ESSAYS IN INDUSTRIAL ORGANIZATION

by

NICHOLAS G WOOD

A DISSERTATION

Presented to the Department of Economics
and the Graduate School of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

June 2019

DISSERTATION APPROVAL PAGE

Student: Nicholas G Wood

Title: Essays in Industrial Organization

This dissertation has been accepted and approved in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Economics by:

Wesley W. Wilson	Chair
Keaton Miller	Core Member
Van Kolpin	Core Member
Joshua Gordon	Core Member
David Guenther	Institutional Representative

and

Janet Woodruff-Borden	Vice Provost and Dean of the Graduate School

Original approval signatures are on file with the University of Oregon Graduate School.

Degree awarded June 2019

DISSERTATION ABSTRACT

Nicholas G Wood

Doctor of Philosophy

Department of Economics

June 2019

Title: Essays in Industrial Organization

I investigate the effect to which consumer heterogeneity impacts varying markets. I find that it has a substantial effect all the markets that I consider. In the US automobile industry, we model demand for the durable goods using a dynamic model. We expand upon the current literature by allow households the flexibility to select between different classes of vehicles. In further analysis, I develop a method to incorporate microlevel data into the analysis of the automobile industry. I show that, with this approach, national policy can be targeted at specific groups and regions of the populous. For policies considering the automobile industry, I consider the Cars Allowance Rebate System to evaluate policy effects. Finally, I consider the boom of the U.S. smartphone industry. Using a random coefficients estimator, I find that ignoring consumer specific effects would cause the researcher to drastically misestimate gains, and therefor implement poor policy, in the industry.

This dissertation includes previously unpublished co-authored material.

CURRICULUM VITAE

NAME OF AUTHOR: Nicholas G Wood

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:
University of Oregon, Eugene, OR
Kalamazoo College, Kalamazoo, MI

DEGREES AWARDED:
Doctor of Philosophy, Economics, 2019, University of Oregon
Master of Science, Economics, 2015, University of Oregon
Bachelor of Arts, Mathematics, Physics, and Economics, 2014, Kalamazoo College

AREAS OF SPECIAL INTEREST:
Industrial Organization
Applied Mircoeconomics
Econometrics

GRANTS, AWARDS AND HONORS:

Dale Underwood Outstanding Graduate Student Scholarship, University of Oregon, 2015

Kleinsorge Fellowship Award, University of Oregon, 2014

Honors, Senior Individualized Project, Kalamazoo College, 2014

ACKNOWLEDGEMENTS

I thank the members of my committee for their guidance in helping me execute my research agenda. In particular, I would like to thank Wesley W. Wilson and Keaton Miller for their invaluable mentorship and patience. Finally, I thank my family and fellow graduate students for their help.

To Gertrude.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER I

INTRODUCTION

There are millions of consumers, each with their own unique characteristics that paricipate in numerous industries across the globe. It is our duty as economists to understand how the interactions of a multitude of vastly different consumers impact different markets in order to recommend effecive policies. This dissertation proposes new means of modelling consumer behavior to better inform policy. In Chapter II Keaton Miller, Wesley Wilson, and I examine policy in the automobile industry. This chapter is co-authored with the above economists and is currently unpublished.

Industrial policies are often sold to the public by appealing to multiple constituencies which are often in opposition. For example, the 2009 Cash for Clunkers program, which provided subsidies to consumers who chose to scrap old automobiles and purchase new vehicles, was promoted under the trifecta of ideas that the exchanges would reduce vehicle emissions, combat rising inequality, and stimulate a troubled manufacturing sector. We evaluate the performance of the policy against these stated goals. As vehicles are durable goods and the policy differentiated between cars and light trucks, we introduce a dynamic, partial equilibrium portfolio model with vertical product differentiation and heterogeneous consumers. We calibrate the model to match data on new and used car purchases from 1998-2011 and find that the program generated roughly \$0.11 in environmental savings per \$1 in subsidies. Since subsidies largely went to middle-income marginal consumers, the program exacerbated income inequality. By pulling forward future truck purchases and incentivizing a switch to cars, the program resulted in small gains to automakers relative to the program's budget. Our results suggest that Cash for Clunkers and similar policies failed to achieve the trifecta.

The most famous vehicle exchange program in history, Cash for Clunkers, ran during the summer of 2009. This program aimed reduce vehicle emissions by removing older and larger vehicles, that were responsible for higher levels of pollution than newer vehicles, from the population. In Chapter III, I develop a structural model that incorporates microlevel data into the estimation procedure. I do this by implementing two machine learning techniques: principle component analysis and neural networks, to reduce the computational burden of including demographic data into a structural model. With a micro-founded structural model, I can examine how deviations in policy would have affected different demographic groups or regions of the country by preforming counterfactual exercises.

Finally, in Chapter IV, I examine the smartphone industry. Over the last decade, the modern smartphone has gone from luxury good, to an integral part of our daily life. During this time, the marketplace has undergone rapid growth and technological improvement. I apply a random coefficient model to measure the welfare effects of the introduction of the smartphone and the resulting technological progress. Using data on phone characteristics from 2010-2014, I find that the average American benefited by $11.50 per year. This comes out to over $16 billion in total consumer surplus over the sample period, making the smartphone industry one of the largest generators of consumer welfare.

