lerman (1977).

Using their notation, let $Q(i)$ equal the share of the decision making population selecting alternative i. Let $H(i)$ equal the share of the choice-based sample population choosing alternative i. If the log-likelihood term of each observation within the choice-based sample is scaled by the quotient $Q(i)/H(i)$, the estimator will be consistent. This quotient essentially adjusts the weight of the choice in the sample to match the weight of the choice in the population.

The scaling of the log-likelihood described by Manski and Lerman (1977) is equivalent to the effective scaling that occurs during the merge of data sources described above. To illustrate this point, consider the following notation:

Let: h_i = The number of people that chose alternative i in the disaggregate sample.

Let: H = The total number of people in the disaggregate sample.

Let: q_{im} = The number of sales of alternative i in the aggregate data in month m.

Let: Q = The total number of sales in the aggregate data.

The objective is to show that the data merge process scales each observation within the disaggregate data by same magnitude as the weight $Q(i)/H(i)$. Given the notation stated above, it is readily seen that $Q(i) = \frac{q_{im}}{Q}$ and $H(i) = \frac{h_i}{H}$. The merge procedure breaks each q_{im} vehicle sales into h_i pieces. The length of each piece is $\frac{q_{im}}{h_i}$, which is the number of sales within the aggregate data represented by a single household observation choosing vehicle i. Each household inevitably represents more people when merged with the aggregate data, but the proportional representation relative to the population can go up or down as compared to the disaggregate sample. The proportion of the entire population that is represented by each household in the merged dataset is $\frac{\frac{q_{im}}{h_i}}{Q}$. This is the proportional contribution of a single household choosing vehicle i to the log-likelihood when merged with the aggregate data. This needs to be compared with the proportional contribution of the household to the log-likelihood function were the coefficients to be estimated using the disaggregate data alone. Within the disaggregate sample, the household represents a share within the sample of $\frac{1}{H}$, and this is the proportional contribution of any household to the log-likelihood when estimated on the disaggregate data. Therefore, the relative change in the proportional contribution of a single household to the log-likelihood function estimated with the aggregate data is the quotient of the two proportions.

$$\frac{\left(\dfrac{\frac{q_{im}}{h_i}}{Q}\right)}{\left(\dfrac{1}{H}\right)} = \left(\frac{q_{im}}{Q}\right)\left(\frac{H}{h_i}\right) = \frac{\frac{q_{im}}{Q}}{\frac{h_i}{H}} = \frac{Q(i)}{H(i)}$$

This simple manipulation shows that the relative change in the proportional representation as a result of the merge is equivalent to the scale established by Manski and Lerman (1977). Essentially, this result implies that the coefficients derived from the merged dataset are consistent. But it also implies that the same estimation results obtained by the merge could be obtained by taking any sample of data, computing the ratio of shares, and multiplying each observation within the sample by this scale factor. When the data are exogenous and representative of the market, the sample share will equal the population share. In all other cases, the scale factor adjusts the representation of the choice to reflect the population. The data requirements for both approaches are the same. Information on the population shares is required for the scaling to take place. Otherwise, any disaggregate dataset applied in choice analysis must be assumed exogenous and representative of the market that the analysis is attempting to model.

3.6 Attributes of the Utility Function

Automobiles have fundamental physical attributes such as horsepower, fuel efficiency, interior space, cargo area, weight, etc., and interact with consumers in very direct ways. These variables directly impact how the car drives as well as its utility and cost. A model built primarily with physical attributes has advantages with respect to robustness over time because such attributes are inherent qualities of personal vehicles now and in the future. In addition, the foundation established by using generalized physical attributes and costs better facilitates the introduction of hypothetical vehicles. The specification of new vehicles based on physical attributes can go some distance in assessing how such a vehicle will fare against its likely competition when considering

purchase price and fuel costs. Finally, research has repeatedly found that physical

attributes are effective in explaining much of market performance. Proxies for style are

rarely included and are often not found to be too important when they are. Recently,

Train and Winston (2007) explored a variety of standing hypotheses put forth to explain

the decline of the American automotive industry to foreign competition. Using data from

the 1990s, they found that much of the loss of American market share could be explained

by relative changes in basic attributes, including price, fuel consumption, and

horsepower. Their research suggested that the decline of the American automotive

industry occurred not because of lack of improvement in vehicles, but because those

improvements did not keep pace with the foreign competition in performance or cost.

They found that relative changes in these attributes between American and foreign

competition could explain roughly 93% of the 6.8% loss of market share during the

1990s.

The utility function developed for this study is built from variables that span

technical specifications, quality and consumer attributes. The attributes available for

coefficient estimation are listed in Table 2.

Table 2: Data Fields Applied to the Model

Data Field	Data Source
Manufacturer suggested retail price	*Automotive News*
Customer incentive by manufacturer	*Automotive News*
Fuel economy (mpg)	*Yahoo Autos*
Horsepower	*Yahoo Autos*
Curb weight	*Yahoo Autos*
Passenger volume	*Yahoo Autos*
Vehicle seating	*Yahoo Autos*
Vehicle body rating	*Consumer Reports*
Climate system rating	*Consumer Reports*
Used car value rating	*Consumer Reports*
Exhaust system rating	*Consumer Reports*
Paint rating	*Consumer Reports*
Power equipment rating	*Consumer Reports*
Age of respondent	CMAP and CEC
Count of household members under 5 years	CMAP and CEC
Count of household members ages 5 to 11	CMAP and CEC
Count of household members ages 12 to 15	CMAP and CEC
Count of household members ages 16 and up	CMAP and CEC
Education of respondent	CMAP and CEC
Income of household	CMAP and CEC
Gasoline price	Energy Information Agency

Some of these attributes are combined together in a variety of ways to produce attributes that more directly correlate with the consumer experience with the vehicle. For example, fuel economy and gasoline prices are combined to generate the monthly fuel cost per mile. The customer incentive is deducted from the MSRP to produce an Adjusted MSRP. Finally, all demographic attributes interact with specific vehicle attributes, which is required to introduce variance across choices for demographic attributes.

4. Results

4.1 Sample Statistics of Data

The complete merged dataset results in a sample of 25,159 observations, weighted in accordance with the merge procedure. The exact count of observations is a function of the number of households that own vehicles with each model year and the number of months during each particular year in which the vehicle was sold. Not all vehicles were sold for the entire twelve months of a year as vehicles within the choice set would enter and exit the market. Thus, not every vehicle count within the disaggregate dataset can be multiplied by twelve to obtain the total count of observations. Some would be multiplied by say three or seven depending on the circumstances of the year in which they might have exited or entered the market. Most of the vehicles were present for all 36 months of the aggregate data. The total number of 2005 model year vehicles in the disaggregate dataset is 781. For the 2006 and 2007 model year, the counts are 736 and 597 respectively. This amounts to 2114 vehicles, which only consists of the 56 sedan models included in the choice set.

The disaggregate dataset is comprised of two separate datasets: the CEC dataset of California collected in 2007, and the CMAP dataset collected in 2007 and 2008. The key component that these datasets have is, make, model and year of each vehicle held within the household. Unfortunately, this is not a data field normally collected in the most common population surveys. For example, it is not currently collected in the Current Population Survey or the American Community Survey by the US Census Bureau (US Census, 2009). Ideally, a broader population survey that collected such data

would be superior to data assembled from available surveys of states and metropolitan regions. To show how the demographics vary within the two datasets, Table 3 illustrates the demographic profile across four key attributes.

Table 3: Demographic Profile of Disaggregate Data Sets

<table>
<tr><td colspan="8" align="center">Summary Statistics
Key Disaggregate Household Attributes</td></tr>
<tr><td>Income</td><td><i>CEC</i></td><td><i>CMAP</i></td><td><i>Total</i></td><td>Age of Respondent</td><td><i>CEC</i></td><td><i>CMAP</i></td><td><i>Total</i></td></tr>
<tr><td>Less than $20,000</td><td>3%</td><td>4%</td><td>4%</td><td>10 to 20</td><td>1%</td><td>0%</td><td>0%</td></tr>
<tr><td>$20,000 - $34,999</td><td>9%</td><td>6%</td><td>7%</td><td>20 to 30</td><td>14%</td><td>6%</td><td>7%</td></tr>
<tr><td>$35,000 - $49,999</td><td>11%</td><td>11%</td><td>11%</td><td>30 to 40</td><td>20%</td><td>12%</td><td>13%</td></tr>
<tr><td>$50,000 - $59,999</td><td>8%</td><td>9%</td><td>9%</td><td>40 to 50</td><td>25%</td><td>20%</td><td>20%</td></tr>
<tr><td>$60,000 to $74,999</td><td>16%</td><td>13%</td><td>14%</td><td>50 to 60</td><td>19%</td><td>30%</td><td>28%</td></tr>
<tr><td>$75,000 to $99,999</td><td>20%</td><td>21%</td><td>21%</td><td>60 to 70</td><td>13%</td><td>20%</td><td>19%</td></tr>
<tr><td>More than $100,000</td><td>34%</td><td>35%</td><td>35%</td><td>70 to 80</td><td>6%</td><td>9%</td><td>9%</td></tr>
<tr><td>Total</td><td>N= 343</td><td>N= 1771</td><td>N= 2114</td><td>Total</td><td>N= 343</td><td>N= 1771</td><td>N= 2114</td></tr>
<tr><td>Education of Respondent</td><td><i>CEC</i></td><td><i>CMAP</i></td><td><i>Total</i></td><td>Household Size</td><td><i>CEC</i></td><td><i>CMAP</i></td><td><i>Total</i></td></tr>
<tr><td>Not a high school graduate</td><td>2%</td><td>1%</td><td>1%</td><td>1</td><td>12%</td><td>21%</td><td>19%</td></tr>
<tr><td>High School graduate</td><td>13%</td><td>15%</td><td>15%</td><td>2</td><td>35%</td><td>43%</td><td>42%</td></tr>
<tr><td>Some college credit</td><td>24%</td><td>15%</td><td>16%</td><td>3</td><td>22%</td><td>18%</td><td>18%</td></tr>
<tr><td>Associate or technical school</td><td>9%</td><td>8%</td><td>8%</td><td>4</td><td>18%</td><td>11%</td><td>12%</td></tr>
<tr><td>Bachelor's or undergraduate</td><td>33%</td><td>30%</td><td>31%</td><td>5</td><td>8%</td><td>5%</td><td>5%</td></tr>
<tr><td>Graduate degree</td><td>18%</td><td>30%</td><td>28%</td><td>6</td><td>3%</td><td>2%</td><td>2%</td></tr>
<tr><td>DK/RF</td><td>0%</td><td>0%</td><td>0%</td><td>7</td><td>2%</td><td>0%</td><td>0%</td></tr>
<tr><td>Total</td><td>N= 343</td><td>N= 1771</td><td>N= 2114</td><td>8</td><td>0%</td><td>0%</td><td>0%</td></tr>
<tr><td></td><td></td><td></td><td></td><td>9</td><td>0%</td><td>0%</td><td>0%</td></tr>
<tr><td></td><td></td><td></td><td></td><td>10</td><td>0%</td><td>0%</td><td>0%</td></tr>
<tr><td></td><td></td><td></td><td></td><td>Total</td><td>N= 343</td><td>N= 1771</td><td></td></tr>
</table>

