

**IMPROVING FUEL ECONOMY THROUGH SUBSIDIES: EVIDENCE FROM “CASH
FOR CLUNKERS”**

by

Daniel Eubanks

Advisor: Mario Crucini

Economics Honors Thesis

April 2013

DEPARTMENT OF ECONOMICS

VANDERBILT UNIVERSITY

NASHVILLE, TN 37235

www.vanderbilt.edu/econ

Improving Fuel Economy through Subsidies: Evidence from “Cash for Clunkers”¹

Daniel Eubanks

Vanderbilt University

Advisor: Professor Mario Crucini

April 15, 2013

Abstract

Policymakers have a clear interest in encouraging American automobile consumers to purchase more fuel efficient vehicles. The fuel efficiency of the automobile stock in the United States has implications for the environment through vehicle emissions and for national security through dependence on foreign energy supplies. We focus on the 2009 “Cash for Clunkers” program, which sought to incentivize the purchase of fuel efficient vehicles through a subsidy that focused on the difference in fuel economy between the trade-in vehicle and the new vehicle. Our analysis of the program indicates that consumers place greater weight on the purchase price of a vehicle than the operating cost; therefore, a subsidy will be more effective than a fuel tax in influencing consumers to purchase fuel efficient vehicles. In the absence of the “Cash for Clunkers” program, purchases of the most inefficient vehicles, defined as vehicles with a fuel economy of less than 20 miles per gallon, would have increased by nearly 15%. But while the

¹ I would like to extend a special thanks to Gregory Huffman, and to my advisor, Mario Crucini, without whose invaluable guidance, feedback, and support I could not have completed this research.

subsidy shifted consumers towards more fuel efficient vehicles, changes in the design of the program could have led to even greater fuel efficiency gains.

1 Introduction

The automobile an individual chooses to purchase is the result of a complex decision making process. The individual must consider not only the physical features of each car, such as size, horsepower, and style, but also the full cost of each vehicle option. The cost of a car extends beyond the purchase price; the new owner must also pay to operate the vehicle. The operating cost is itself complex, the result of how often the individual drives, the fuel economy of the vehicle, and the price of gas. Furthermore, the characteristics of each individual consumer will shape that consumer's decision.

The policy levers available to influence consumer choice seem crude in light of the complexities of the car purchasing decision. Most government policy has aimed to steer consumers to fuel efficient vehicles through decreasing the costs associated with purchasing efficient vehicles. One such policy, the 2009 "Cash for Clunkers" program, provided a subsidy to encourage consumers to trade-in older, inefficient vehicles and purchase new, more fuel efficient vehicles. Although the timing of the program suggests that economic stimulus and support to ailing Detroit automakers were the primary goals, the program was also sold as a step towards a cleaner environment and reduced energy dependence on foreign nations. This paper focuses on the environmental aspects of the program. Our goal is to determine how the "Cash for Clunkers" subsidy affected consumer choice of automobile with respect to fuel efficiency and how the outcome of the program would have changed under alternative subsidy schemes.

Section 1.1 provides background information on “Cash for Clunkers,” including legislative history and eligibility criteria. Section 1.2 presents an overview of the existing literature on discrete choice modeling with an emphasis on studies that examine automobile choice. Section 2 discusses our data and provides summary statistics. Section 3 models consumers’ utility functions, provides an overview of discrete choice methodology and details the construction of our variables. Section 4 presents empirical results of the model estimation and discusses counterfactuals and policy implications for “Cash for Clunkers.” Finally, Section 5 provides a concluding summary, caveats, and suggestions for further research.

1.1 Background

President Obama signed the Consumer Assistance to Recycle and Save (CARS) Act into law on June 24th, 2009, during the trough of the Great Recession and less than a year after the automobile industry bailouts of late 2008 and early 2009. The scrappage program, which the press dubbed “Cash for Clunkers,” had dual goals: to provide economic stimulus and to improve the fuel economy of America’s automobile stock. Under the program, buyers of qualifying new vehicles would receive a rebate for trading in an older, less fuel efficient “clunker.” To prevent abuse, the National Highway Traffic Safety Administration, the organization tasked with administering the program, set limits on which vehicles would qualify as clunkers. To be eligible, the NHTSA required the trade-in to be less than twenty-five years old, have a combined fuel economy of eighteen miles per gallon or less, be continuously registered and insured to the same owner for a full year prior to the trade in, and be in drivable condition.

To encourage consumers to purchase the most fuel-efficient vehicles, the amount of the rebate increased with the difference in fuel efficiency between the trade-in and the new vehicle.

Consumers who traded in a passenger car received a rebate of \$3,500 if the new vehicle was at least 4 miles per gallon more fuel efficient than the old vehicle and \$4,500 if the new vehicle was at least 10 miles per gallon more fuel efficient. The fuel efficiency requirements were relaxed for consumers who traded in a light truck: the \$3,500 rebate required a 2 miles per gallon gain while the \$4,500 rebate required only a 4 miles per gallon gain. To avoid the possibility of the “clunkers” making their way back onto the used car market, a controversial condition of the CARS program required trade-in vehicles to be crushed or shredded within six months of the transaction. Salvage facilities were permitted to sell select parts of the trade-in vehicle, but those parts could not include the engine or drive train. “Cash for Clunkers” began July 24th and ended only a month later, far in advance of the planned November 1st end date, after the exhaustion of the program’s \$3 billion allocation.

1.2 Related Literature

The analysis in this paper modifies and combines several of the approaches taken by previous researchers to examine automobile choice. Particularly, we adapt the methodology of Lave and Train (1979) to create hypothetical representative vehicles in the choice set and employ average socioeconomic data to represent consumer heterogeneity as proposed by Petrin (2002).

The application of the discrete choice framework to the automobile market has a long history. Over the past several decades, researchers have focused on two different but related applications of the model: aggregate discrete choice models and disaggregate discrete choice models. Both employ an identical theoretical framework; the distinction between the two comes from the level of data used to estimate the model.

Aggregate discrete choice models rely on market share data to estimate the aggregate demand for different makes and models of automobile. These models explore how prices and attributes of different vehicles relate to market shares and then use these relationships to estimate the weights of different attributes in a representative utility function.

In an influential paper, Berry, et al. (1995) develop a comprehensive discrete choice model of aggregate automobile demand. This paper made important theoretical contributions to discrete choice modeling and demonstrated the potential scope of such models. Berry, et al. address many issues that arise in these types of models, including the inability to observe or quantify many of the product-specific characteristics that determine an individual's choice, such as brand reputation, brand loyalty, and style attributes. To account for these unobserved characteristics, Berry includes in the representative consumer's utility function a constant term to capture the average utility derived from un-measurable vehicle attributes. Berry, et al. find that consumers of small, fuel-efficient cars have highly elastic demand with respect to the fuel economy of competing vehicle models. Particularly relevant to the present paper, their results also indicate that consumers of larger vehicles lose utility with increasing fuel efficiency. In theory, if all other attributes are held constant, all consumers should gain utility from increased fuel efficiency through the reduction in vehicle operating costs. The results of Berry, et al. illustrate the difficulty in specifying a model that captures the often unobservable attributes for which consumers sacrifice fuel efficiency, such as size, luxury, power, and brand loyalty. As other researchers have noted (Allcott and Wozny, 2012), the negative correlation between these attributes and fuel efficiency make it difficult to disentangle the competing impacts on utility and achieve coefficients with the anticipated sign.