CHAPTER II

STIMULUS, ENVIRONMENTALISM, AND INEQUALITY REDUCTION THROUGH INDUSTRIAL POLICY: DID CASH FOR CLUNKERS ACHIEVE THE TRIFECTA?

The contents of this chapter were developed by several members, including Keaton Miller, Wesley Wilson, and myself. Keaton Miller and Wesley Wilson contributed to this work by participating in the writing and formatting of final draft that was sent out for publication, as well as providing guidance towards the implementation of the economic analysis. I was the primary contributor toward obtaining and modifying the data, solving the economic relations, and preforming the computational analysis.

> *"[Cash for Clunkers] holds the promise of performing a remarkable public policy trifecta – stimulating the economy, improving the environment, and reducing income inequality all at the same time." (Blinder, 2008)*

Introduction

Legislators often face a trade-off between policies designed to support or stimulate industrial employment and productivity and policies designed to promote environmental protection, as industrial activity generally leads directly to environmental harms (Ayres and Kneese, 1969; Mäler, 1974; Helfand et al., 2003). Often, policymakers search for programs which can be promoted as "win-win" in the sense that both the employment and environmental preferences of constituents and interest groups can reasonably be said to be satisfied (Becker, 1983; Aidt, 1998).

One class of policies designed to balance these concerns is "tied-scrappage" programs, in which individuals are offered subsidies to scrap older, highly-polluting durable goods and replace them with newer, more efficient models. For example, the 2009 U.S. Cash for Clunkers program (hereafter C4C) offered consumers who replaced personal automobiles (cars and light trucks) up to $4,500. Indeed, C4C was promoted under not only employment and environmental rationales, but also under the idea that the subsidies may allow low-income consumers to purchase newer automobiles, thus reducing the purchasing-power impacts of income inequality (Blinder, 2008). Tied scrappage programs have been implemented or proposed in other jurisdictions as well (The Economist, 2009): one program in Germany has led to over a million new vehicle registrations (Böckers et al., 2012; Gürtler et al., 2016), and India's government has proposed a program designed to take 28 million vehicles off of the road starting in 2020 (PTI, 2018).

In this paper, we evaluate C4C against these multiple goals. The degree to which any tied-scrappage program is able to realize equality, environmental, or stimulus benefits depends in part on the mix of consumers who receive subsidies, the realized environmental differences between the scrapped and new goods, and, since these goods are durable, the counterfactual sequence of purchases which would have taken place if the subsidies were not available. In particular, C4C offered the highest subsidies to consumers who replaced trucks with cars. As trucks generally emit a greater amount of carbon dioxide per mile travelled than cars, these trades may be particularly environmentally friendly on a mile-for-mile basis. However, as trucks have higher prices than cars, substitution between a future purchase of a truck and a present purchase of a car may represent lost present value to a risk-neutral firm.

Understanding these intertemporal linkages is crucial to evaluating the impacts of C4C and thus we develop a dynamic model of portfolio choice. In the model, consumers are differentiated by income and tastes, and choose a portfolio of (potentially multiple) vehicles to own. Vehicles are vertically differentiated by age and horizontally differentiated across cars and trucks. We calibrate the model to match data on new and used vehicle purchases from the Consumer Expenditure Survey from 1998-2011. We capture the impact of the 2009 program by simulating the sequence of purchases from 2009 onward both with and without subsidies.

We find C4C generated $317 million in environmental benefits through reduced CO_2 emissions from travel under a budget of $3 billion. By pulling forward vehicle purchases and encouraging substitution from trucks to cars, the program induced $8.7 billion of additional private expenditures during its lifespan, but reduced private expenditures on automobiles in the five years following its expiration by $5.5 billion. We find the program also had negative consequences for inequality. Most consumers participating in the program were consumers who would have purchased a vehicle in the next five years even without the subsidies. Therefore, from the perspective of the "inside share" of vehicles, the subsidies largely went to infra-marginal consumers, who are generally of medium-to-high income.

Our framework allows us to consider alternative designs for the C4C program. We find that focusing on the incentive to trade in trucks for cars could increase the environmental benefit by up to 30%, but at the cost of exacerbating the other tradeoffs of the program. Focusing on immediate stimulus, an alternative program design could increase consumer spending to \$12.8 billion with similar tradeoffs. However, we are unable to find a version of the policy which successfully accomplishes the “policy trifecta” of simultaneously reducing inequality, environmental improvements, and incentivizing increased spending.

Our work builds upon the literature examining interactions between primary and secondary markets both with and without transaction costs (Rust, 1985; Stolyarov, 2002). Gavazza et al. (2014), hereafter GLR, introduce a dynamic model of the used car market which allows for variable transaction costs. Barahona et al. (2015) apply the GLR model to show that Chile’s day-of-week driving restrictions increase fleet turnover. In a working paper, Quan and Singer (2015) use the model to compare the ‘tied subsidy’ design of C4C to a simple subsidy design (in which no scrappage is required) and find the C4C mechanism design is an improvement over the simple subsidy with respect to environmental goals. Gavazza and Lanteri (2018) extend the model to incorporate aggregate credit shocks and show that such shocks affect the vehicle market by flowing through the decisions of households of different income levels who play different roles within the secondary market. Relative to these other papers, we allow for horizontal, as well as vertical, product differentiation and corresponding heterogeneity in preferences for vehicle portfolios across households.