The distribution of household income shows that there is little distinction between the incomes of California and Chicago respondents. Overall, the distribution shows that respondent households earned relatively high incomes, with 50% of the sample earning $75,000 or more. The distribution of respondent education was not quite as congruent. The response variable is the education of the respondent, and does not consider other members of the household, which is a source of variance. Overall, the data shows that the Chicago respondents have slightly higher levels of education. But relative to the US, both sets of respondents are well educated, with more than 50% holding a bachelor's degree or higher. The distribution of age shows that the California respondents were

slightly younger than the Chicago respondents. Nearly 60% of Chicago respondents were older than 50, whereas only 40% of California respondents were older than 50. Finally, the distribution of household size shows that California households were on average bigger than Chicago households. For instance, 30% of the California households were comprised of four persons or more, whereas not even 20% of Chicago households were at least of four people.

   The differences between the Chicago dataset and the California dataset likely stem from the fact that Chicago is a cosmopolitan city and California is a very diverse state with urban and rural environments. Thus Chicago's smaller households and higher relative education are expected given the strictly urban environment. However, it is surprising that the income distribution between the two populations are not very distinct. The dataset overall is well-educated with comparatively high incomes. Therefore, it is not representative of the population as a whole. But the overall population of the country is not representative of new car buyers. It is important to emphasize that these are households that own a vehicle with a model year of 2005, 2006 or 2007 in a survey taken soon after these years. That is, these are households that own very new cars, and likely bought them new. Given the data challenges outlined earlier, there currently is no publicly available national demographic profile of new car buyers within the sedan or any other car market. The public at large faces the same data challenges encountered by this study.

4.2 The Choice Set

The choice set includes vehicles that entered, exited or remained in the market during the full three-year period. Table 4 presents the vehicles that are included in the choice set.

Table 4: Sedan Models Considered in the Choice Set

Vehicle Class	Vehicles			
Subcompact	Chevrolet Aveo	Ford Focus	Honda Fit	Hyundai Accent
	Nissan Versa	Pontiac Sunfire	Pontiac Vibe	Subaru Impreza
	VW Eos	Toyota Yaris		
Compact	Chevrolet Cobalt	Dodge Charger	Honda Civic	Honda Civic Hybrid
	Kia Spectra	Mazda Mazda3	Mitsubishi Galant	Nissan Sentra
	Hyundai Elantra	Pontiac Grand Am	Saturn Ion	Toyota Corolla
	VW Jetta			
Midsize	Buick Century	Chevrolet Impala	Dodge Stratus	Ford Fusion
	Honda Accord Hybrid	Hyundai Sonata	Kia Optima	Mazda Mazda6
	Nissan Altima	Nissan Altima Hybrid	Pontiac G6	Subaru Legacy
	Honda Accord	Mercury Milan	Toyota Camry	
	Toyota Camry Hybrid	Toyota Prius	VW Passat	
Fullsize	Buick LaCrosse	Buick LeSabre	Buick Lucerne	Buick Park Avenue
	Chevrolet Malibu Hybrid	Chrysler 300	Ford Five Hundred	Ford Taurus
	Chevrolet Malibu	Hyundai XG350/Azera	Toyota Avalon	Pontiac Grand Prix
	Kia Amanti	Mercury Sable	Nissan Maxima	

The estimation of models in this study was done with software called BIOGEME 1.8 (Bierlaire, 2003).[1] BIOGEME is capable of estimating multinomial logit (MNL), nested logit (NL), cross-nested logit (CNL), mixed logit models, and Network GEV models. The Network GEV is a new versatile model structure that can be used to construct any GEV model including the GNL model.

Several model structures were tested on the data. The logit, nested logit, cross-nested logit, and generalized nested logit, were all considered with similar utility

---

[1] http://transp-or.epfl.ch/page63023.html

specifications. The logit estimation was the most successful. It produced coefficient estimates of theoretically consistent sign and statistical significance. The nested logit yielded reasonable coefficients, but the logsum coefficients could not universally converge to values that were within an acceptable range. The logsum coefficient must be between 0 and 1 for the model to be consistent with random utility theory. The cross-nested logit model and the more demanding generalized nested logit model had too many variables to solve with the available estimation algorithms. Even with simple specifications of the utility function, the estimation procedure could not settle on consistent proportional allocation parameters. This suggests a few ideas for further work with GEV models of automotive choice. The network GEV models (such as the CNL and GNL) that permit proportional allocation of alternatives across nests offer a theoretically sound approach to addressing the complex substitution patterns that can arise in choice settings in which alternatives are grouped, yet ordinal in nature. Within the automotive market, there are few definitive divisions between vehicle types, but many more subtle divisions associated with vehicle class. The network GEV structures that permit overlapping nests offer a promising means to represent this fluid pattern of substitution within the context of GEV models. But the computational cost is high and grows rapidly when the number of alternatives and nests grow large. The development of network GEV models is relatively new, and hence there is still considerable exploration to be done with specific applications. Given existing estimation procedures, model structures containing a large number of alternatives can be extremely difficult to solve or to prove that a global optimum has in fact been found.

4.3 Estimation Results of the Logit Model

The logit estimation provides a useful perspective on the attributes that affect automotive choice. The complete merged dataset constitutes 25,159 observations for all data spanning 2005 - 2007. The utility function is specified to be linear and the estimated coefficients are shown in Table 5.

Table 5: Estimation of Logit Model

Description	Coefficient	Robust T-Statistic	p-value
Compact Dummy X Income X Age	0.0029	23.61	0.00
Midsize Dummy X Income X Age	0.0047	26.48	0.00
FullSize Dummy X Income X Age	0.0039	35.85	0.00
CR Fuel System Rating	0.1120	12.39	0.00
CR Rattle Rating	-0.1560	-16.36	0.00
CR_Body Rating	.0529	5.01	0.00
CR_Climate Rating	.0446	4.55	0.00
CR_Exhaust Rating	.2510	10.31	0.00
CR_Paint Rating	.0075	0.93	0.35
CR_Power Equipment Rating	.0493	5.00	0.00
CR Used Car Rating X Education of Respondent	0.0112	4.07	0.00
Vehicle Seating X Children Under 5	.1560	2.98	0.00
Power:Weight Ratio	10.2000	7.82	0.00
MSRP minus Customer Incentive	-0.6930	-24.83	0.00
Cost per mile (cents per mile)	-0.0362	-4.62	0.00

* "X" implies the interaction between attributes.

The results offer several important insights into vehicle attributes that influence automotive choice. All variables within this model are highly significant with the exception of the *Consumer Reports* rating for Paint. All but one of the variables are of the theoretically expected sign. The *Consumer Reports* rating for Rattle should be

positive, since all of the organization's ratings attribute higher values to superior quality (less of a negative attribute). The remaining variables show that vehicles that perform well in the *Consumer Reports* rating criteria generally perform better in the market. A set of variables incorporate the interaction between income age and vehicle class. Dummy variables are established for vehicles in the Compact, Midsize and Fullsize class. These dummy variables are set relative to the Subcompact class and are interacted with Income and Age. That is, the higher the combined age and income, the larger the vehicle. As expected, dummy coefficient estimates are positive and significant. But the magnitude is unexpected, as the midsize dummy is larger than the full size dummy. The compact dummy is smaller than both. This reversal in size, when interpreted literally suggests that the combined effects of age and income are correlated with a greater attraction to the midsize sedan over the full size sedan. This seems contrary to what one might expect. There are several possible explanations for this result. The midsize sedan is a larger class and contains hybrid vehicles, which do not currently exist in the full size class. The full size sedan class may attract consumers with a particular taste that is less identifiable with simple demographics, while other people with similar age and income find vehicles in the midsize category to be satisfactory. The intent of the vehicle class dummy variables is to permit the model to predict the probability that an older, wealthier consumer purchases a larger vehicle. Another interaction variable combines the vehicle seating capacity with the number of children in the household, to reflect the notion that households with more children require vehicles with greater maximum seating capacity. This is in contrast to household size, which can result in additional vehicles, if all members are adults. Finally, the model contains an interaction variable that combines education and the *Consumer*

55

Reports used car rating. The proposition here being that increased education raises the propensity of the consumer to consider the long-term value of the vehicle.