Other studies of discrete choice problems using aggregate data have sought to address issues with the multinomial logit model. The multinomial logit (MNL), the basis for the model used in this paper (alternative-specific conditional logit) and in many others to model discrete choice problems, imposes unrealistic restrictions on substitution patterns. In the MNL, cross-price elasticities depend only on the average level of utility provided by each vehicle and not the characteristics of the vehicle. As a result, any two vehicles that have the same market share will have the same cross-price elasticity with a given third vehicle. This property is known as IIA: independence of irrelevant alternatives. When the price of a vehicle increases, consumers are likely to substitute towards a different vehicle with similar characteristics. The MNL model, however, does not account for this.

Boyd and Mellman (1980) address the unrealistic substitution patterns in the MNL model by assuming that preferences vary among consumers, such that the coefficients in the utility function follow a random distribution. This is known as the mixed logit model. This modification allows for more realistic substitution patterns: a consumer who purchased a fuel efficient vehicle will be modeled as more likely to substitute to another fuel efficient make or model given a price increase. Their utility function includes price, fuel economy, repair frequency, and several other vehicle attributes. The researchers find that a doubling of gasoline prices at the time of their study would lead to a 6% increase in average fuel economy of new vehicles.

Cardell and Dunbar (1980) also employ the mixed logit to model automobile demand. Their application of the aggregate discrete choice framework focuses on the welfare implications of Corporate Average Fuel Economy (CAFE) standards as compared to changes in fuel prices. They find that policy aimed at increasing fuel economy by increasing fuel prices would have a lower social cost than CAFE reductions that achieved the same improvement in fuel economy.

Petrin (2002) further improves aggregate discrete choice methodology by including socioeconomic factors in the mixed logit estimation. Using market level data and average socioeconomic characteristics of consumers of different products taken from the Consumer Expenditure Survey, Petrin's model allows for more realistic substitution patterns without requiring individual-level data for each purchase.

The research cited above focused aggregate data; alternatively, disaggregate discrete choice models employ individual or household characteristics and purchase decision data to estimate the demand for new vehicles for a given individual. These models relate individual or household level decisions to vehicle prices and attributes, and then use this information to estimate the coefficients in a representative utility function.

Disaggregate discrete choice models arose in the 1970s with the development of the discrete choice framework (McFadden, 1974). Lave and Train (1979) applied the discrete choice methodology to examine household vehicle choice, given that a household has already made the decision to purchase a vehicle. The researchers create a choice set of ten fictitious representative vehicles by averaging vehicle attributes within a market class, and then use the MNL model to estimate the probabilities that a household will purchase a vehicle in one of the ten classes. The "representative vehicle" approach is adapted in this paper.

Later researchers focused on capturing the heterogeneity of consumer preferences. Berkovec and Rust (1984) developed a sequential choice framework in which a household first chooses the class of vehicle and then chooses the make and model. While this approach does allow the coefficients in the utility function to vary depending on the class of vehicle selected, the inflexibility of the choice structure is a limiting factor. Mannering and Mahmassani (1985)

account of consumer heterogeneity by estimating one set of coefficients for consumers who purchased domestic vehicles and another for consumers who purchased foreign vehicles. Both studies revealed the importance of accounting for heterogeneity in consumer preferences when modeling vehicle purchase decisions.

Several studies rely on consumer survey data to develop a nested logit choice model of individual automobile demand. In a nested logit model, households choose several characteristics simultaneously, but the model is organized in a hierarchical way and utilizes conditional probabilities. Goldberg (1995) uses data from a survey of American consumers between 1983 and 1987 to model a five-stage decision process: households decide to purchase a vehicle (or not), new or used, vehicle class, domestic or foreign, and the model of vehicle. Gold includes household characteristics only in the final stage. McCarthy and Tay (1998) use data from a 1989 consumer survey to specify a model where the “nests” of the nested logit model are ranges of fuel economy. Their results show the importance of fuel economy class in the determination of automobile demand. Mohammadian and Miller (2003) use the results of a Canadian survey containing vehicle transaction data over a nine year span to build a nested logit model that includes used vehicles. In their model, the household chooses the vehicle class and then vehicle age.

The preceding section provides context for the present paper and precedent for our methodology. With these studies as a foundation, the goal of this paper is to model the vehicle choice of consumers who have chosen to participate in the “Cash for Clunkers” program using the characteristics of the previous vehicle as a proxy for consumer heterogeneity.

2 Data

2.1 Sources

Our primary data was obtained from the National Highway Traffic Safety Administration's (NHTSA) database of CARS transactions. For each transaction, the dataset includes location information and details of both the trade-in vehicle and the purchased vehicle. The location information includes the city, state, and zip-code of the dealership where the transaction occurred.

The dataset contains several relevant details about the trade-in vehicle, including the vehicle category (Passenger Vehicle or Category 1, 2, or 3 Truck), make, model, year, drive train, fuel economy, and odometer reading. Similarly, for the purchased vehicle, the dataset contains vehicle category, make, model, drive train, fuel economy, and Manufacturer's Suggested Retail Price (MSRP). Summary statistics are provided in Section 2.2.

Supplementary income data comes from the American Community Survey (Table S1903). The ACS data, obtained through the US Census Bureau, contains five-year estimates (2007-2011) of median household income by zip code. To generate estimates of consumer income in "Cash for Clunkers," we match the zip code in the CARS transaction database to the income data in the ACS. However, the consumers in CARS likely did not live in the zip code in which they purchased their new vehicle. If the median income in the zip code of the dealership differs greatly from the median income in the consumer's home zip code, these could be inaccurate estimates.

Finally, in estimating operating costs for vehicles, no attempt was made to predict or model gasoline prices. Instead, we assume the price of gasoline is constant at \$3.75 per gallon.

This figure comes from the Environmental Protection Agency's fueleconomy.gov website, which uses \$3.75 per gallon to estimate cost savings from improving fuel efficiency.

2.2 Summary Statistics

Table 1 below presents summary statistics for the “Cash for Clunkers” program. Note that the fuel economy classes that define our representative vehicles are broken into ranges of fuel efficiency. The low category is defined as vehicles that get less than 20 miles to the gallon. Medium low vehicles fall in the range 20 to 25; medium, 25 to 30; medium high, 30 to 35. High fuel efficiency vehicles can travel more than 35 miles on one gallon of gas.

The NHTSA data places each new vehicle purchased into one of four classes: passenger automobile (P), category one truck (1), category two truck (2), or category three truck (3). The truck categories correspond to weight classes and include pickup trucks, sport utility vehicles, and vans. A category one truck has a gross vehicle weight rating (GVWR) from 0 to 6,000 pounds. This category covers lighter pickups, such as the Toyota Tacoma and Dodge Dakota and smaller SUVs, such as the Toyota RAV4. Category two trucks have a GVWR from 6,001 to 10,000 pounds. This category includes heavier pickups, such as the Ford F-150 and Dodge Ram 1500, and larger SUVs, like the Chevrolet Suburban. Finally, category three trucks were the largest vehicles sold under CARS. These vehicles have a GVWR 10,001 to 14,000 pounds. Included in this category are large pickups, such as the Ford F-350 and GMC Sierra 3500. Large SUVs, such as a Hummer H1, would also fall under this category. The remaining category, passenger automobiles, covers all other vehicles sold under CARS: coupes, sedans, luxury cars, and so on.