Our work also contributes to the literature studying the tradeoffs between the automobile industry and environmental conditions. Berry et al. (1996) estimate a hedonic marginal cost function and find that tightened emissions standards significantly increased the cost of automobile production. Munk-Nielsen (2015) analyzes tax policies around diesel automobiles in Denmark with a model where households decide both which vehicles to purchase and how much to drive. He finds that re-optimizations along both margins in response to changes in policy imply that emissions reductions programs may be less effective than they might seem. Knittel et al. (2016) use data from California to show that the ambient air pollution generated by automobile congestion significantly increases infant mortality rates, particularly for premature or low birth-weight infants.

In addition to the work mentioned above, C4C has received considerable attention from other perspectives. Mian and Sufi (2012) and Li et al. (2013) examine consumer decisions using difference-in-differences frameworks and find that most of the participants would have purchased a new car in the immediate future even in the absence of the program. Hoekstra et al. (ming) use a regression discontinuity design to show that more than half of the subsidies went to households who would have purchased even without the program. Copeland and Kahn (2011) consider vehicle production and that almost all C4C sales came from previously overstocked inventories, implying that the program may not have induced any new production. Abrams and Parsons (2009) estimates the environmental benefits from the program under the assumption that traded-in vehicles would have remained on the road for three more years. They estimate a benefit of $600 per vehicle, which is much less than the total subsidy. Lenski et al. (2010) find that the estimated cost of a metric-ton reduction of greenhouse gas under the program was about $600, which is significantly larger than the EPA's $13 estimate of the cost of CO2 reduction through more direct channels. Li and Wei (2013) evaluate the tradeoff between the stimulus and environmental objectives with a focus on the range of vehicles purchased under the program, as opposed to our focus on multi-vehicle households and the interplay between the primary and secondary markets. Relative to this literature, we make two primary contributions.

First, we focus on the horizontal differentiation between cars and trucks which is relevant due to the specific design of C4C. Second, instead of focusing on any single policy goal, we develop and apply a framework that captures each of the stated policy goals in a consistent way.

We proceed by providing details on the C4C program and our data in Section 2.2. Section 2.3 introduces our model of dynamic vehicle portfolio choice. We calibrate the model to the data on the U.S. primary and secondary automobile industry in Section 2.4. In Section 2.5 we present our evaluation of C4C through counterfactual exercises. Section 2.6 concludes.

The Cash for Clunkers Program

The Cars Allowance Rebate System, also known as Cash for Clunkers (C4C), was introduced by the United States federal government in 2009, and ran from July through August of that year. C4C offered participants a credit for scrapping their old vehicle when buying a new vehicle under certain conditions. Only vehicles with a combined city/highway miles-per-gallon (MPG) rating of 18 or fewer, as measured by the Environmental Protection Agency (EPA), were eligible to trade-in for scrap. If the new vehicle purchased had an EPA combined rating of at least 4 MPG higher than the trade-in, participants would receive a $3,500 credit. If the new vehicle represented an improvement of 10 MPG or greater, the credit increased to $4,500.[1]

These thresholds were modified for owners of trucks. Consumers trading in light-duty trucks needed only to purchase a new vehicle with a two MPG rating improvement and were eligible for the increased credit with an improvement of

[1] Consumers could view these rules via a government website, archived via: https://web.archive.org/web/20100101064549/http://www.cars.gov/faq

only five MPG.[2] Given the option set in the new vehicle market in 2009, these thresholds made it almost impossible to receive the large credit while trading in an old car for a new truck. Conversely, it was almost guaranteed that participants would be eligible for the large credit it they traded in a truck for a new car. These features underscore the need for a model of differentiated choice between the categories of cars and trucks.

The program was originally signed into law as part of the Supplemental Appropriations Act of 2009 and began on July 1. Congress initially appropriated $1 billion for the program, though funds were nearly exhausted by July 30. An additional $2 billion was appropriated by August 7. All funds were spent by August 24 (Bunkley, 2009).

To understand the direct impact of C4C in terms of trade-ins and purchases, we obtain data on the program from the National Highway Transportation Safety Administration (NHTSA). The NHTSA data covers all trade-ins and new-vehicle purchases related to the Cash for Clunkers program. An observation in this data is a vehicle model and type, with data provided on the number of times the vehicle was traded-in or purchased (if it was new in 2009) as part of the program, as well as other descriptors of the vehicle. Summary statistics for the program are in Table 1 below. The first column summarizes the data for models classified as "cars," and the second column summarizes the data on "truck" models. Trucks consisted of roughly 85% of the vehicles traded-in, but only 41% of the vehicles purchased. This suggests that the program generated substitution from trucks to cars.

---

[2]There were slight differences between the thresholds for light-duty (e.g. GMC Sierra 1500), medium-duty (e.g. Ram 3500), and heavy-duty (e.g. Ford F-750) trucks. As light-duty trucks comprised roughly 80% of the trucks traded-in and purchased under the program, we abstract away from differences within the truck category.

TABLE 1.
Cash for Clunkers Summary Statistics by Vehicle Type

	Cars	Trucks
Trade-ins	102,638	574,443
New Vehicles	397,182	279,899
Average Trade-in MPG	17.5	15.5
Average New Vehicle MPG	27.9	20.7
Average Age of Trade-ins (years)	16.6	15.5

Notes: Data from National Highway Traffic Safety Administration. MPG is the EPA combined highway/city rating.