The final two coefficients listed are the cost parameters of the vehicle. As is expected, both are negative and significant. The price variable is coded as the difference between the MSRP and the average reported customer incentive for that vehicle. This allows the price variable to vary over each month. This variance is important because customer incentives shift over the course of the year, which changes the competitive position of vehicles. During this period covered by the data, domestic automakers would discount their vehicles throughout the year, essentially adjusting the effective price to be lower than the list price.

The cost per mile parameter is strictly a function of the combined fuel economy and the average national gasoline price. It is technically the fuel cost per mile and it shows that vehicle purchase activity shifts to more efficient vehicles during months in which the gasoline price is higher. The in-year variance of gasoline prices during this period is considerable. Figure 3 shows a plot of the average real monthly national gasoline prices from 2005 through 2007.

Figure 3: Average Monthly National Gasoline Prices from 2005 to 2007
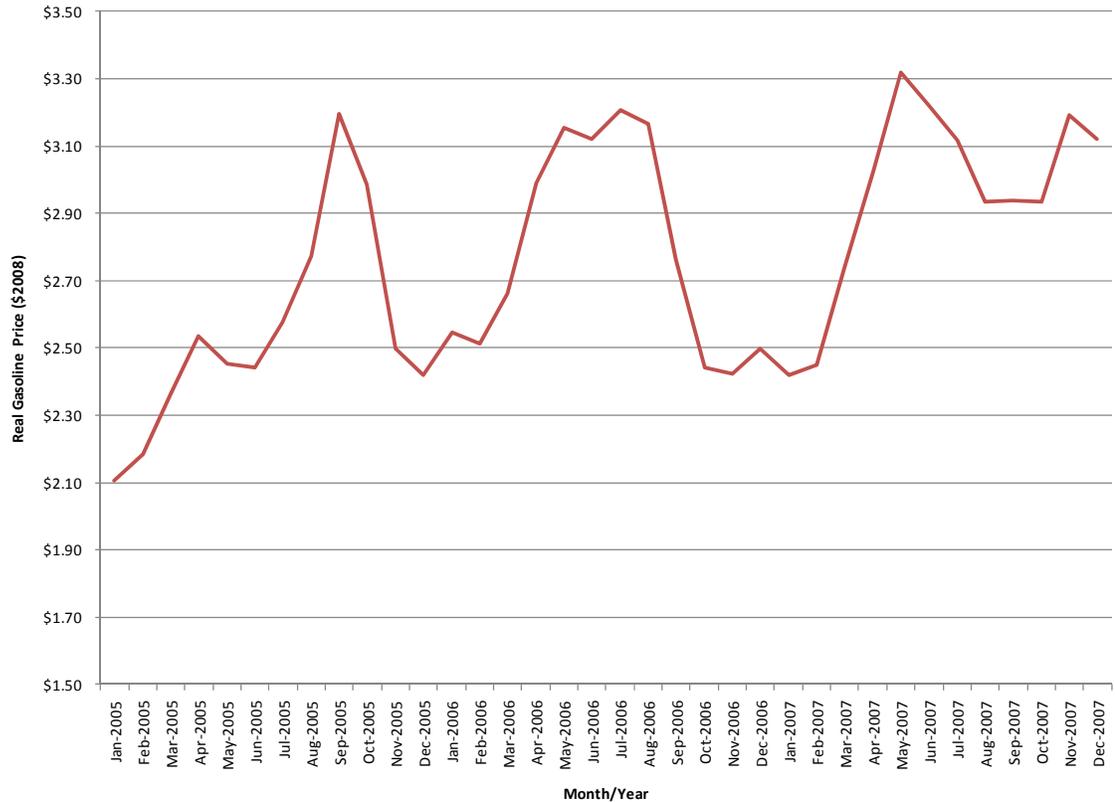
While the negative sign on the cost parameter for gasoline prices is an expected result, it is not always obtained. The automotive market contains many attributes that are desirable to consumers at the expense of fuel economy. These attributes, such as horsepower and size, comprise what the consumer is actually paying for. This model captures these effects well enough such that the optimization procedure does not attribute the higher cost per mile of larger cars as the attribute that draws consumers to those cars.

The estimation of the model illustrates the feasibility of developing a choice model using aggregate sales data with supporting disaggregate data. The combined dataset contains a revealed preference of nationwide choice behavior. Supporting

information supplied by a disaggregate sample informs the model on how distinct demographic attributes influence choice. While the model has several core strengths, it also has a variety of weaknesses that should be targeted for improvement. The section that follows outlines some of these weaknesses and areas for future research within the model structure.

4.4 Model Weaknesses

The model shows that a dataset assembled with combination of aggregate and disaggregate data can produce a collection of coefficients that are mostly of consistent theoretical sign. However, a number of weaknesses are present that ideally could be improved. The income by age dummy variable is established to capture the interactive effect that older people with higher incomes are attracted to specific vehicle classes. The intuition of this construction is that a person interested in a full size sedan is generally both older and of high income, but not uniquely one or the other. But an undesirable aspect of this specification is that it uses two demographic variables in a single dummy. Often, it is better to maximize the interaction of these scarce variables across as many model alternative attributes as possible. To be sure, many specifications were tried in anticipation of the specification presented, and several of those specifications offered comparable results with a slightly different selection of attributes and interacted terms. The model presented here had a final log-likelihood higher than most, and was selected for this reason. But by no means does this imply that it is the best model for automotive demand or that it is even the best model that could be constructed by this dataset. There are several modifications to the specification that are worth exploring for likely

improvement. For example, this model is linear in attributes, but several of the attributes are promising in a non-linear specification. The influence of age on choice could be quadratic, and even the effect of fuel cost on choice may find an improved fit in a non-linear specification. The improperly signed rattle variable deserves additional consideration as it may capture a dummy effect associated with a few vehicles that perform poorly in the market, but received favorable scores by *Consumer Reports* on this attribute. In addition, there may be opportunities to explore data reduction techniques that reduce the parameters presented by *Consumer Reports* through principal components analysis. This analytical technique can sometimes extract a simple and deep structure that isolates the core information of a large set of parameters into a small number of variables.

5. Policy Analysis

This section presents a policy analysis that explores several important questions with respect to new vehicle purchases and policy design. The model is used to perform policy simulations that explore the effectiveness of hybrid tax credits in reducing greenhouse gases through a shift in new vehicle purchases. In addition, hybrid tax credits and changes to the gasoline tax are compared in their ability to shift the average fuel economy of the new car fleet. Finally, the model is used to evaluate the average willingness to pay for any attribute as represented by the ratio of coefficients. In particular, there is considerable interest among automakers in consumer willingness to pay for fuel economy. This section ends with a discussion of the average willingness to pay for fuel economy that is implied by the coefficient estimates.

5.1 The Impact of Hybrid Tax Credits on Greenhouse Gas Emissions

A policy simulation incorporates several ingredients to develop conclusions on the effectiveness of policy. The estimated parameters of the model constitute the primary ingredient. But the choice model will only produce probabilities. To forecast impacts, a total number of sales is required.

The period of data in which the model is estimated spans a time when hybrids started to gain appreciable market share as well as a time when federal incentives for hybrid vehicles were introduced. A natural question that arises from this period pertains to how effective the tax incentives were in shifting demand towards hybrids and subsequently reducing greenhouse gases. The Internal Revenue Service (IRS) computed

the incentive based on a formula that considers the increment in fuel economy offered by the vehicle. These tax credits were phased out gradually as manufacturers hit a sales target of 60,000 vehicles. Table 6 shows the schedule of the tax credits initiated in 2006.

Table 6: Schedule of Sedan Hybrid Incentives During Study Period

Table of Incentives						
Start Period	End Period	Accord Hybrid	Civic Hybrid	Altima Hybrid	Camry Hybrid	Prius
January-06	September-06	$650	$2,100	$2,350	$2,600	$3,150
October-06	March-07	$650	$2,100	$2,350	$1,300	$1,575
April-07	December-07	$650	$2,100	$2,350	$650	$788

The data used to estimate the model includes the tax credits as a deduction from the MSRP. The policy simulation that evaluates the impact of the tax incentives is historical in nature. It proceeds by simulating the sales of automobiles using the estimated parameters with the original data, thus providing a baseline forecast of sales over the three year span. The simulation is then run with tax credits removed, while the remaining events within the data stay the same (e.g., manufacturing incentives, gasoline prices). The difference in sales between the baseline forecast and the forecast with tax credits removed illustrates how the market shifted as a result of the tax credits. The hybrids will see a drop in sales, while the remaining vehicles will experience an increase in sales. Given this information, the impact and cost effectiveness of the policy can be measured. One weakness of the tax incentive policy is the ubiquitous disbursement of incentives to all buyers, including those who would have bought the vehicle without them. The environmental impact of the policy only considers those sales that are derived

as the marginal change between scenarios. The tax credit was only applied in 2006 and

2007 and the resulting change in sales during these years is presented in Table 7.

Table 7: Change in Sales of Hybrid Vehicles Due to Federal Tax Credits

Hybrid Sedan	2006	2007	Total
Accord Hybrid	7823	7983	15806
Civic Hybrid	11144	10663	21806
Altima Hybrid	0	19955	19955
Camry Hybrid	26255	7652	33907
Prius	39413	8412	47825
Total	84635	54664	139299

Because the simulated change is removing an incentive, hybrid sales fall.

Therefore, the numbers in the table above reflect forecasted sales that happened because

the tax credits were applied. The simulation of environmental impact proceeds by

placing assumptions on the vehicle life of all vehicles sold (not just hybrids) as well as

assumptions regarding annual miles driven. A baseline assumption is that vehicles have a

working life of ten years and are driven 10,000 miles per year during this period. Given

these assumptions, a general quantification of the impacts of the hybrid vehicle tax

credits can be evaluated simply by calculating the change in the total gasoline

consumption that results from the two different scenarios. Table 8 illustrates some of the

basic metrics of the policy given these assumptions.