Table 1: CARS Summary Statistics

	Trade-in Vehicles		Purchased Vehicles	
Top 5 Models	Ford Explorer (4WD)		Toyota Corolla	
	Ford F150		Honda Civic	
	Jeep Grand Cherokee		Toyota Camry	
	Ford Explorer (2WD)		Ford Focus	
	Dodge Caravan		Hyundai Elantra	
Fuel Economy Class	Low	579,023	Low	81,081
	Medium Low	96	Medium Low	202,883
	Medium	1	Medium	242,242
	Medium High	0	Medium High	33,801
	High	0	High	19,113
Vehicle Category	P	102,638	P	397,182
	1	447,505	1	230,220
	2	119,394	2	47,425
	3	7,544	3	2,254
Average Fuel Efficiency (MPG)	15.81		24.97	
Average Age (years)	13.78		--	
Average Odometer Reading (miles)	159,950		--	
Average MSRP (\$)	--		\$22,403.15	

2.3 Data Cleaning

The accuracy of the data in the NHTSA database relies on the ability and willingness of thousands of employees at thousands of dealerships across the country to correctly enter information about hundreds of thousands of transactions. Not surprisingly, many features of the data suggest errors in data entry. In the interest of reproducibility, this section briefly documents the criteria we used to reject observations.

The eligibility criteria for CARS required trade-in vehicles to be older than one year and younger than twenty five years. Thus, we rejected observations that indicated a trade-in vehicle age outside of this range. The database indicates an odometer reading of above one million for

many observations; in hundreds of cases, the odometer reading was listed as 9,999,999 or 8,888,888. To ensure realistic data, we rejected observations with odometer readings greater than 500,000 miles or less than 1,000 miles. Many vehicle price entries also suggest inaccuracies. We reject observations where the MSRP of the new vehicle is given as below \$7,500. Entries of zero for any of the numerical fields led to the rejection of that observation. Finally, if the dealership zip code could not be matched to Census income data, we did not include that observation in our model estimation.

Of the 677,081 transactions listed in the CARS database, the criteria outlined above led us to reject 103,443 observations for a final dataset of 573,638 transactions.

3 Model

3.1 Discrete Choice Framework

To model the behavior of consumers in the “Cash for Clunkers” program, we will employ the discrete choice framework. In a discrete choice model, decision-makers choose from among a finite set of mutually exclusive alternatives. The decision-maker is assumed to choose the single alternative that provides the highest level of utility. Because utility is not directly observable, several challenges arise when implementing the theoretical model.

In our model of the CARS program, consumers with different needs and characteristics face an array of vehicles with different attributes. Each vehicle, with its unique combination of attributes, provides a certain level of utility to the consumer. We denote the utility that consumer i obtains from vehicle j as U_{ij} , $j = 1, 2, \dots, J$, where J is the total number of vehicle options. The consumer purchases the vehicle that provides the greatest level of utility. Thus, consumer i will choose vehicle j if and only if $U_{ij} > U_{ik}$ for all $j \neq k$.

We cannot observe consumer i 's utility, but we can observe many of the key attributes of the vehicle options and several of the individual characteristics of consumer i . Denote the vector of attributes for vehicle j as \mathbf{x}_j and the vector of characteristics of consumer i as \mathbf{s}_i . Then we can specify a function that maps the attributes of vehicle j and the characteristics of consumer i to a “systematic” level of utility, denoted $V_{ij} = V(\mathbf{x}_j, \mathbf{s}_i)$. There are unobservable characteristics and attributes that also contribute to utility, so we write $U_{ij} = V_{ij} + \epsilon_{ij}$, where the error term ϵ_{ij} captures all elements of utility that are not included in V_{ij} . The terms $\epsilon_{ij}, j = 1, 2, \dots, J$ are modeled as random, with the joint density of the vector $\boldsymbol{\epsilon}_i = [\epsilon_{i1}, \epsilon_{i2}, \dots, \epsilon_{iJ}]$ denoted $f(\boldsymbol{\epsilon}_i)$. Using this joint density, we can make probabilistic statements about the decision of consumer i . Let P_{ij} be the probability that consumer i purchases vehicle j . Then,

$$\begin{aligned} P_{ij} &= \text{Prob}\{U_{ij} > U_{ik} \text{ for all } j \neq k\} \\ &= \text{Prob}\{V_{ij} + \epsilon_{ij} > V_{ik} + \epsilon_{ik} \text{ for all } j \neq k\} \\ &= \text{Prob}\{\epsilon_{ik} - \epsilon_{ij} < V_{ij} - V_{ik} \text{ for all } j \neq k\} \end{aligned}$$

Using the joint probability density of the error terms $f(\boldsymbol{\epsilon}_i)$, we can rewrite the above expression as a cumulative probability. Let $I(\cdot)$ be a function that takes the value 1 when then the argument of the function is true and 0 otherwise. Then we have:

$$P_{ij} = \int_{\boldsymbol{\epsilon}} I(\epsilon_{ij} - \epsilon_{ik} < V_{ik} - V_{ij}) f(\boldsymbol{\epsilon}_i) d\boldsymbol{\epsilon}_i$$

Different specifications of $f(\boldsymbol{\epsilon}_i)$ result in different discrete choice models. The assumption that each ϵ_{ij} is independently, identically distributed (iid) according to the extreme

value distribution (Type I) results in a closed form solution to the integral. Returning to our first expression for P_{ij} , we have:

$$\begin{aligned}
P_{ij} &= \text{Prob}\{U_{ij} > U_{ik} \text{ for all } j \neq k\} \\
&= \text{Prob}\{V_{ij} + \epsilon_{ij} > V_{ik} + \epsilon_{ik} \text{ for all } j \neq k\} \\
&= \text{Prob}\{\epsilon_{ik} < V_{ij} - V_{ik} + \epsilon_{ij} \text{ for all } j \neq k\}
\end{aligned}$$

If we take ϵ_{ij} as given, the above equation is the cumulative density function for each ϵ_{ik} evaluated at $V_{ij} - V_{ik} + \epsilon_{ij}$. Because of the assumption that the ϵ_{ij} are iid, the cumulative distribution over all $k \neq j$ is just the product of each individual cumulative distribution function. Then, by the total probability theorem, the probability that individual i chooses vehicle j is the integral of $P_{ij}|\epsilon_{ij}$ over all values of ϵ_{ij} and weighted by the probability density of ϵ_{ij} . Due to the functional form of the extreme value distribution, this integral has the closed form solution:

$$P_{ij} = \frac{e^{V_{ij}}}{\sum_{k \neq j} e^{V_{ik}}}$$

By inspection, this solution meets the criteria for a probability distribution: the P_{ij} are nonnegative and sum to one.