While these data are illustrative, they must be placed in the context of the steady-state pattern of vehicle purchase, ownership, and replacement decisions by U.S. households. To capture these patterns, we obtain data from the Consumer Expenditure Survey (CEX). The CEX is a publicly available rolling-panel interview survey conducted every quarter by the Bureau of Labor Statistics and contains detailed data on American household characteristics, consumption habits, and vehicle holdings.

We collect the CEX survey responses from 1996-2011. We use the Q1 sample of each year. We observe the number of vehicles owned, and the age and price paid for each. The data are summarized in Figure 1 for the truck and car vehicle types. Panel (a) reports the stock of vehicles by vintage, normalized to the number of new cars purchased by households in the original model year. The stock temporarily rises above one as commercial vehicles (e.g. rental fleets) are sold to households. This effect is more prominent for cars than trucks. Panel (b) reports the average price of cars by vintage. Trucks are more expensive than cars, but both depreciate at roughly the same rate.

FIGURE 1.
Distributions of Vehicle Stocks and Prices by Age

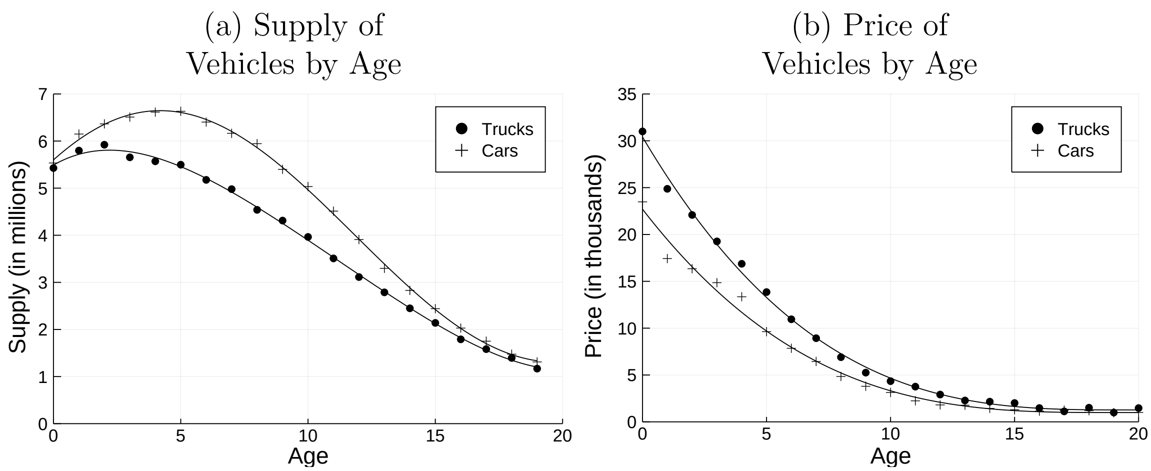

Notes: Data from Consumer Expenditure Survey, 1996-2011 Q1 sample. Points denote sample averages, lines are best-fit 5th-order polynomials. Age refers to years since model year.

Model

Our dynamic partial equilibrium model of vehicle purchase, ownership, and resale has features that are similar to the approach of GLR but includes horizontal differentiation between vehicles. There is a measure of infinitely-lived households denoted by i with common discount factor β, each of which may own up to two vehicles. Each vehicle has a type $m \in \{C, T, 0\}$, denoting cars, trucks, and no vehicle, respectively. Vehicles also have a vintage $a \in \mathbb{Z}$ with $0 \leq a < A^*$ where 0 represents the most recent vintage and A^* is the last vintage in the stock of vehicles.[3]

Households enter each period t with a mix of zero, one, or two vehicles of different ages. They earn constant income y_i drawn from a log-normal distribution with mean μ_y and variance σ_y. Households incur operating costs $C^{m \times m}$ that

[3]Imposing a finite A^* implies that all vehicles are scrapped once they reach a certain age and allows us to ignore vehicles which make up a minuscule portion of the fleet.

depend on the mix of vehicles they own that period. They receive flow utility from their vehicles and may participate in the primary or secondary vehicle markets. After purchasing and/or selling vehicles, they receive log utility from their net income, and enter the next period with their updated vehicle inventory.

To construct the flow utility from vehicle ownership, we let q_a^m be the quality of each vehicle. As this quality is fully flexible across vehicle types and ages, we can without loss of generality set the marginal utility of owning a single vehicle with respect to that vehicle's quality to one. Households that own two vehicles receive utility from the quality of both. Without loss of generality, we defined the "second" vehicle as the vehicle which is older (i.e. the vehicle with a higher a). To allow for differences in vehicle usage patterns across vehicle inventories, we allow for the possibility that the marginal utility of a second vehicle may differ from the (normalized) marginal utility of the first vehicle. Furthermore, to allow for differences in usage patterns across households, we allow households to differ in their marginal utility of second vehicle ownership by the type of the vehicle. Let γ_{iT} be the marginal utility coming from the quality of a second truck when the household already has a truck and let γ_{iC} analogously be the marginal utility coming from the quality of a second car. Let γ_i be the marginal utility of the second vehicle's quality when the two vehicles are of different types. Let m and n be the types of the two vehicles and let a and b be the ages of the two vehicles. The flow utility of vehicle ownership is therefore time-invariant and can be written as

$$u_i(m,n,a,b) = \qquad (2.1)$$

$$\begin{cases} 0 & \text{if } m = 0 \\ q_a^m & \text{if } m \neq 0, n = 0 \\ q_a^m + q_b^n(\gamma_{iT} 1^{\{m=n=T\}} + \gamma_{iC} 1^{\{m=n=C\}} + \gamma_i 1^{\{m \neq n\}}) & \text{if } m \neq 0, n \neq 0 \end{cases}$$

In each period, households may participate in the vehicle market by purchasing a new vehicle, selling an old vehicle, or both. We take vehicle markets to be competitive in the sense that in each period households observe a schedule of vehicle prices that vary by type and age $\{p_a^m\}$ and take those prices as given. Households which sell a vehicle in the secondary market incur a transaction cost that may vary with the price of the vehicle. We define the effective re-sale price of a vehicle via $\Psi(p) = \psi_0 + \psi_1 p$. For simplicity, we allow households to purchase and/or sell only one vehicle in each period.