Table 8: Baseline Policy Impacts

Category	2006	2007	Total	Units
Impact				
Total Change in Gallons	113,013,000	53,510,000	2,533,171,268	gallons of gasoline
Total Change in Emissions	1,047,000	496,000	1,543,000	t of GHG
Cost				
Policy Expenditure	1,414,103,000	950,999,000	2,365,102,000	$
Dollar per Metric Ton	1,351	1917	1533	$ / t of GHG

Table 8 shows the total change in the consumption of gasoline gallons and greenhouse gas emissions under the strict baseline assumption that each car sold is driven 10,000 miles per year for ten years. The column describing the total impacts suggest that the tax credits will displace 1.5 million metric tons of GHG emissions based on sales over the two year period at a cost of $1533 per metric ton. The separation of impacts across years shows an asymmetric contribution to the total. In 2006, the change in emissions and the policy expenditure is larger than in 2007, but the cost effectiveness of the policy is also better. The reason for this shift in impact is due to the changes in the incentives and the market. As Toyota hit its sales benchmark in late 2006, the phase out of the Prius incentive began. This led to a halving of the incentive in October 2006 and in April of 2007; the incentive fell to a quarter of its original value. This dramatic drop on the incentive portends a drop in the cost-effectiveness of the incentive. This result is driven by the "S" shape of the cumulative logistic curve described in Chapter 3. All vehicles within the choice set capture shares that are far below 50%. Because the logistic curve is non-linear and exponentially increasing at all probabilities lower than 50%, any policy

that improves the utility of the vehicle will have increasing returns to scale until a share of 50% is reached. For this reason, larger incentives induce a relatively higher proportion of people to shift vehicles. A policy with higher incentives is more expensive overall, but the cost-effectiveness is improved. This dynamic, which is strictly a function of the shape of the logistic curve, illustrates one of its main conceptual attractions. The intuition of consumer decision making is supported by the mathematics. For example, a tax credit policy that gives each consumer $100 for buying a hybrid is likely to have no impact on new vehicle purchase decisions given the small incentive relative to the purchase cost. But because all buyers receive the $100 anyway, the small incentive policy is very expensive with almost no results. The public simply pays $100 for each hybrid sold. At the extreme opposite, an incentive of $10,000 per hybrid will be far more expensive but also more successful in drawing demand simply because it creates such a large opportunity cost for not selecting the vehicle endorsed by public policy.

The policy conclusions stated thus far are based on fixed assumptions that are applied equally and universally across all vehicles. But in reality, vehicles have different lifetimes and are driven different average annual distances over those lifetimes. A sensitivity analysis can illustrate how the distribution of impacts varies with assumptions on the variance of vehicle factors. This sensitivity analysis simulates different assumptions of vehicle life and miles driven simultaneously across all vehicles. This simulation can be repeated many times with different assumed values drawn for each instance.

The simulation assumes that both the values of vehicle life and vehicle miles are drawn from a normal distribution. The vehicle miles are drawn from a distribution of $N(10000, 3000)$, and the vehicle life parameters are drawn from a distribution of $N(10, 1)$. The simulation is run through a program known as Crystal Ball®, which can record the selected output values with each new draw of the input parameters (Oracle, 2008). Figure 4 shows the distribution of the average impact per year as a result of the sensitivity analysis.

Figure 4: Distribution of Tax-Credit Cost Effectiveness (t GHG avoided/policy year)
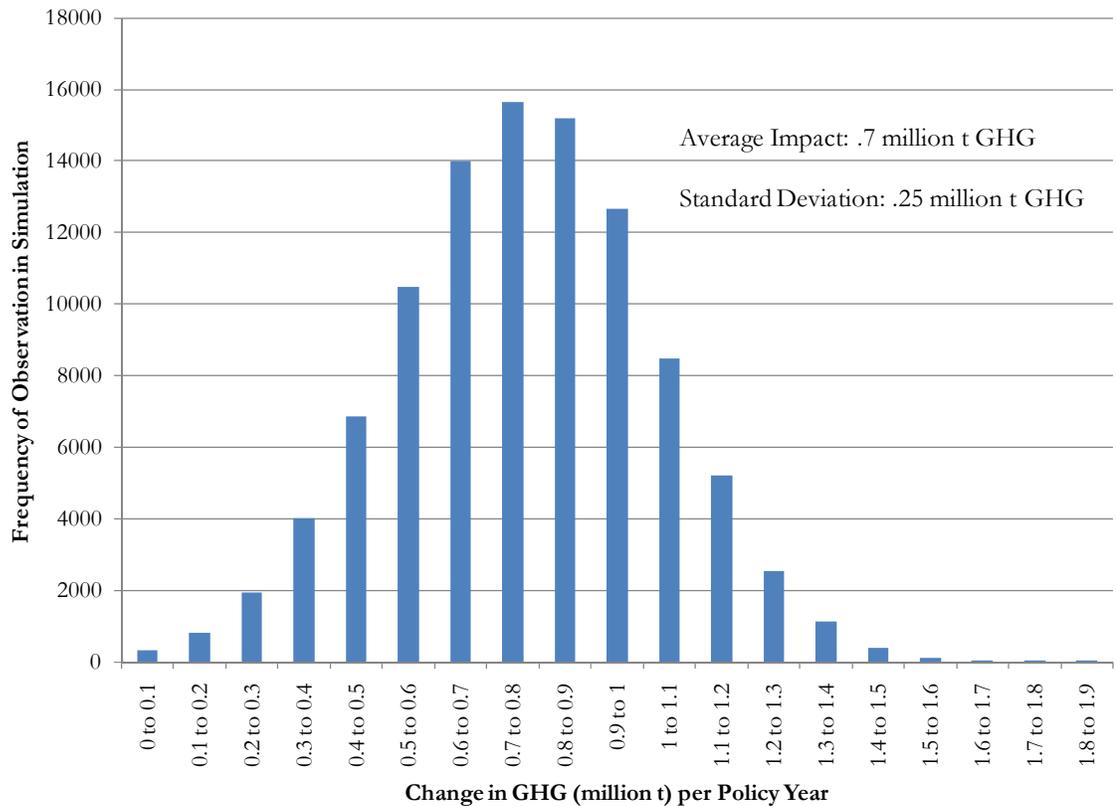
Figure 4 shows the average impact per policy year. The distribution shows that the assumptions governing vehicle age and miles driven can introduce considerable

variance in the impact of tax credit policy. If the vehicles that are endorsed by a policy live a short time and are driven a low annual mileage, the impact of the policy will be small. However, only 7% percent of the observations are less than 400,000 t GHG per policy year, and there is a little circumstantial evidence to suggest that hybrids are universally driven a short distance. This claim is supported by an evaluation of the distribution of the annual mileage placed on a sample of hybrid drivers. Figure 5 shows the distribution of annual mileage that is driven by respondents of a survey of carsharing members in North America that are part of households that own the Toyota Prius or the Honda Civic Hybrid.

Figure 5: Distribution of Annual Miles Driven by Hybrid Civic and Prius Drivers
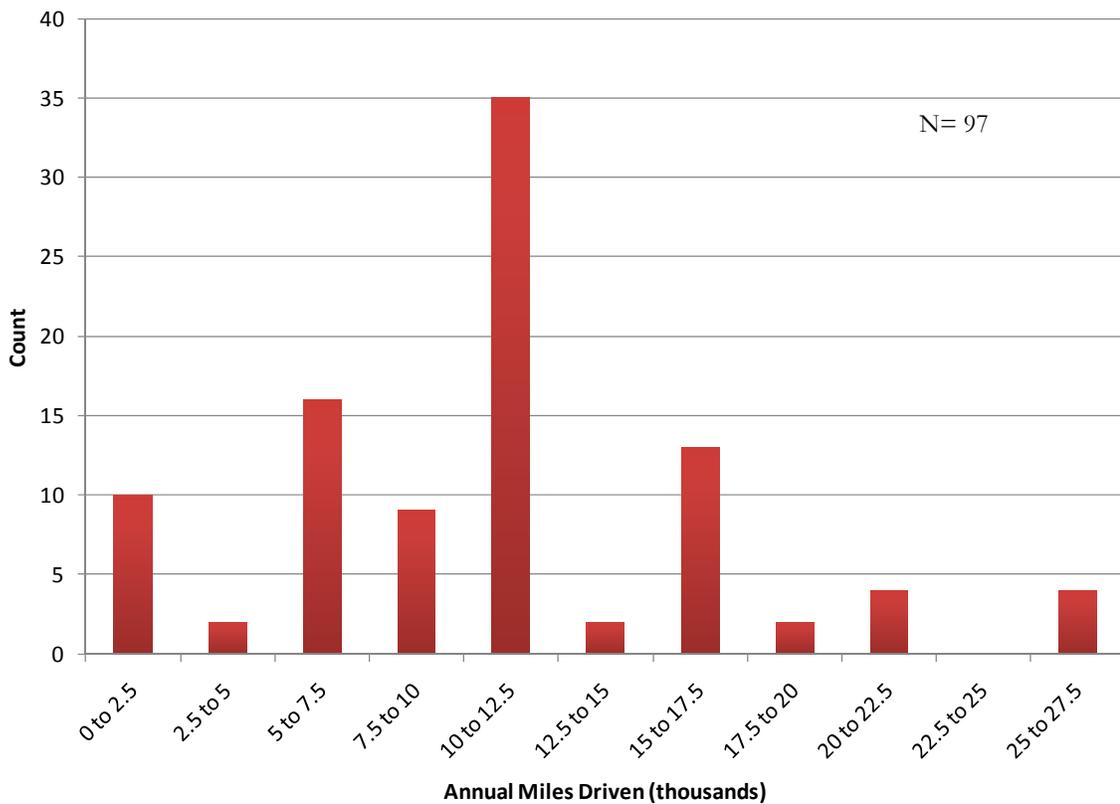
The distribution exhibits the outline of a normal distribution with a mean close to 10,000 miles per year. The sample is derived from the population of carsharing members, which may exhibit a bias towards lower mileages due to their propensity to live in regions with high population density and better transit. But even households in this population show considerable variance in the annual miles driven. The variance roughly matches the range and distribution of annual miles driven as a result of the sensitivity analysis.