We improve the model by including alternative specific constants. The constant terms will capture the average impact on utility of all vehicle characteristics not included in the model. By construction, ϵ_{ij} will have mean zero when alternative specific constants are included. The inclusion of constants ensures that the average probabilities equal the observed shares of each vehicle in the data. Given how we have defined choice probabilities,

$$\begin{aligned}
P_{ij} &= \text{Prob}\{U_{ij} > U_{ik} \text{ for all } j \neq k\} \\
&= \text{Prob}\{U_{ij} - U_{ik} > 0 \text{ for all } j \neq k\},
\end{aligned}$$

only differences in utility are relevant to the decision. Therefore, the magnitudes of the alternative specific constants are not important; only the differences in the constants matter. Two models with different constants but the same *difference* in constants are equivalent and result in identical choice probabilities. Thus, there are an infinite number of constants that could be used in any given model. To address this issue, we choose one constant to normalize to zero. In a model with J alternative vehicle choices, then, we will estimate $J - 1$ constants. All other constants are then interpreted relative to the constant that we normalized to 0.

A similar issue arises with the inclusion of individual-specific characteristics, such as previous ownership of a domestic vehicle. Individual-specific characteristics do not vary over alternatives. These characteristics create differences in utility over the choice set, and we would expect the effect on utility of a given consumer characteristic to differ for each alternative that the consumer faces. For example, we would expect an increase in income to increase a consumer's probability of buying certain vehicle types and decrease the probability of buying other vehicle types. But again, only differences in utility matter; the absolute levels of the coefficients on individual-specific variables are meaningless. Indeed, because there are an infinite number of coefficients that will result in the same differences, we cannot estimate absolute levels of coefficients for each alternative. As above, we normalize the coefficient for one of the alternatives to zero. For J alternatives, we will estimate $J - 1$ coefficients. We then interpret the coefficients as the impact of the consumer characteristic on the utility of each alternative relative to the alternative we normalized to zero.

3.2 Choice Set

To populate our choice set, we create five representative vehicles based on fuel economy class. This representative vehicle approach is adapted from Lave and Train (1979). The choice set includes five hypothetical vehicles defined by their fuel economy: low, medium low, medium, medium high, and high. The fuel efficiency ranges for low, medium low, medium, medium high, and high are less than 20 mpg, 20-25 mpg, 25-30 mpg, 30-35 mpg, and greater than 35 mpg, respectively. The attributes of each representative vehicle are determined by averaging the attributes of all vehicles that fall into that class. Although the representative vehicles are determined by fuel economy class, consumers are still choosing a bundle of attributes when they select a vehicle and not just the fuel efficiency. Thus it remains a discrete choice problem. Table 2 contains descriptive statistics for the representative vehicles. Note that different configurations of the same model may fall into different fuel economy classes.

The representative vehicle approach has advantages and disadvantages. Vehicles with outlier attributes may disproportionately shift the average value of an attribute for a fuel economy class. By averaging vehicle characteristics, we lose information about vehicle attributes that lead consumers to purchase one vehicle of a given class over another vehicle of that same class. However, a choice set of five vehicles is computationally simpler and easier to interpret. In addition, the CARS program was designed to encourage the purchase of fuel efficient vehicles; creating representative vehicles based on fuel economy allows us to examine this aspect directly.

Table 2: Choice Set Summary Characteristics

Vehicle Class	Average MSRP (\$)	Average MPG	Average Engine Volume (L)	Selected Models
High	\$24,024.59	46.49	1.75	Ford Fusion Hybrid, Honda Insight, Honda Civic Hybrid, Toyota Prius
Medium High	\$18,186.66	31.02	1.75	Chevy Cobalt, Honda Fit, Kia Rio, Mini Cooper, Toyota Corolla, VW Jetta
Medium	\$19,054.97	27.24	2.07	Acura TSX, Chevy Malibu, Ford Focus, Honda Civic, Hyundai Elantra, Mazda 3, Toyota Corolla
Medium Low	\$24,095.21	22.28	2.67	Dodge Caliber, Ford Escape, Honda CR-V, Honda Accord, Hyundai Santa Fe, Jeep Patriot, Toyota Highlander, Volvo C30, Subaru Forester
Low	\$29,514.12	17.35	4.19	Chevy Silverado, Chrysler Town and Country, Dodge Ram 1500, Ford F150, GMC Sierra, Lincoln MKX, Mercedes-Benz GLK350, Subaru Tribeca, Toyota Tacoma

3.3 Variables

Our specification uses two alternative-specific variables and three individual-specific variables. The alternative-specific variables are initial cost and operating cost, and the individual specific variables are income, previous ownership of a domestic vehicle, and previous ownership of a Category 1, 2, or 3 truck.

3.3.1 Alternative-Specific Variable: Initial Cost

In the context of “Cash for Clunkers,” the effects of initial cost on consumer behavior are of vital importance. The program sought to incentivize the purchase of fuel efficient vehicles through subsidies that reduce the initial cost of the vehicle to the consumer. The sign and magnitude of the coefficient on this variable will indicate the effectiveness of subsidies in altering consumer behavior. As the initial cost of an automobile increases, consumers will have

to sacrifice additional consumption of alternative goods. Holding all other vehicle attributes constant, we would expect an increase in the initial cost of a vehicle to decrease the utility a consumer gains from its purchase through an increase in the opportunity cost. This decrease in utility should decrease the probability that any given consumer purchases that vehicle. Therefore, we would expect a negative coefficient on initial cost.

In our model, the initial cost of a new vehicle is the Manufacturer’s Suggested Retail Price (MSRP) minus the “Cash for Clunkers” subsidy. The size of the subsidy is a function of the increase in fuel economy between the trade-in vehicle and the new vehicle. For consumers who traded in a passenger car, the subsidy was \$3,500 if the new vehicle was at least 4 miles per gallon more fuel efficient than the old vehicle and \$4,500 if the new vehicle was at least 10 miles per gallon more fuel efficient. For consumers who traded in a light truck, the \$3,500 rebate required a 2 miles per gallon gain while the \$4,500 rebate required only a 4 miles per gallon gain. Specifically,

$$initial_cost_{ij} = MSRP_j - subsidy_{ij}$$

where, if the consumer is trading in a passenger car,

$$\begin{aligned} subsidy_{ij} &= \$4500 \text{ if } (mpg_i - mpg_j) \geq 10 \\ &= \$3500 \text{ if } 4 \leq (mpg_i - mpg_j) \leq 10 \\ &= 0 \text{ otherwise} \end{aligned}$$

and if the consumer is trading in a light truck (category 1, 2, or 3),

$$\begin{aligned} subsidy_{ij} &= \$4500 \text{ if } (mpg_i - mpg_j) \geq 4 \\ &= \$3500 \text{ if } 2 \leq (mpg_i - mpg_j) \leq 4 \\ &= 0 \text{ otherwise} \end{aligned}$$

The variables mpg_i and mpg_j are the miles-per-gallon of the trade-in vehicle and the potential new vehicle, respectively. The size of the subsidy will vary for each consumer and each vehicle option, so *subsidy* receives an *ij* subscript. For the same reason, *initial_cost* also takes an *i* and a *j* subscript.