We now have nearly all the ingredients needed to write household i's recursive value function. For ease of notation, we re-frame the household's choice problem in terms of choosing a portfolio of vehicle types and vintages to own tomorrow, where the cost of transitioning between portfolios depends on the household's current inventory. Let $R(m', n', a', b'; m, n, a, b)$ be the revenue (or cost) which comes from transitioning from portfolio (m, n, a, b) to portfolio (m', n', a', b'). There are fourteen possible ways in which a household can interact with the market and update its portfolio based on their current state. We therefore define $R(\cdot)$ piecewise via

$$
R(m', n', a', b'; m, n, a, b) = \tag{2.2}
$$

$$
\begin{cases}
0 & \text{if } m = n = m' = n' = 0 \\
0 & \text{if } m' = m, a' = a + a, n' = n = 0 \\
0 & \text{if } m' = m, n' = n, a' = a + 1, b' = b + 1 \\
\Psi(p_a^m) & \text{if } m' = n' = n = 0, m \neq 0 \\
\Psi(p_a^m) & \text{if } m' = n, n' = 0, a' = b + 1 \\
\Psi(p_b^n) & \text{if } m' = m, n' = 0, a' = a + 1 \\
-p_{a'}^{m'} & \text{if } m = n = n' = 0, m' \neq 0 \\
-p_{b'}^{n'} & \text{if } n = 0, n' \neq 0, m' = m, a' = a + 1 \\
-p_{a'}^{m'} & \text{if } n' = m, b' = a + 1, n = 0, m' \neq 0 \\
\Psi(p_a^m) - p_{a'}^{m'} & \text{if } \neg(m' = m, a' = a + 1), m' \neq 0, m \neq 0, n = 0 \\
\Psi(p_a^m) - p_{a'}^{m'} & \text{if } a' \leq b', n' = n, b' = b + 1, \neg(m' = m, a' = a + 1) \\
\Psi(p_b^n) - p_{b'}^{n'} & \text{if } a' \leq b', m' = m, a' = a + 1, \neg(n' = n, b' = b + 1) \\
\Psi(p_a^m) - p_{b'}^{n'} & \text{if } a' \leq b', a' = b + 1, m' = n, n' \neq 0 \\
\Psi(p_b^n) - p_{a'}^{m'} & \text{if } a' \leq b', n' = m, b' = a + 1, m' \neq 0 \\
-\infty & \text{else}
\end{cases}
$$

In this equation, the first case occurs when the household owns no vehicles and does not participate in the vehicle market. The second and third cases occur when the household allows their vehicle(s) to age without buying or selling. The fourth case occurs when a single vehicle household sells their only vehicle. The fifth and sixth cases occur when a two vehicle household sells their newer or order vehicle, respectively. The seventh case occurs when the household buys a first vehicle and the eighth case represents the household buying a second vehicle that is older than their first. The ninth case captures households who buy a second vehicle which is younger than the first at the time of purchase so that the first vehicle becomes the second vehicle in the ordering of the model. The tenth case occurs when a household with a single vehicle trades it in for a new one. The eleventh and twelfth cases occur when a household replaces their first or second vehicle, respectively, while maintaining the order between them. The thirteenth case occurs when households sell their first vehicle and buy a vehicle that is older than their current second vehicle, and the fourteenth case captures the converse; a household which sells its second vehicle and buys a vehicle which is newer than their current first vehicle. All other portfolio choices involve transactions over more than one vehicle and are assumed to be impossible.

With $R(\cdot)$ in hand, we can write the recursive value function as

$$\begin{aligned} V_i(m,n,a,b) &= u_i(m,n,a,b) \\ &\quad + \max_{m',n',a',b'} \left[\ln\left(y_i - C^{m,n} + R(m',n',a',b';\cdot)\right) + \beta V_i(m',n',a',b')\right]. \end{aligned} \tag{2.3}$$

Steady State Equilibrium

Quan and Singer (2015) show that if transaction costs are sufficiently high, households can be separated into groups which differ in the number and types of vehicles owned. Within a group, households choose to trade-in and purchase replacement vehicles in such a way to maintain their number-type portfolio (though the vintages may be different). We therefore focus on a steady-state equilibrium for the purposes of solving the model and recovering parameters from the data, and introduce C4C in Section 2.5 as an unforeseen temporary shock.