Figure 4 shows that the impact of the policy is sensitive to assumptions, but it also shows that the true impact of the policy is likely between .5 and 1 million t GHG / per policy year. As the impact of the policy has a distribution, so does the cost effectiveness of the policy as a result of the sensitivity analysis. The distribution of the cost effectiveness is a function of the change in the GHG impact of the policy as compared to the policy expense. The policy expense is entirely a function of the choice model and does not change as a result of the sensitivity analysis. The cost effectiveness of the policy is computed as before, simply a division of the expense of the policy by the computed GHG impact. This distribution is illustrated by Figure 6.

Figure 6: Distribution of Cost-Effectiveness of the Hybrid Tax Credit Policy
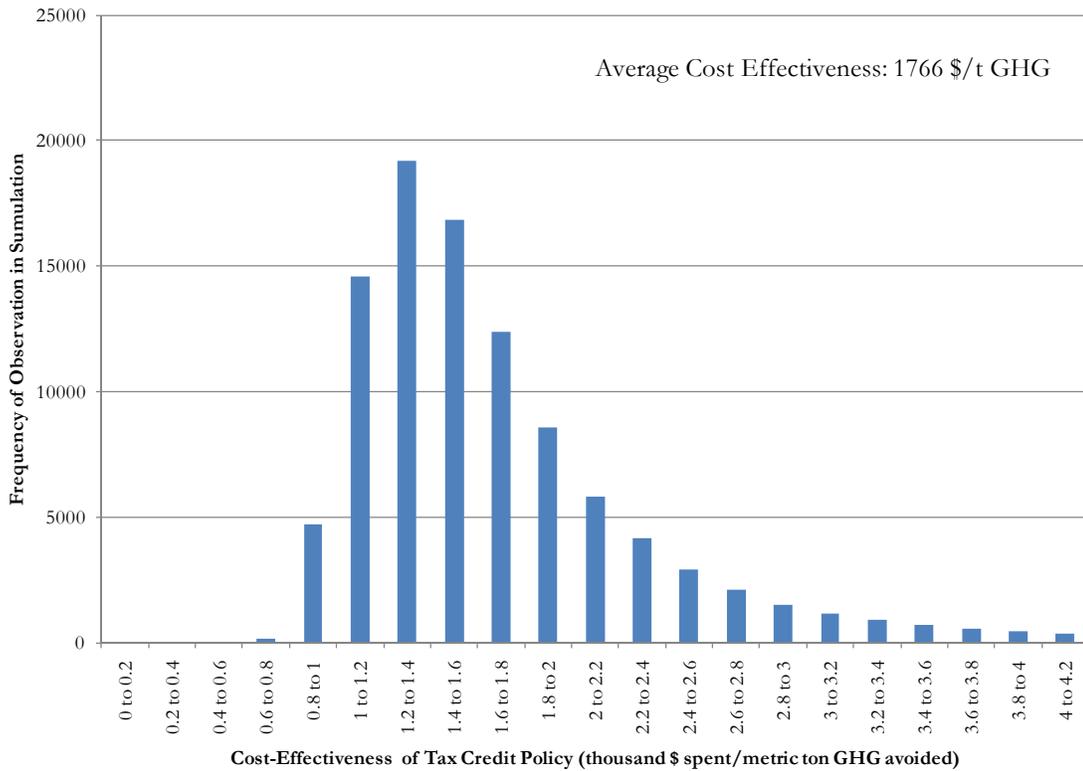
The distribution of cost-effectiveness falls within a far tighter range than the distribution of the GHG impact. This distribution also exhibits a different shape, which is more representative of a chi-squared distribution as opposed to a normal distribution. The reason for this shape and tightness is due to the fact that the numerator is fixed while the denominator varies as a normal distribution. For draws in the sensitivity analysis that generate a high policy impact, the resulting cost-effectiveness improves (is lower). But as the GHG impact of the policy increases, the marginal improvement in the cost-effectiveness of the policy declines. The opposite occurs at the other end of the distribution. As the impact of the policy declines, the marginal worsening of the cost-effectiveness increases. This generates the long tail towards less attractive amounts of

dollars spent per metric ton GHG reduced. The shape of this distribution suggests that the cost effectiveness of the hybrid tax credit is bounded. It appears to be not better than $600 per metric ton of GHG avoided. But the majority of values (91%) derived from the sensitivity analysis are contained between $1000 and $3000 per metric ton. The model suggests that the cost-effectiveness of the hybrid tax credit policy during 2006 and 2007 falls within this range.

5.2 Gasoline Taxes and the Reduction of Greenhouse Gases

The model can evaluate the impact of raising the gasoline tax for inducing a shift towards more efficient vehicles. The gasoline tax is often a policy that is academically debated but rarely considered due to political resistance. The current gasoline tax is broken up into a state component and a federal component. The federal component is 18.7 cents per gallon and the state tax naturally varies by state. On average across states, the combined gasoline tax is currently 47 cents.

The policy simulation evaluates how sales of vehicles would have changed had the gasoline tax been higher. The simulation is conducted on the historical data over 2005 to 2007. For each simulation, the gasoline tax is raised by 10 cents per gallon. The simulated tax increments by 10 cents from 10 up to 40 cents per gallon, which is likely at or above any politically feasible gasoline tax at this time. There are two ways in which an increased gasoline tax can reduce emissions. Recent rises in gasoline prices have crossed a threshold at which Americans have reduced their driving due to high fuel costs (Krauss, 2008). A rise in gasoline prices can also cause a shift in vehicle purchases. The

GHG emissions resulting from this shift in purchases can be evaluated by the policy simulation.

The simulation establishes a baseline forecast of sales that occur over the course of the three year estimation period. The simulation is then repeated with gasoline prices incremented 10 cents higher, until the gasoline tax is a total of 40 cents over baseline taxes. Each vehicle is impacted by the policy such that vehicles that are relatively more efficient gain with respect to the broader choice set. The gasoline tax policy also differs from the hybrid tax credits in that the policy earns revenue for the government. In addition, the costs of the policy are born by society, as people not in the market for new vehicles pay the cost of the policy. For these reasons, the cost effectiveness metric does not apply in the same way as with the hybrid tax credits. But the policies can be compared according to GHG emissions and shifts in the average fuel economy of the new vehicle fleet. Table 9 shows a baseline comparison of the gasoline tax and hybrid tax credits. This comparison is done with the standard assumptions that all new vehicles have a ten year vehicle life and are driven 10,000 miles per year.

Table 9: Comparison of Gasoline Tax and Hybrid Tax Credit Policy

Policy	Average Fuel Economy	GHG Emissions Change per policy year
Gasoline Tax		
Base	30.54	0
10 cent tax	30.55	-56022
20 cent tax	30.56	-112131
30 cent tax	30.58	-168328
40 cent tax	30.59	-224612
Hybrid Tax Credit Policy		
No Tax Credits	30.34	0
Tax Credits	30.54	-771266

The comparison shows that the magnitude of change in the average fuel economy is roughly four times greater with the hybrid tax credits than the gasoline tax. The change in average fuel economy is not large for either policy, but moving average fuel economy is inherently difficult given the large and diverse vehicle pool. The change in the greenhouse gases also shows that the tax credit policy is more effective in inducing reductions through the shift of new vehicle purchases. It is important to emphasize that this comparison only considers impacts as a result of this shift. Impacts resulting from a general reduction in driving that might also accompany an increase in the gasoline tax are not estimable with this model. Thus, from the perspective of adjusting the vehicle fleet towards more efficient vehicles, hybrid tax credits are superior to the gasoline tax. However, if the effect of VMT is strong enough, then the gasoline tax may be more effective in reducing overall emissions through the primary effect of reducing driving.

This impact would include the behavior of the entire population and would need to be considered separately for a comprehensive comparison of the two policies.

5.3 Willingness to Pay for Fuel Cost Reductions

Based on the utility function, the valuation of an attribute is computed as the ratio of the attribute coefficient to the vehicle price coefficient. The market valuation for any attribute k is simply:

$$W_k = \frac{\beta_k}{\beta_p}$$

Where β_p is the coefficient of the vehicle purchase price. This ratio expresses the change in purchase price that is necessary to equalize consumer utility given a unit change in attribute k. The sign of the price coefficient is negative. Hence, the valuation quotient will always be positive if the unit change of the attribute is in the direction that increases utility. For an attribute with a positive coefficient, this constitutes an increase in that attribute. For a negative attribute, a drop in the attribute increases utility, and will yield a positive valuation. Since both fuel cost and purchase price are negative, the willingness to pay for reductions in fuel cost is positive and illustrated as follows:

Coefficient	Estimated Value	Units	Scale in Model Estimation
Adjusted MSRP	-0.693	$ / vehicle	(1 / 10,000)
Fuel Cost	-0.0362	$ / mile	(100)

$$\frac{\beta_{FuelCost}}{\beta_{MSRP}} = \frac{-0.0362}{-0.0000693} = 522 \ \frac{\$}{\cent \, fuel \, cost \, per \, mile}$$

The interpretation of this value should be understood as the willingness to pay for fuel cost reductions. This is conceptually close to the willingness to pay for fuel economy, and perhaps as close as any metric could get. The willingness to pay for fuel economy will always depend on the price of fuel. The changing price of fuel provides an important input into the trade-off faced by a consumer in considering the value of fuel economy. This money saved by increased fuel economy constitutes the private benefit of fuel economy. There is also a public benefit to fuel economy in the form of reduced pollution that some people value privately. Previous research of hybrids has suggested that consumers do not pay exclusively for the monetary benefits of fuel savings (Heffner, 2007). These two components come together to characterize the total willingness to pay for fuel economy. If the price of gasoline was close to zero, then the willingness to pay for fuel economy would consist entirely of the private value of a public benefit. If gasoline also had zero external costs, then fuel economy would have no value.