To illustrate, if a CARS participant traded in a Ford F150 light truck with a fuel economy of 15 miles per gallon and purchased a new Toyota Corolla sedan with an MSRP of \$24,500 and a fuel economy of 30 miles per gallon, we estimate the initial cost to that participant as \$20,000: the MSRP of the Corolla minus the subsidy of \$4,500.

3.3.2 Alternative-Specific Variable: Operating Cost

The operating cost of a vehicle is also highly relevant to our analysis of “Cash for Clunkers.” The magnitude of the coefficient on operating cost will have implications for the relative effectiveness of a subsidy scheme versus an increased fuel tax or other policy lever designed to encourage the purchase of fuel efficient vehicles through higher operating costs. As the operating cost of a vehicle increases, consumers must give up additional consumption of alternative goods. A high operating cost implies a high opportunity cost. As above, if we hold all other vehicle attributes constant, we would expect an increase in operating cost to decrease the utility a consumer gains from the purchase of a vehicle, and thus decrease the probability of purchasing that vehicle. Therefore, we expect a negative coefficient on operating cost.

To determine operating cost requires the fuel economy of the vehicle in miles per gallon, an estimate of how many miles a given consumer will drive that vehicle, a fuel cost, and a time horizon. In our model, we estimate the number of miles a consumer will drive in a year as the reading on the odometer of the trade-in vehicle divided by the age of the trade-in. We use the

EPA-standard \$3.75 per gallon as our fuel cost. For our time horizon, we use automotive research firm R.L. Polk's estimate for the length of time Americans keep new vehicles in 2009: 59.4 months, or 4.95 years. Thus, operating cost is given by:

$$operating_cost_{ij} = \left(\frac{odo_i}{age_i}\right) \left(\frac{1}{mpg_j}\right) (fuel\ cost)(time)$$

where odo_i is the odometer reading on individual i 's trade-in vehicle, age_i is the age in years of individual i 's trade-in, mpg_j is the miles-per-gallon of the new vehicle, and $fuel\ cost$ and $time$ are the constants given above. $Operating_cost$ varies for each consumer i and each vehicle option j , so it receives an i and a j subscript.

To illustrate, if a CARS participant traded in a 10 year old vehicle with 100,000 miles on the odometer and purchased a new vehicle with a fuel economy of 25 miles per gallon, we would estimate the operating cost to that participant over the projected life of the new vehicle to be \$7,425:

$$operating_cost_{ij} = \left(\frac{100000\ miles}{10\ years}\right) \left(\frac{1\ gallon}{25\ miles}\right) \left(\frac{\$3.75}{1\ gallon}\right) (4.95\ years) = \$7,425.$$

3.3.3 Individual-Specific Variables

An individual's income undoubtedly shapes their automobile purchasing behavior. Higher income individuals may be less sensitive to operating costs and may be more influenced by characteristics unobserved in our data, such as luxury, style, power, and brand prestige. As previously noted, several of these characteristics are negatively correlated with fuel efficiency. Consumers therefore face a trade-off between fuel efficiency and luxury, power, and similar characteristics. If they are less sensitive to operating costs, high income individuals may be more

willing to sacrifice fuel efficiency for these attributes. Therefore, we would expect high income individuals to be more likely to purchase vehicles on the lower end of the fuel economy spectrum. However, we might expect the most fuel efficient vehicles, hybrids, to also appeal to high income individuals. Such vehicles are often prohibitively expensive, with purchase prices that may not be justified by their lower operating costs. High income individuals may be attracted to these vehicles as status symbols; they are willing to pay a premium to signify their status as environmentally conscious. Therefore, we anticipate high income to increase the probability that an individual purchases a high fuel efficiency vehicle.

Characteristics of the vehicle a consumer traded in may indicate an individual's preferences. For instance, previous ownership of a domestic vehicle may indicate a preference for models traditionally considered "American." In our model, we do not consider Toyotas and other makes often manufactured in the United States to be domestic vehicles. To the extent that domestic automakers manufacture vehicles of all fuel efficiency levels, it is not clear how loyalty to American brands would impact the probability of purchasing a vehicle of a given fuel economy class. However, domestic automakers may have offered smaller lineup of high fuel efficiency models at the time of the program, in which case loyalty to American brands would drive many consumers towards lower fuel efficiency vehicles. This would lead us to anticipate an increase in the probability of a domestic vehicle owner purchasing a vehicle at the lower end of the fuel economy spectrum and a decrease in the probability of purchasing a more fuel efficient vehicle. In addition, domestic automakers produced many of the largest vehicles on the market in the decade prior to CARS. Previous ownership of a domestic vehicle may indicate a preference for SUVs and pickup trucks. In this case, we would expect a positive coefficient for lower fuel efficiency vehicles and a negative coefficient for higher fuel efficiency vehicles.

Likewise, previous ownership of a light truck may indicate a preference for larger vehicles. The light truck classification includes SUVs, pickups, and vans. Individuals who previously owned such a vehicle could be expected to seek out similar vehicles under the CARS program. Additionally, because it would be easier for these individuals to receive the larger rebate, there was less incentive for them to purchase a high fuel efficiency vehicle. We anticipate previous ownership of a light truck to increase the probability of purchasing a low or medium low fuel efficiency vehicle and decrease the probability of purchasing a medium high or high fuel efficiency vehicle.

4 Results

4.1 Base Specification

The estimated coefficients for our linear utility function are given in Table 3. These results offer insight into the decision process of a consumer participating in the CARS program. With the exception of income for vehicles in the medium high fuel efficiency class, all of the variables in the model are statistically significant at the .001 level. However, not all estimated coefficients are of the anticipated sign. We would expect both of the alternative specific variables, initial cost and operating cost, to have negative coefficients (see Section 3). Contrary to our expectations, operating cost has a positive and significant coefficient. This is likely due to the negative correlation between fuel efficiency and unobserved desirable vehicle attributes such as size, luxury, and horsepower. The positive coefficient on initial cost is as expected, and indicates that an increase in the initial cost for one class of vehicles would decrease the probability that a consumer would purchase a vehicle of that class and would increase the probability of purchasing a vehicle of any other class.

Table 3: Determinants of Automobile Choice In “Cash for Clunkers”

Regressor	Mixed Conditional Logit			
Initial_Cost (in 1000s)	-0.449*** (2.40E-3)			
Operating_Cost (in 1000s)	0.023*** (1.78E-3)			
	Reference Class: Medium			
	Low	Medium Low	Medium High	High
Income (in 1000s)	-0.003*** (1.79E-4)	0.002*** (1.15E-4)	0.000 (2.24E-4)	0.009*** (2.40E-4)
Prev_Domestic	0.786*** (1.68E-2)	0.027*** (8.42E-3)	-0.198*** (1.51E-2)	-0.543*** (1.77E-2)
Prev_Truck	2.298*** (3.47E-2)	0.353*** (8.43E-3)	0.370*** (1.49E-2)	0.300*** (1.86E-2)
Constant	1.863*** (5.23E-2)	1.586*** (1.86E-2)	-2.565*** (2.49E-2)	-0.814*** (3.09E-2)

* $p < 0:10$; ** $p < 0:05$; *** $p < 0:01$ (two-tailed). Standard errors are given in parentheses.