Let $h_a^m(i)$ be the number of vehicles of type m and age a owned by household i, and x_a^m be the stock of vehicles, determined exogenously. Given the set of vehicle stocks and qualities as well as the distribution of household preferences, an equilibrium is a set of decision rules for each household $\{m_i^*, n_i^*, a_i^*, b_i^*\}$ and prices p_a^m such that the decision rules solve the maximization problem of Equation 2.3 and markets clear:

$$\int h_a^m(i)di = x_a^m \quad \forall a, m \tag{2.4}$$

We adopt this partial equilibrium setting due to the complexity of modeling general equilibrium in the primary market for vehicles, particularly in a setting where manufacturers are forward-looking. Furthermore, as seen in Figure 1, the stock of vehicles owned by households is increasing in the vintage in the first few years, indicating that a general equilibrium model of the secondary market would need to consider the decisions of commercial fleet owners to resell their vehicles. Finally, as the C4C program as implemented influenced a small number of vehicle purchases – approximately 677,000 – relative to the size of the primary market – 10.4 million in 2009 according to an industry group (NIADA, 2017) – it is unlikely that manufacturing or commercial fleet decisions would be altered by the presence of the program. Indeed, Copeland and Kahn (2011) conclude that the program in the U.S. may not have induced new production. Furthermore, Gürtler et al. (2016) studied a similar program in Germany that affected over two-thirds of new vehicles purchases in that year and found that the price of used vehicles only rose by 6%.

Separating households into groups allows us to characterize the decision problems of those groups independently. Since households play a repeated pattern of purchasing vehicles of a certain vintage, holding them until they age a certain amount, and then replacing them with another vehicle of the original vintage, the recursive value function for each household type, and therefore the maximization problem for consumers, can be rewritten in a non-recursive form. For example, households which own a single car follow decision rules that can be represented as $\{\tilde{a}_i, \tilde{b}_i\}^C$, where $\tilde{a}_i$ is the vintage of car that the household purchases when it decides to trade-in its current vehicle, and $\tilde{b}_i$ is vintage at which the trade-in occurs. In steady state, the household's value function can be rewritten in the form of these decision rules with

$$V^C(y) = \max_{a,b} q_a^C + \ln(y - p_a^C - C^C) + \tag{2.5}$$
$$\frac{1}{1-\beta^{b-a+1}} \left[\sum_{j=1}^{b-a} \beta^j [q_{a+j}^C + \ln(y - C^C)] + \beta^{b-a+1} \left[q_a^C + \ln(y - p_a^C + \Psi(p_b^C) - C^C) \right] \right].$$

These reformulations of the value function allow us to reduce computational effort as instead of computing a recursive value function through e.g. value function iteration, we can directly solve for the decision rules as a function of the qualities and prices of the relevant vehicles and then test the market clearing condition.[4]

Calibration

To evaluate C4C, we must first fit the model to the steady-state vehicle portfolio equilibrium in the data. We calibrate some of the parameters directly

[4]The non-recursive value function for two-vehicle households is considerably messier and more difficult to compute. To reduce the computational burden, we restrict one-car, one-truck households to replace their vehicles at the same age (e.g. their car is replaced when it reaches age 15, and their truck is replaced when *it* reaches age 15). As we allow households to re-optimize if they participate in C4C (see Section 2.5), this simplification should not qualitatively affect the results.

from the data and literature, and solve for the rest through matching the model's output to the data. Table 2 summarizes the model parameters. We proceed by detailing each of these parameters.

TABLE 2.
Summary of Model Parameters and Values

Variable	Description	Value
β	Discount rate	.95
μ_y	Mean of household income distribution	10.5
σ_y	Variance of household income	1.06
A^*	Maximum vehicle age	20
x_a^m	Vehicle stock by age and type	See Figure 1
p_a^m	Vehicle prices by age and type	See Figure 1
$C^{m\times m}$	Vehicle operating cost	See Table 3
ψ_0	Constant transaction cost	-624 (see Figure 2)
ψ_1	Proportional transaction cost	0.83 (see Figure 2)
q_a^m	Vehicle quality by age and type	Solved for below
γ_{iT}	Distribution of marginal utility of trucks	Solved for below
γ_{iC}	Distribution of marginal utility of cars	
γ_i	Distribution of marginal utility of variety	

We set the yearly household discount rate β to 0.95 (Gavazza et al., 2014), and calibrate the log-normal income distribution from the CEX data to $\mu_y = 10.5$ and $\sigma_y = 1.06$. Figure 1 shows the stock of vehicles falls below 1% of its original value after 20 years, and so we set the maximum age of vehicles, A^*, to 20. We also use the CEX data as seen in Figure 1 to calibrate the stock and price of vehicles for each age and type.

Operating costs $C^{m\times m}$ are the sum of the vehicle maintenance, registration, and fuel costs incurred by households and vary by the household's vehicle portfolio of vehicle types. Table 3 summarizes the median annualized operating costs captured by the CEX in Q1 of 2009 by household type.

TABLE 3.
Annualized Vehicle Operating Costs by Vehicle Portfolio Type

Portfolio	Obs.	Operating Cost ($)
No Vehicles	442	0
One Car	777	1,200
One Truck	435	1,760
Two Trucks	360	2,840
Two Cars	403	2,332
One of Each	1,176	2,520

Notes: Data from Q1 CEX 2009. Operating costs are in 2009 dollars and are annualized from median quarterly measure.