In some applications, improved efficiency results in increased consumption. But recent research has found a declining rebound effect resulting from improved efficiency of vehicles over the past few decades. According to Small and Van Dender (2007), increased efficiency over the decades has not led to increased driving by individuals. Although aggregate VMT has continually climbed, average VMT has remained steady with rising efficiency. This is likely because vehicle speeds have not changed significantly, and there is a natural limit to how much time a person can spend driving during a 24-hour period.

The valuation derived from the ratio of coefficients is inclusive of both the compensated and uncompensated willingness to pay for fuel cost reduction. In other words, the degree to which consumers pay for fuel economy to save money versus lower pollution is not immediately apparent with this number. The degree to which the willingness to pay metrics represent these separate components can be elucidated by analyzing the conditions under which the metric implies compensated or uncompensated payments. Three factors are key in this evaluation. They are the miles driven annually, the life of the vehicle, and the consumer's discount rate. The higher the miles driven, the more the savings from improved fuel economy. The higher the discount rate, the lower the value of future savings. While the annual miles driven is a collectable and reportable data point, the discount rate of an individual consumer is a very difficult number to pin down and can change over time and circumstances. Thus at best, an evaluation can illustrate the degree to which the uncompensated willingness to pay is positive over a likely range of discount rates. The savings that a consumer receives is calculated as the present value of a stream of annual savings received throughout the vehicle life. This present value is defined by the equation $s\frac{1+(1+r)^l}{r}$, where s is the value of the annual savings, l is the implied vehicle life, and r is the implicit discount rate. As the discount rate rises, the value of savings in the distant future declines. Table 10 shows a compilation of two figures that illustrate how the market willingness to pay varies with annual miles driven, the consumer discount rate, and a fixed vehicle life of 10 years.

Table 10: Sensitivity Analysis of Willingness to Pay for Fuel Cost Reduction

Vehicle Life 10 years		Sensitivity Analysis of Expected Savings from Reduced Fuel Cost						
Annual Miles	Annual Savings	Discount Rate						
		≈ 0	5%	10%	15%	20%	25%	30%
5000	$50	$500	$386	$307	$251	$210	$179	$155
6000	$60	$600	$463	$369	$301	$252	$214	$185
7000	$70	$700	$541	$430	$351	$293	$250	$216
8000	$80	$800	$618	$492	$402	$335	$286	$247
9000	$90	$900	$695	$553	$452	$377	$321	$278
10000	$100	$1,000	$772	$614	$502	$419	$357	$309
11000	$110	$1,100	$849	$676	$552	$461	$393	$340
12000	$120	$1,200	$927	$737	$602	$503	$428	$371
13000	$130	$1,300	$1,004	$799	$652	$545	$464	$402
14000	$140	$1,400	$1,081	$860	$703	$587	$500	$433
15000	$150	$1,500	$1,158	$922	$753	$629	$536	$464

Vehicle Life 10 years		Sensitivity Analysis of Willingness to Pay for						
Annual Miles	Annual Savings	Discount Rate						
		≈ 0	5%	10%	15%	20%	25%	30%
5000	$50	$22	$136	$215	$271	$313	$344	$368
6000	$60	-$78	$59	$154	$221	$271	$308	$337
7000	$70	-$178	-$18	$92	$171	$229	$272	$306
8000	$80	-$278	-$95	$31	$121	$187	$237	$275
9000	$90	-$378	-$173	-$31	$71	$145	$201	$244
10000	$100	-$478	-$250	-$92	$20	$103	$165	$213
11000	$110	-$578	-$327	-$154	-$30	$61	$130	$182
12000	$120	-$678	-$404	-$215	-$80	$19	$94	$151
13000	$130	-$778	-$481	-$276	-$130	-$23	$58	$120
14000	$140	-$878	-$559	-$338	-$180	-$65	$22	$90
15000	$150	-$978	-$636	-$399	-$230	-$107	-$13	$59

The two components of Table 10 illustrate the same information in two different ways. The table at the top of Table 10 entitled "Expected Savings from Reduced Fuel Cost" shows the total expected savings that would be obtained from a reduction in the fuel cost per mile of a vehicle by 1 cent. For example, if a vehicle were driven 5000 miles for 10 years, the nominal value of the total reduction in the fuel cost would be $500 at the end of the ten years. For a consumer with a discount rate of zero, in which the value of money is the same today as it is ten years from now, the appreciated savings is

the full $500. But rarely do consumers carry a discount rate of zero, simply because there is an opportunity cost of storing money in an asset that does not appreciate. Hence the remaining columns of the table describe how the appreciated savings are adjusted with higher discount rates. The rows of the table illustrate how these values change with increases in the annual miles driven. The table immediately below the computed savings subtracts the savings from the computed willingness to pay of $522. The cells shaded in grey delineate the contiguous region of values that are positive. They illustrate the conditions under which a consumer is willing to pay for fuel economy more than savings in fuel cost over the life of the vehicle. The shape and values within the shaded region suggest that on average consumers that drive low mileage are more likely to exhibit a positive uncompensated willingness to pay for fuel cost reductions. Intuitively, these consumers receive less monetary benefit from fuel cost savings. Only those driving 5000 miles or less have a positive uncompensated willingness to pay for fuel cost reductions at a zero discount rate. As the discount rate rises, a greater range of mileage yields positive values. At discount rates of 5% to 10%, which may be within a reasonable range for new car buyers, those driving up to 8000 miles a year will have a positive uncompensated willingness to pay for fuel cost reductions. That is, these drivers are paying for the reductions in the externalities of fuel use out of their own pocket.

The results give policy makers a sense of the degree to which the market is paying to save private costs or public costs. The results suggest that most consumers consider the private costs primarily. But it also shows that the market exhibits a reasonable set of circumstances under which consumers may be paying for the reduction of public costs.

6. Summary and Conclusion

This study contributes to both methods and policy, applying several innovations in the generation of data to estimate choice models for the evaluation of policy in the automotive market. The research developed programs that collected data from various sources on the internet to assemble a unique collection of variables associated with automobiles of interest. The data are collected for a recent period from 2005 to 2007, in which hybrids had gained prominence in the general sedan market. In addition to the aggregate data of sales, incentives and vehicle attributes, the study also collected disaggregate data derived from household surveys in California and Chicago, Illinois. This disaggregate sample provided an opportunity to explore the consumer-specific attributes that govern automotive choice. This led to the development of a procedure to merge the aggregate and disaggregate data, and produced a dataset that generates estimates that are consistent. This consistency arises from the fact that each observation within the disaggregate data is effectively scaled by a weight that is equivalent to that established by Manski and Lerman (1977) for estimation of choice models with endogenous datasets. The application of this approach to generating datasets for choice model estimation has the potential to save resources in terms of data collection costs. More importantly, it permits an efficient evaluation of the large collective decision processes that would otherwise be extremely challenging to model. For example, individual observations of new vehicle purchases are inherently difficult to collect and often constitute proprietary data. Household surveys alone are inefficient because random sampling will inevitably cover populations outside those of traditional new car

buyers. These samples will have to filter or discard the many households that hold cars that were bought used. Both of these approaches require extensive resources allocated to data collection. For these reasons, modeling automotive choice from the ground up is an extremely challenging task with expensive data needs. This data needs to be regenerated frequently because the market contains a choice set that changes rapidly. The aggregate-disaggregate data combination simplifies the problem significantly. That is, disaggregate data collection is simplified to the task of generating an endogenous dataset, in which as many observations as possible are collected of a particular choice. As demonstrated in this study, the observations can be assembled from existing datasets that are collected for separate purposes as long as those datasets are random in automotive choice. Ideally though, a complete dataset representative of the country overall would be preferred. Unfortunately, many household surveys conducted by federal and state governments do not collect information on household vehicles. Still, this approach to data generation has a natural application to the automotive market, but it could find useful applications to other circumstances in which the market is large, diverse and met with similar data challenges.

While the aim of the choice model estimation was to conduct policy analysis, the research also explored the degree to which recent advances in GEV model structures could capture some of the complicated substitution patterns that are known to exist within the automotive market. Structures such as the CNL, GNL and various other network GEV models permit alternatives to straddle more than one nest, allowing alternatives within similar nests to be more closely associated than alternatives in more distant nests.

The drawback is the large number of additional parameters, which increases the time and complexity of the estimation process. Although this study does not rule the possibility that estimable models of automotive choice could be executed in the form of a network GEV, efforts here found that too many local optima were introduced to produce a reliable set of estimates. It is possible, perhaps likely, that a higher order GEV model of some specification does exist, and that this model is capable of reproducing substitution patterns within the market while maintaining theoretically consistent estimates. The search for better models within this data is certainly an area worthy of further research.