Log likelihood = -698940.51

The alternative specific constants, which capture the average utility of all unobserved vehicle attributes and are interpreted relative to the reference class, the medium category, show that vehicles in the low fuel efficiency class possess unobserved attributes that contribute significantly to utility. Vehicles in the medium low class also have unobserved attributes that add a relatively smaller amount to utility. However, relative to medium class vehicles, the unobserved attributes of medium high vehicles subtract from utility. The same reasons that explain the positive coefficient on operating cost explain the relationship among the constants for low, medium low, medium, and medium high fuel efficiency vehicles: unobserved attributes that add to utility are negatively correlated with fuel efficiency. However, we see that the trend begins to reverse from medium high to high fuel efficiency vehicles. While high fuel efficiency vehicles still have unobserved attributes that on average subtract from utility relative to medium vehicles, the negative effect is much smaller in magnitude than for medium high vehicles. This

suggests that after a certain point, high fuel efficiency becomes desirable in and of itself and begins to outweigh the attributes that are negatively correlated with efficiency.

Previous vehicle ownership contributed significantly to consumer behavior in CARS. Relative to medium vehicles, consumers who previously owned a domestic automobile were more likely to purchase a low fuel efficiency vehicle. Similarly, these same consumers were far less likely to purchase a high fuel efficiency vehicle.. The relationship among the coefficients for previous ownership of a domestic vehicle suggest that, as we move up the spectrum of fuel efficiency classes, vehicles possess fewer of the attributes that owners of domestic vehicles find attractive.

The role of income is more difficult to determine. As income increases, consumers were more likely to purchase a medium class vehicle than a low vehicle. However, the opposite relationship exists between medium low and medium vehicles. As income increases, the probability of purchasing a medium low fuel efficiency vehicle increases relative to the probability of purchasing a medium vehicle. The coefficient on income for medium high vehicles was not significantly different from zero at the .005 percent level. The coefficient on income for high fuel efficiency vehicles, however, is positive, significant, and relatively large in magnitude, suggesting that an increase in income increased the probability of the purchase of a high fuel efficiency vehicle relative to a medium vehicle. This suggests that the vehicles that higher income individuals find attractive belong to either the medium low category or the high category. It may be that the low category consists predominantly of trucks and SUVs, while the medium low category contains luxury sedans. Likewise, the high fuel efficiency class contains many hybrid models that are often prohibitively expensive.

Similarly, the type of vehicle a consumer traded in acted as a predictor for which class of vehicle they would purchase. Consumers who traded in a light truck (Category 1, 2, or 3) were far more likely to purchase a low fuel efficiency vehicle under the CARS program. Likewise, the effects of previous truck ownership on the probabilities of purchasing a medium low, medium high, or high vehicle are positive and significant relative to medium, though smaller in magnitude than for low fuel efficiency vehicles.

To gain a better understanding of these results, we can look at marginal effects. Because the “Cash for Clunkers” program sought to influence consumer behavior through subsidies to reduce the purchase price of vehicles, it is informative to focus on the effects of changes in initial cost. Table 4 summarizes the marginal effects of changes in initial cost for each vehicle class if all variables are set at their means. The entry in the i -th row and j -th column can be interpreted as the change in the probability of purchasing a vehicle of class j given an increase in the price of vehicles in class i . The table is symmetric.

As required by the negative coefficient on initial cost, an increase in the price of a class of vehicle decreases the probability of purchasing that class of vehicle for all consumers and increases the probability of purchasing all other classes of vehicle. That is, the diagonal elements in the table capture the reduced probability of purchasing a vehicle when its initial cost rises.

Each row in the table provides information about the implications of an initial cost increase on the spectrum of choices. For instance, consider the first row, where a \$1000 average increase in the purchase price of low fuel efficiency vehicles would decrease the probability of purchasing a low vehicle by 3.6%. That same \$1000 price increase would increase the probability of purchasing a medium low vehicle by 1.5% and a medium vehicle by 1.8%.

Table 4: Marginal Effects at the Mean for a Change in Initial Cost

	Low	Medium Low	Medium	Medium High	High
Low	-.036*** (2.28E-4)	.015*** (9.70E-5)	.018*** (1.14E-4)	.002*** (2.00E-5)	.001*** (1.30E-5)
Medium Low		-.105*** (5.75E-4)	.074*** (4.41E-4)	.010*** (8.10E-5)	.005*** (5.20E-5)
Medium			-.111*** (5.98E-4)	.012*** (9.60E-5)	.007*** (6.20E-5)
Medium High				-.026*** (1.94E-4)	.001*** (9.70E-06)
High					-.014*** (1.31E-4)

* $p < 0:10$; ** $p < 0:05$; *** $p < 0:01$ (two-tailed). Standard errors are given in parentheses.

Table 4 shows how changes in initial cost, such as an increase or decrease in the “Cash for Clunkers” subsidy for a given class of vehicle, would impact the choice probabilities for an average consumer. We see that a decrease in price of high fuel efficiency vehicles would yield relatively small changes in behavior. In contrast, changes in purchase price of medium class vehicles would produce relatively large effects. For instance, a \$1000 average price decrease for medium vehicles would increase the probability of purchasing a medium vehicle by 11.1%, while decreasing the probability of purchasing a low fuel efficiency vehicle and medium low fuel efficiency vehicle by 1.8% and 7.4%, respectively. These values are large relative to the other values in the table.

The previous table examined marginal effects at the mean of a change in initial cost; that is, the table lists the changes in probabilities given an increase in price with initial cost of each vehicle, operating cost of each vehicle, income of each individual, previous ownership of a domestic vehicle, and previous ownership of a light truck set to their average values. Ideally, this information tells us how a price change will affect the purchase probabilities for an individual with average characteristics. However, it is possible that no real individual and no actual vehicle

option meet the criteria for “average.” Further, it is possible that the marginal effects change drastically as we move away from the mean of each attribute. Therefore, it is useful to calculate average marginal effects. To do so, we calculate a marginal effect of a price increase on the purchase probabilities for each observation and then average the marginal effects. Table 5 summarizes the average marginal effects of a change in the initial cost of a vehicle.

We see that in many cases the average marginal effects differ significantly from the marginal effects at the mean. Particularly, an increase in the average price of low fuel efficiency vehicles on average decreases the probability of purchasing a low fuel efficiency vehicle by 4.8%. This is considerably larger than the decrease of 3.6%, the marginal effect at the mean. Indeed, the average marginal effects of an increase in the purchase price of low efficiency vehicles are consistently greater in magnitude than the marginal effects at the mean. The same pattern does not hold true for a change in the price of high fuel efficiency vehicles.