To calibrate the transaction cost function $\Psi(p)$, we obtain data from the Kelly Blue Book (KBB), July-September 2009 edition. KBB reports prices of used vehicles which were originally produced from 1994-2008. We observe four different values for each vehicle: the trade-in value for vehicles in fair condition; the trade-in value for vehicles in good condition; the private-property value of vehicles in good condition; and the retail value of vehicles in excellent condition. Following (Quan and Singer, 2015), we regress the trade-in value of vehicles in good condition on the private-property value for vehicles in good condition to estimate the parameters ψ_0 and ψ_1. We limit the observations used in this regression to those models for which we observe at least 250 trade-ins as part of C4C. The data and estimated function are displayed in Figure 2. The results show that a linear fit is a reasonable approximation to the difference between the value of a vehicle and what a consumer can expect to receive when trading that vehicle in. This relationship extends to both extremes of the distribution of vehicle values.[5]

[5]We also experimented with varying transaction costs by vehicle type and found nearly identical results.

FIGURE 2.
Estimated Transaction Costs from Kelley Blue Book

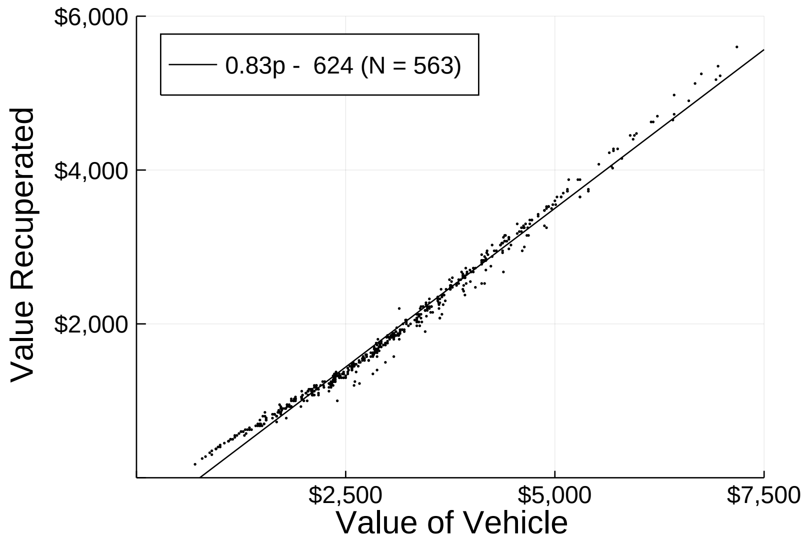

Note: The horizontal axis is the private-property value of a vehicle of a specific model-year in good condition. The vertical axis is the trade-in value of the same model-year in good condition. The difference, therefore, is the effective transaction cost for that model-year. Each dot represents a model-year observation that had at least 250 trade-ins as part of the program.

The remaining parameters capture vehicle quality by age and type, and the distribution of marginal utilities for second vehicles across households. We set these parameters by solving the model and matching the stock of vehicles x_a^m and the relative frequencies of the household types in Table 2.

Since γ_{iT}, γ_{iC}, and γ_i capture the marginal utility of a second vehicle, it is reasonable to assume that they should be drawn from a distribution with support between zero and one. The NHTSA data imply that the C4C program successfully incentivized consumers to switch their portfolio of vehicle types, which implies that there may be some correlated between the values of γ_T, γ_C, and γ for a given household. Though we have six moments to match – the six types of households seen in Table 3 – we have only three degrees of freedom as the relative frequencies must sum to one, and the sum of the households with different vehicle portfolios must match the number of cars and trucks in the stock. To pin down three potentially correlated distributions with three free parameters, we use

$$\gamma_{iT} \sim \text{Beta}(.5, \beta_T),\ \gamma_{iC} = \gamma_{iT} + \varepsilon_C,\ \gamma = \gamma_T + \varepsilon; \qquad (2.6)$$
$$\varepsilon_C \sim \text{Exp}(\lambda_C),\ \varepsilon \sim \text{Exp}(\lambda).$$

This specification roughly imposes the assumption that marginal utilities are between zero and one. The support of the Beta distribution is the interval $(0, 1)$, so draws of γ_{iT} fall in this range with probability 1. Draws of γ_{iC} and γ_i will fall in the same range with some probability that depends on the magnitudes of λ_c and λ.[6] The degree of correlation depends on the magnitude of λ_c and λ. If the means of the shocks are quite low, the degree of correlation is high, and γ_c, and γ should almost always remain smaller than one.

With this specification in mind, we match the data via a nested fixed point algorithm. In the inner loop, we take the parameters $\{\beta_T, \lambda_C, \lambda\}$ as given and solve for the schedule of vehicle qualities q_a^m to match the stocks in the data. We do so by starting with a guess of the qualities, solving for the optimal portfolio decision rules m^*, n^*, a^*, b^*, and aggregating over the households to find the stock of each vehicle predicted by the model.[7] We then adjust the qualities in the direction of the difference between the stock in the data and the model-predicted stocks. In the outer loop, after solving for the qualities, we aggregate the households by vehicle type portfolios and compare the model-predicted portfolio shares to the relative frequencies in the data as represented in Table 3. We calculate the distance

[6]At the estimated values, the probability of the γ draws falling outside of the interval $(0, 1)$ is close to zero. As a consequence, our choice of γ_{iT} as the 'baseline' preference is not likely to affect the results.

[7]We use 5,000 household draws and integrate-through-simulation.

between these vectors and use the Nelder-Mead hill-climbing procedure (Nelder and Mead, 1965) to find the parameters $\{\beta_T, \lambda_C, \lambda\}$ that minimizes this distance.