The policy analysis simulated what would have happened to vehicle sales had the hybrid tax credits not been applied. This analysis suggests that from 2006 to 2007, hybrid tax credits will prevent 1,500,000 t of GHG emissions over the lives of the policy-induced vehicles. On a yearly basis, this amounted to an average 770,000 t GHG per year. However, the impacts are uneven across years. The simulated impact is larger in 2006 than in 2007. The reason for this discrepancy is a result of the larger incentives that were applied in 2006. Because all probabilities predicted are less than 50%, increased incentives provide increasing returns in drawing demand away from other models. This is a property of the logistic curve, but it is also anchored in real human behavior. Large incentives tend to catch more attention and can create large shifts in behavior. The recent cash-for-clunkers policy may be a good example of this in action. This policy's minimum incentive of $3500 per clunker is larger than the tax incentive on the Prius at its peak. The clunker incentive could be as large as $4500 if a vehicle as efficient as the Honda Civic, Honda Fit, or Toyota Prius is bought. Thus far, the program has proven to

be incredibly popular, with preliminary evidence suggesting that the policy is causing

shifts in vehicle purchases from gas guzzling SUVs to smaller more efficient cars (Scott,

2009). While the shift in vehicle type is observable, there exist factors in the policy

analysis that always require assumptions and simulation of the variation in those

assumptions. For instance, it is unknown exactly how far the vehicles replaced by any

incentive policy are driven. In addition, vehicle life is another variable that varies across

the population of affected vehicles. These factors have a direct effect on the emissions

impacts induced by these policies. They will always be characterized by a distribution,

and there will always be some degree of approximation applied in their measurement.

These factors are relevant both for the vehicle purchased and the vehicle removed or

avoided as a result of the policy. The results suggest that the cost effectiveness of the

hybrid tax credits is within a range of $1000 to $3000 per metric ton of GHG avoided.

While new policy options are being developed all the time, it is likely that the hybrid tax

credits are somewhere in the middle range of policy cost-effectiveness. That is, it may

not be considered a "low hanging fruit" among GHG policies. Still, metric tons of GHG

avoided is but one measure criterion of policy effectiveness. The policy is rewarding

technological innovation within the automotive industry, and the initial success of

hybrids alongside their endorsement has led to laggards in the industry working diligently

to bring products to the market. It is likely such technology would not have been

introduced by the industry laggards independently. The hybrid tax credits of 2006 and

2007 are also effective in reducing gasoline consumption, which is a geo-political

objective. The analysis suggests that the hybrid tax credits can save an average of 2.5

billion gallons of gasoline (producing the 1.5 Mt GHG) from being consumed over a ten

year period. In addition, hybrid tax credits were found to be more effective in moving the

average fuel economy of the new vehicle fleet than a near doubling of the gasoline tax.

This is an important finding because the gasoline tax has long been touted as a desirable

yet politically infeasible policy objective. There are many good reasons to still consider

the gasoline tax for future analysis. But the political climate for the foreseeable future

does not bode well for the gasoline tax, even if the policy were found to be

overwhelmingly beneficial for the environment. Thus, taken together with the impacts on

GHG emissions, petroleum consumption, and a likely (though qualitative) spur in

automotive innovation, the hybrid tax credits can be considered a successful policy.

With the expansion of data sources, this area of research contains considerable

opportunities for future work. The data challenges in evaluating new vehicle choice will

be ever present, as choice sets will always be diverse and shifting. But one of the greatest

limitations of this study was the basic choice model applied to the data. Although more

advanced specifications were attempted in an effort to better model substitution patterns,

several barriers were encountered. These limitations are likely the result of a lack of

computing power, limitations of optimization algorithms available, and limits on the

number of model permutations that could conceivably be attempted given the long

computation time required for each attempted model run. There is certain room for

improvement in this respect, as many of the limiting factors are technical, and bound to

be overcome by improvements in computing speed and discrete choice estimation

software such as BIOGEME and NLOGIT. Further research opportunities also exist

within the realm of policy. The government has recently become more involved with the

automotive industry during the current recession. It is unclear at this time whether and to what degree government involvement at this level will persist over the long-term. But it is clear that government will retain a keen interest in influencing the products and standards of the industry. As the automotive technologies available to consumers diversify, the government will create new programs such as the very recent cash-for-clunkers, which was initiated with the idea that older cars are among those that are the least efficient and most polluting. The program was introduced rapidly, motivated primarily by the objective to stimulate new vehicle sales during an episode of crisis throughout much of the economy. But the program is controversial, and to evaluate its effectiveness, a similar set of methodological tools and data design are necessary to understand how people behave when faced with such a program. With this and other policies likely on the horizon, there is a considerable need for policy research in understanding which structures produce the most cost-effective set of benefits for all parties involved. It is likely that as automotive technology advances, the mix of technology and fuels will increase the diversity of options available to consumers. With each option carrying its own implications on foreign oil dependence and greenhouse gas emissions, the importance of understanding how effective incentives can be in influencing the composition of the automotive fleet is likely to grow. It may be that direct consumer incentives are not the most efficient decisions that the public sector can make in mitigating public costs. But without a mechanism to answer this question, it is likely that the question will never be asked at all. Thus, with improved estimation techniques of choice models and the increased availability and accessibility of data, there

is a promising future for the development of routine applications of choice models to

answering critical policy questions for the automotive sector and beyond.

References

Arguea, N. M.; Hsiao, C. T. A. (1994), 'Estimating Consumer Preferences Using Market Data - An Application to US Automobile Demand', *Journal of Applied Econometrics* **9**(1), 1-18.

Beggs, Steven; Cardell, S. (1980), 'Choice of Smallest Car by Multi-Vehicle Households and the Demand for Electric Vehicles', *Transportation Research A* **14A**, 389-404.

Ben-Akiva, M., Lerman, S. (1985), 'Discrete Choice Analysis' *The MIT Press*. Cambridge, MA.

Berkovec, James. (1985), 'Forecasting Automobile Demand using Disaggregate Choice Models', *Transportation Research B* **19B**(4), 315-329.

Berkovec, James; Rust, J. (1985), 'A Nested Logit Model of Automobile Holdings for one Vehicle Households', *Transportation Research B* **19B**(4), 275-285.

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

Berry, Steven, Levinsohn, J., and Pakes (2004), 'Differentiated Product Demand Systems from a Combination of Micro and Macro Data: The New Car Market', **112**(1), 68 – 105.

Bhat, Chandra; Sudeshna, S. (2006), 'Household vehicle type holdings and usage: an application of multiple discrete-continuous extreme value model', *Transportation Research B* **40**, 35-53.

Bierlaire, Michel. (2003), 'BIOGEME: A free package for the estimation of discrete choice models' Proceedings of the 3rd Swiss Transportation Research Conference.

Bierlaire, M. (2009), 'Estimation of discrete choice models with BIOGEME: 1.8', biogeme.epfl.ch.

Boyd, Hayden; Mellman, R. (1980), 'The Effect of Fuel Economy Standards on the US Automotive Market: An Hedonic Demand Analysis', *Transportation Research A* **14A**, 367-378.

Brownstone, D., Bunch, D., Train, K. Joint mixed logit models of stated and revealed preferences for alternative-fuel vehicles. Transportation Research Part B. 34. pp 315-338.

California Air Resources Board (2009), 'Key Events in the History of Air Quality in California', http://www.arb.ca.gov/html/brochure/history.htm.

California Assembly (2006) 'Assembly Bill 32: California Global Warming Solutions Act of 2006' California Assembly, Signed September 27, 2006.

Cardell, Scott; Dunbar, F. (1980), 'Measuring the Societal Impacts of Automobile Downsizing', *Transportation Research A* **14A**, 423-434.

Cowling, K., Cubbin, J. (1972) Hedonic Price Indexes for United Kingdom Cars. The Economic Journal, Vol. 82, No. 327, pp. 963-978.

Choo, Sangho; Mokhtarian, P. (2004), 'What type of vehicle do people drive? The role of attitude and lifestyle in influencing vehicle choice type', *Transportation Research A* **38**, 201-222.

Espey, M., and Nair, S., (2005) Automobile Fuel Economy: What is it worth? *Contemporary Economic Policy*. Vol 23, No. 3. pp. 317 – 323.

Greene, D. (1985), 'The market share of diesel cars in the USA, 1979 - 83', *Energy Economics* **8**(1), 13-21.

Greene, David; Patterson, P. S. M. L. J. (2005), 'Feebates, rebates and gas-guzzler taxes: a study of incentives for increased fuel economy', *Energy Policy* **33**, 757-775.

de Haan, Peter; Peters, A. M. M. (2006), 'Comparison of Buyers of Hybrid and Conventional Internal Combustion Engine Automobiles: Characteristics, Preferences, and Previously Owned Vehicles', *Transportation Research Record: No. 1983*, 106-113.

Heffner, R. (2007), 'Semiotics and Advanced Vehicles: What Hybrid Electric Vehicles Mean and Why it Matters to Consumers', PhD thesis, University of California, Davis.

Herriges, J; Kling, C. (1996), 'Testing the consistency of nested logit models with utility maximization', *Economic Letters* **50**, 33-39.

Hogarty, T. (1975), 'Price-quality relations for automobiles: a new approach', *Applied Economics* **7**(1), 41-51.

Hughes, Jonathan; Knittel, C., Sperling, D. (2008), 'Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand', *The Energy Journal* **29**(1).

Kavalec, C. (1996), 'CALCARS: The California Conventional and Alternative Fuel Reponse Simulator', Technical report, Califronia Energy Commission.

Kling, C.; Herriges, J. (1995), 'An empirical investigation of the consistency of nested logit models with utility maximization', *American Journal of Agricultural Economics* **77**, 875-884.

Krauss, Clifford, 'Driving Less, Americans finally react to the sting of high gas prices, a study says', *New York Times*, June 19, 2008.

Mannering, Fred; Mahmassani, H. (1985), 'Consumer Valuation of Foreign and Domestic Vehicle Attributes: Econometric Analysis and Implications for Auto Demand', *Transportation Research A* **19A**(3), 243-251.

Mannering, Fred; Train, K. (1985), 'Recent Directions in Automobile Demand Modeling', *Transportation Research B* **19B**(4), 265-274.

Mannering, Fred; Winston, C. S. W. (2002), 'An exploratory analysis of automobile leasing by US households', *Journal of Urban Economics* **52**, 154-176.