Table 5: Average Marginal Effects for a Change in Initial Cost

	Low	Medium Low	Medium	Medium High	High
Low	-0.048 (3.50E-2)	0.020 (1.51E-2)	0.023 (1.67E-2)	0.003 (2.25E-3)	0.002 (1.16E-3)
Medium Low		-0.101 (7.22E-3)	0.066 (1.59E-2)	0.009 (2.54E-3)	0.005 (2.77E-3)
Medium Medium			-0.107 (5.42E-3)	0.011 (4.19E-3)	0.006 (3.68E-3)
Medium High				-0.025 (4.80E-3)	0.001 (6.26E-4)
High High					-0.014 (6.35E-3)

4.2 Policy Implications

Our results have a number of implications for the “Cash for Clunkers” program. First, the coefficient on initial cost is considerably larger in magnitude than the coefficient on operating cost. This indicates that the purchase price of a vehicle is of greater importance in determining a consumer’s vehicle choice than the operating cost of the vehicle over its lifetime, which implies that a subsidy scheme, such as “Cash for Clunkers,” could be more effective at influencing consumer behavior than a policy with the same goal that acts on operating cost, such as a fuel tax. However, the positive sign on operating cost in our results indicates that this variable is picking up desirable attributes that are negatively correlated with fuel efficiency. Therefore, our model will not produce meaningful predictions with respect to changes in operating cost. We can, however, use the results of our model to make predictions about the outcome of “Cash for Clunkers” under alternative subsidy schemes.

As required by the inclusion of alternative-specific constants, the predictions of our model match the observed data exactly. Table 6 below shows the predictions of the model alongside the observed data. The entries in the table are the average probabilities that an individual will purchase a vehicle of each class. By construction, the individual probabilities of purchasing each type of vehicle are equivalent to that vehicle’s proportion of total sales.

Table 6: Base Model Predictions

Vehicle Class	Observed	Prediction
Low	.140	.140
Medium Low	.351	.351
Medium	.418	.418
Medium High	.058	.058
High	.033	.033

First, we predict the results of the program in the absence of subsidies. Table 7 shows the result of this exercise.

Table 7: Model Predictions in the Absence of Subsidies

Vehicle Class	Observed	Prediction	Difference
Low	.140	0.287	0.147
Medium Low	.351	0.294	-0.057
Medium	.418	0.348	-0.070
Medium High	.058	0.045	-0.013
High	.033	0.026	-0.007

As the large alternative-specific constant for low fuel efficiency vehicles would lead us to suspect, consumers flock to low fuel efficiency vehicles in the absence of subsidies. The effects of eliminating subsidies entirely are equivalent to the effects of instituting a subsidy that lowers all vehicle prices by the same amount. The large alternative-specific constants for low and medium low vehicles indicate that a scheme that lowers vehicle purchase prices across the board will simply enable consumers to buy low fuel efficiency vehicles with the desirable unobserved characteristics reflected in the constants.

In order to discourage the purchase of inefficient vehicles, an effective subsidy scheme must require new vehicles to meet a certain level of fuel efficiency or make the amount of the subsidy conditional on the increase in fuel efficiency between the trade-in and new vehicle, as in CARS. The CARS program itself, however, had lax requirements for improvement in fuel efficiency, particularly for consumers who traded in light trucks. If the requirements for CARS had been stricter, such that buyers of low fuel efficiency vehicles would not receive a subsidy, our results (see marginal effects in Table 5) suggest that we would see a considerable reduction in the purchase of low fuel efficiency vehicles and increases in the purchases of all other

categories. Below, we predict the results of the program if subsidies for low fuel efficiency vehicles were eliminated entirely.

Table 8: Model Predictions in the Absence of Subsidies for Low Fuel Efficiency Vehicles

Vehicle Class	Observed	Prediction	Difference
Low	.140	0.054	-0.086
Medium Low	.351	0.387	0.037
Medium	.418	0.459	0.041
Medium High	.058	0.064	0.006
High	.033	0.036	0.003

But a subsidy scheme that focuses primarily on reducing the price of the most fuel efficient vehicles might also not be effective. Relative to the other vehicle categories (see Table 4), changes in the initial cost of high fuel efficiency vehicles had the smallest marginal effects at the mean, and much smaller average marginal effects than those for low fuel efficiency vehicles. Our results suggest that a subsidy that only reduces the prices of vehicles with fuel economy of greater than 35 miles per gallon would only slightly increase the probability of purchasing a high fuel efficiency vehicle and would leave the probabilities of purchasing lower fuel efficiency vehicles relatively unchanged. The low price elasticity of high fuel efficiency vehicles, particularly hybrids, may reflect the appeal of these vehicles to high income consumers who are less sensitive to changes in price. Below, we predict the results of increasing the subsidy for high fuel efficiency vehicles by \$1,000.

Table 9: Model Predictions with Increased Subsidies for High Fuel Efficiency Vehicles

Vehicle Class	Observed	Prediction	Difference
Low	.140	0.138	-0.002
Medium Low	.351	0.344	-0.007
Medium	.418	0.410	-0.008
Medium High	.058	0.057	-0.001
High	.033	0.051	0.018

Our results also show that individuals who previously owned a light truck may require more inducement to purchase a fuel efficient vehicle. These consumers make up the majority of the participants in CARS (see Section 2.2). The “Cash for Clunkers” policy had relaxed fuel efficiency improvement requirements for individuals trading in a light truck, which could be seen as an effort by policymakers to provide additional incentive for these individuals. However, the lower requirement led many of these individuals to use the subsidy to purchase a low or medium low fuel efficiency vehicle. Below, we consider an increased subsidy for consumers trading in a light truck. Table 10 shows our model’s predictions in the case of a \$1,000 increased subsidy towards the purchase of a medium, medium high, or high fuel efficiency vehicle for these consumers.

Table 10: Model Predictions with Increased Subsidies for Truck Owners to Purchase a Medium, Medium High, or High Fuel Efficiency Vehicle

Vehicle Class	Observed	Prediction	Difference
Low	.140	0.113	-0.027
Medium Low	.351	0.285	-0.066
Medium	.418	0.495	0.076
Medium High	.058	0.069	0.011
High	.033	0.039	0.006

5. Conclusion

5.1. Caveats

Several issues suggest that we should interpret our results cautiously. First, data constraints led us to use proxy variables that may not accurately reflect the item of interest. For instance, our estimate of the initial cost of a vehicle relies on the Manufacturer’s Suggested Retail Price (MSRP). While this provides a baseline estimate for what each vehicle should cost,

it does not convey what the actual consumer paid. Car buying often involves extensive negotiation over price, and the final price paid might differ substantially from the MSRP. Our estimate of operating cost also relies on variables that may not be an accurate representation of the characteristics we are trying to measure. Our estimate uses the odometer reading on the trade-in vehicle and the age of the trade-in vehicle to determine the average miles per year that the consumer will drive. This measure will only be truly accurate if the consumer purchased the trade-in vehicle new and owned it continuously over its lifetime. If the consumer purchased the trade-in vehicle used, the vehicle's previous owner may have put the majority of the miles on it. In this case, the odometer reading does not capture the driving tendencies of the consumer. Furthermore, we estimate the length of time the consumer will own the new vehicle using the national average for new vehicles. Consumers who participated in CARS may be prone to keeping their vehicles for longer or perhaps shorter periods of time than the average consumer. The length of time a consumer anticipates keeping a vehicle undoubtedly plays a role in the decision, but we were unable to model this type of consumer heterogeneity in our estimate. Additionally, gas prices can differ dramatically from state to state. CARS participants who live in states with systematically higher gasoline prices, such as New York, likely factor higher gas prices into their decision. We do not account for this. Finally, the use of income data for the zip code of the dealership as a proxy for the incomes of consumers who purchased vehicles from that dealership may lead us to misestimate the effect of consumer income on vehicle choice.