Manski, Charles; Lerman, S. (1977) 'The Estimation of Choice Probabilities From Choice-Based Samples,' *Econometrica,* **45**(8), pp.1977-1988. Reinstated 2013-12.

Manski, Charles; Sherman, L. (1980), 'An Empirical Analysis of Household Choice Among Motor Vehicles', *Transportation Research A* **14A**, 349-366.

McCarthy, P. (1996), 'Market Price and Income Elasticities of New Vehicle Demands', *The Review of Economics and Statistics* **78**(3), 543-547.

McFadden, D. (1974) 'Conditional Logit Analysis of Qualitative Choice Behavior' in P. Zarembka (ed.), *FRONTIERS IN ECONOMETRICS*, 105-142, Academic Press: New York, 1974.

McFadden, D; Train, K. (2000), 'Mixed MNL models for discrete response', *Journal of Applied Econometrics* **15**(5), 447-470.

McManus, W. (2007) The Link Between Gasoline Prices and Vehicle Sales: Economic Theory Trumps Conventional Detroit Wisdom. *Business Economics*. January 2007. 53-60.

Oracle. (2008) Oracle Crystal Ball Enterprise Performance Management. Oracle Data Sheet. Oracle, Redwood Shores, CA.

Page, J. Malachi Weng-Guitierrez, Gordon Schremp. Overview of the proposed energy analyses for the 2007 integrated energy policy report. Staff Report California Energy Commission. CEC-600-2007-005, 2007.

Revelt, David; Train, K. (1998), 'Mixed Logit with Repeated Choices: Households' Choice of Appliance Efficiency Level', *The Review of Economics and Statistics* **80**(4), 647-657.

Scott, C. (2009) "Clunkers = SUVs" *San Francisco Chronicle*. San Francisco, CA. http://www.sfgate.com/cgi-bin/blogs/green/author?blogid=49&auth=264. August 6[th], 2009.

Small, K. (1994), 'Approximate generalized extreme value models of discrete choice', *Journal of Econometrics* **62**, 351-382.

Small, K. (1987), 'A discrete choice model for ordered alternatives', *Econometrica* **55**, 409-424.

Small, K; Van Dender, K. (2007), 'Fuel Efficiency and Motor Vehicle Travel: The Declining Rebound Effect', *Energy Journal* **28**, 25 - 51.

Train, K. (2003) Discrete Choice Methods with Simulation. Cambridge University Press. Cambridge, UK.

Train, Kenneth; Davis, W. L. M. (1997), 'Fees and Rebates on New Vehicles: Impacts on Fuel Efficiency, Carbon Dioxide Emissions, and Consumer Surplus', *Transportation Research Part E* **33**, 1 - 13.

Train, Kenneth; Hudson, K. (2000), 'The Impact of Information on Vehicle Choices and the Demand for Electric Vehicles in California', Technical report, National Economic Research Associates, 444 Market Street, Suite 910, San Francisco, CA 94111.

Train, Kenneth; Lohrer, M. (1983), 'Vehicle Ownership and Usage: An Integrated System of Disaggregate Demand Models'' Transportation Research Board Annual Meeting', Note: This constitutes Chapter 8 of Train's 1986 book.

Train, Kenneth; Winston, C. (2007), 'Vehicle Choice Behavior and the Declining Market Share of U.S. Automakers', *Internation Economic Review* **48**(4), 1469-1496.

Turrentine, Thomas; Kurani, K. (2007), 'Car buyers and fuel economy?', *Energy Policy* **35**, 1213-1223.

United States Census Bureau. (2009) Current Population Survey, American Community Survey. United States Census Bureau. http://www.census.gov/cps/, http://www.census.gov/acs/www/SBasics/. Accessed August, 2009.

Vovsha, P. (1997), 'The cross-nested logit model: Application to mode choice in the Tel Aviv metropolitan region' 76th Transportation Research Board Conference.

Appendix A: The American Automotive Market

As of 2007, there were nearly 400 unique vehicle models that could be purchased by a new car buyer in the US automotive market. As a choice set, this includes vehicles with a purchase price range of $10,000 to over $100,000 and vehicle types that range from 2-seater sports cars to full size SUVs and pickup trucks. The car market serves a variety of lifestyles as each type of car buyer has different fundamental needs and different preferences on what attributes should be emphasized within their vehicle. Consumers also enter the market with different incomes, facing different constraints in evaluating attribute trade-offs within their vehicle choice set.

While the automotive market may appear to have a large number of choices, modeling the consumer decision becomes easier if a simple, yet realistic conceptual structure is applied to the decision process. As an example, assume that every new car buyer asks two fundamental questions of him or herself before engaging the market. They are 1) "What kind of car do I want?" and 2) "What car of that kind can I afford?" For those with middle-class income constraints, who can afford at least one vehicle within each market class, the questions would be likely asked in this order. The first question is shaped by lifestyle needs or desires and the second by personal budgets.

But after self-evaluating these two questions for any consumer, the choice set narrows down considerably from the entire market. It would be unusual to find a consumer who is seriously considering either the purchase of a two-seater sports car or an SUV as their final decision given that the two vehicles serve very different needs. As vehicle types become more alike, some choices sets that straddle similar vehicle classes

may arise. A 2-door sports coupe may on occasion compete with a family sedan, and a large family sedan may compete with a wagon or small SUV. However it is likely that these "cross-class" choice sets developed by certain individuals are not the norm.

Automakers may take a similar view, and this is suggested by casually examining the pricing structure of the sedan market. Sedan buyers exhibit a standard set of needs and do not have special preferences for off-roading, towing, hauling large items, or driving excessively fast. Their vehicle may exhibit some properties in favor of one or two of these needs, but if any of these needs were a serious lifestyle consideration, then the buyer would look towards another type of vehicle. In addition, the demand for non-luxury sedans (< $ 30,000) is less subject to idiosyncratic preferences for status bearing components. Non-luxury sedans in this respect are relatively simple to characterize, constitute roughly 20% of the total automotive market, and offer a good starting point in terms of evaluating consumer preferences for specific vehicle attributes in more detail. Extending analysis into other markets is more plausible once a model for this simple market is established and understood.

A closer examination of pricing structures within the sedan market suggests additional substructure that if adequately captured by attribute variables, may be effective in controlling for the true cause of increased premiums and the nature of substitution patterns. This structure is evident by the existence of well-defined price echelons within sedans on the market under $30,000. There is a difference between the sedan buyer that looks for a car with a maximum budget of $15,000 versus a buyer with the same lifestyle needs but a budget of $20,000 or $30,000. To illustrate this point, consider the following

set of graphs in Figure 1A, which show the range of Manufacturer Suggested Retail

Prices (MSRPs) of the 2007 vehicle models across four major competitors within the

sedan market. Note that the combined fuel economy of each vehicle listed below the

model name, and the passenger volume (in ft^3) is listed across the top.

Figure 1A: MSRP Ranges of Selected Sedan Lineups
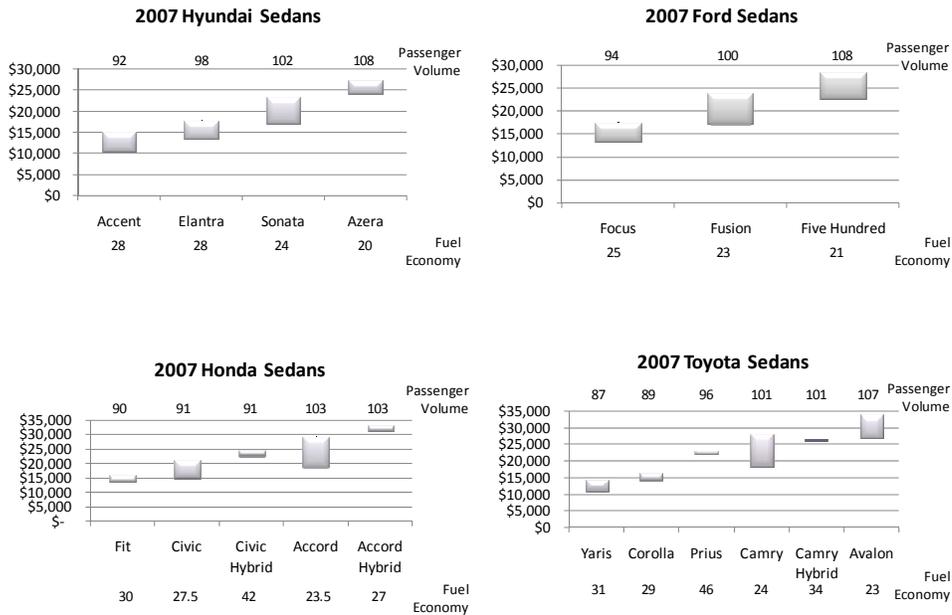
This set of graphs is useful for conceptually framing this segment of the

automotive market and the general choice structure that may exist within the broader

sedan market. The graphs clearly show segregated price tiers in vehicle models within

the selection offered by each automotive manufacturer. Conceptually, it suggests that

automakers are less inclined to consider vehicles within their sedan line-up to be

differentiable in terms of unique attribute packages that serve different purposes. Rather,

their vehicles are distinguishable by price, size and amenity, so as to not compete with

one another. The competition occurs across automakers by vehicles in the same price tier.

The combined perspective of sedan MSRP and fuel economy illustrates an important dynamic with respect to consumer choice and the value placed on fuel economy. Prior to the advent of hybrids, emphasizing fuel economy as a priority attribute of the sedan purchase meant *paying less* for a vehicle. Other attributes such as improved power and increased size have long commanded direct premiums. Because variation in fuel economy has been introduced in a fashion that ties premiums to more efficient, but otherwise similar vehicles, there exists a unique opportunity today that has not existed before to evaluate the consumer willingness to pay for fuel economy from new vehicle choice data over recent years.