Next, our model uses several major simplifying assumptions. First and foremost, we do not include an "outside option." That is, in our counterfactual estimates, we do not cover the possibility that a change in the subsidy scheme could reduce or increase participation in the program. Our results require the assumption that the agents in our data have predetermined to

participate in “Cash for Clunkers.” Including the outside option would require extending the model to cover all potential new car buyers, which in turn would require knowledge of the size of the population of potential car buyers and characteristics of that population. This task is beyond the scope of the current project. An additional issue in our model is the lack of discounting when determining the operating costs of the vehicle over its lifetime. In our model, the operating cost is simply the cost per year multiplied by the expected number of years the consumer will operate the vehicle.

5.2. Summary and Suggestions for Further Research

Our results indicate that a subsidy program, such as “Cash for Clunkers,” will more effectively shift consumers towards fuel efficient vehicles than a fuel tax or similar policy that acts on operating costs. Additionally, our model predictions lead to the conclusion that a subsidy scheme that focuses exclusively or predominately on hybrids and other highly efficient but prohibitively expensive vehicles will be of limited effectiveness relative to a subsidy that targets mid-range efficiency vehicles that have a greater elasticity with respect to price. The greater appeal of these vehicles to the average consumer also increases the potential impact of a subsidy focusing on mid-range fuel efficiency vehicles. We also find that a subsidy should be carefully structured to avoid incentivizing the purchase of low fuel efficiency vehicles. The average utility from unobserved characteristics is greatest for the lowest fuel efficiency vehicles; consumers will gravitate towards these vehicles both in the absence of a subsidy and in the presence of a subsidy that lowers vehicle prices by the same amount across the board. If increasing fuel efficiency is the goal, a sharp decrease or elimination of the subsidy for low fuel efficiency vehicles is necessary to curb the attraction of these vehicles.

Given our data limitations and the construction of our model, we conclude that “Cash for Clunkers” successfully improved the fuel efficiency of the vehicle fleet relative to the purchases that would have occurred in the absence of a subsidy (see Table 7). Furthermore, the outcome could have been improved by eliminating the subsidy for vehicles with a fuel economy of less than 20 miles per gallon. A smaller improvement would also have occurred if the subsidy amount for high fuel efficient vehicles had been increased. This research, however, could be expanded and improved upon in several aspects. A larger choice set would provide a more realistic picture of vehicle choice. Our choice set places many luxury vehicle models in the same category as large trucks and SUVs. Consumers, however, likely do not consider these vehicle choices to be equivalent. We could also improve our results by applying our model and methodology to an expanded dataset that includes additional demographic variables. An ideal dataset would match vehicle purchases to the characteristics of the consumer and free us from the constraint of having to rely on aggregate Census data for information on income and other variables. Finally, the construction and inclusion of variables that capture desirable vehicle attributes that are negatively correlated with fuel efficiency would lead to more accurate results regarding the effect of operating cost on vehicle choice. Achieving the anticipated sign on operating cost would result in more accurate counterfactuals and allow us to better predict the effects of a fuel tax versus a subsidy.

As long as policymakers and the public remain concerned with the environment and with achieving energy independence, the ability of policy to shift the purchasing behavior of consumers towards more efficient vehicles remains relevant. While this research suggests that subsidies can be a powerful tool to increase the fuel efficiency of America’s vehicle stock, it also highlights the challenges of creating a subsidy scheme that achieves the desired results. Concerns

over the environment and the national security implications of energy independence ensure that automobile choice will be an important area of research for years to come.

References

- Allcott, H., & Wozny, N. (2012). "Gasoline Prices, Fuel Economy, and the Energy Paradox" (Working Paper No. 18583). National Bureau of Economic Research.
- Berkovec, J. and J. Rust. (1984). "A Nested Logit Model of Automobile Holdings for One-Vehicle Households", *Transportation Research B*, 19 (4): 275-286.
- Berry, S.T., J. Levinsohn and A. Pakes. (1995). "Automobile Prices in Market Equilibrium", *Econometrica*, 63(4): 841-890.
- Boyd, J.H. and R.E. Mellman. (1980). "The Effect of Fuel Economy Standards on the U.S. Automotive Market: An Hedonic Demand Analysis", *Transportation Research A*, 14A: 367-378.
- Cameron, Colin A., & Pravin K. Trivedi. (2005). Multinomial Models. *Microeconometrics: Methods and Applications* (pp. 491-533). New York: Cambridge University Press.
- Cardell, S.N., F.C. Dunbar (1980). "Measuring the Societal Impacts of Automobile Downsizing", *Transportation Research A*, 14A: 423-434.
- Economics for the Environment Consultancy Ltd. 2008. *Demand for Cars and their Attributes*, report prepared for the Department of Transport, UK, London, January.
- Goldberg, P.K. (1995). "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry", *Econometrica*, 63(4): 891-951.

- Lave, C. A., K. Train. (1979). "A Disaggregate Model of Auto-type Choice", *Transportation Research A*, 13A: 1-9.
- Mannering, F., H. Mahmassani. (1985). "Consumer Valuation of Foreign and Domestic Vehicle Attributes: Econometric Analysis and Implications for Auto Demand", *Transportation Research A*, 19A (3): 243-251.
- Martin, Elliott William. (2009). "New Vehicle Choice, Fuel Economy and Vehicle Incentives: An Analysis of Hybrid Tax Credits and the Gasoline Tax". UC Berkeley: University of California Transportation Center.
- McCarthy, P.S. and R.S. Tay. (1998). "New Vehicle Consumption and Fuel Efficiency: A Nested Logit Approach", *Transportation Research E*, 34(1): 39-51.
- McFadden, D. (1974). "Conditional Logit Analysis of Qualitative Choice Behavior", in P. Zarembka, ed., *Frontiers in Econometrics*, New York: Academic Press.
- Mohammadian, A., E. J. Miller, (2003). "An Empirical Investigation of Household Vehicle Type Choice Decisions", *Journal of the Transportation Research Record*, 1854: 99-106.
- National Highway Traffic Safety Administration. (2009). Transportation Secretary Ray LaHood Kicks Off CARS Program, Encourages Consumers to Buy More Fuel-Efficient Cars and Trucks [Press release]. Retrieved from <www.nhtsa.gov>.
- National Highway Traffic Safety Administration. (2009). Cash for Clunkers Wraps up with Nearly 700,000 car sales and increased fuel efficiency, U.S. Transportation Secretary

LaHood declares program “wildly successful” [Press release]. Retrieved from <www.nhtsa.gov>.

Petrin, A. (2002). “Quantifying the Benefits of New Products: The Case of the Minivan”, *Journal of Political Economy*, 110 (4): 705-729.

R.L. Polk. (2010). Consumers Continuing to Hold onto Vehicles Longer, According to Polk [Press release]. Retrieved from <www.polk.com>.

Train, Kenneth E. 2007. *Discrete Choice Models with Simulation*. New York: Cambridge University Press.