The results can be seen in Figure 3. Panel (a) plots the densities of the γ_i terms under the parameters found through the procedure above: $\beta_t = 1.06, \lambda_c = 17.24, \lambda = 13.69$. Because the estimated value of the variety parameter is lower than that of the preference for a second car, the average consumer has a stronger preference for variety (the second vehicle being different than the first) than either a second car or truck. The associated qualities that form the steady-state equilibrium from these γ_i terms are shown in Panel (b) of the Figure. The qualities are decreasing in vintage; trucks have a higher quality than cars of the same age.

FIGURE 3.
Estimated Household Preferences and Vehicle Qualities

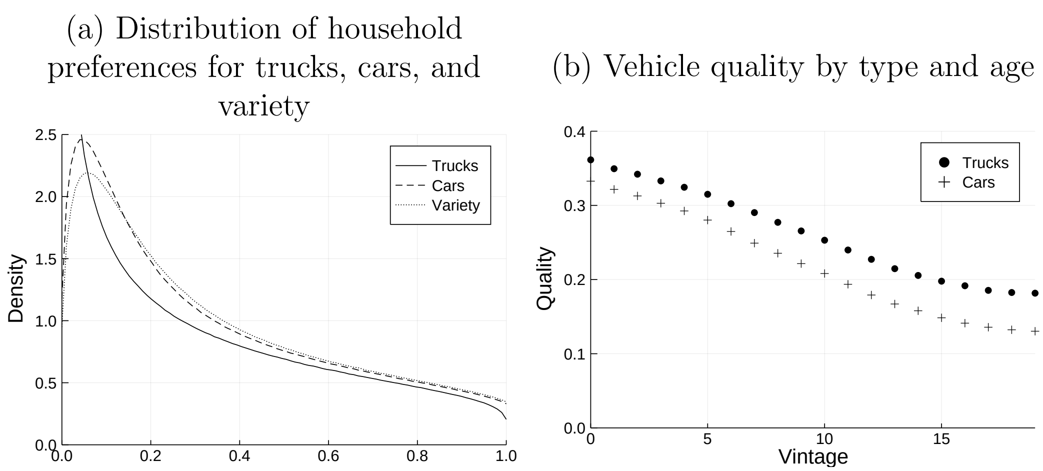

To evaluate the fit of the model, we examine the extent to which the model matches moments from the CEX data which were not used to calibrate parameters. In particular, we examine the rate at which households with different income levels select different types of vehicle ownership. The results from this exercise can be found in Table 4. In general, the model captures the relative share of each portfolio type well. The main difference comes from the existence of a cutoff point in the income distribution, below which zero households participate in the vehicle market and above which all households participate in the vehicle market. While the percentage of households with zero vehicles is declining in income in the data, the pattern is not as extreme as predicted by the model.

TABLE 4.
Household Vehicle Type Portfolios by Income Quartile,
Model-Predicted and CEX Averages

Income Quartile	Zero Vehicles	One Vehicle	Two Trucks	Two Cars	One Each
Lowest	53.8%	34.4%	2.3%	0.4%	9.1%
	37.1%	43.2%	3.2%	5.6%	10.9%
Mid-Low	0.0%	49.4%	15.9%	5.8%	28.9%
	14.2%	48.4%	6.5%	8.5%	22.4%
Mid-High	0.0%	30.9%	9.7%	15.4%	44.0%
	7.9%	28.9%	11.9%	13.8%	38.0%
Highest	0.0%	14.1%	12.7%	17.9%	55.2%
	0.3%	14.7%	16.5%	15.9%	52.6%

Notes: Model-predicted values are reported first; values from the data are reported underneath.

Evaluating Cash for Clunkers

With the steady-state equilibrium in hand, we now turn to evaluating the C4C program and other programs in a similar vein. We model C4C as a one-time shock in which certain eligible households receive a subsidy if they trade-in their vehicle and purchase a new one. Vehicles in the model differ only by type and vintage; C4C required a certain MPG for eligibility. To determine who is eligible for the program, we assume the trade-in's MPG is exogenously determined and probabilistically assign an MPG rating to the household's oldest vehicle based on EPA MPG estimates for vehicles of that vintage.

To determine which households participate in the program, we must first calculate the subsidy available to each household. Unlike the trade-in's MPG, the MPG of the purchased vehicle is clearly endogenous to the household's decision of whether to participate and which vehicle to buy, based on the credit received. Our model does not differentiate between new vehicles of different MPG ratings and so it is not possible to determine the size of the credit received. The data provided by the NHTSA shows that roughly three-quarters of households received the $4,500 credit and so therefore we set the subsidy received by all households to the maximum credit of $4,500.

As households in the steady-state equilibrium follow a path that fixes the vehicle type portfolio, we determine participation rates separately for households with different portfolio holdings. For eligible households with a single vehicle of vintage a, we allow the household to stay on its equilibrium path or trade-in the vehicle while purchasing a new vehicle. If the household elects to participate, it re-optimizes its future purchasing and holding behavior. Since C4C is a one-time shock to the effective value of trade-ins, the household will thereafter follow a new steady-state equilibrium path.

A household with two vehicles of vintages (a, b), where $b > a$, may choose to stay on its current path or it may elect to trade in its oldest vehicle for either a new car or new truck while receiving the credit. If the household elects to participate, the vintages of its vehicles update and it re-optimizes over its future decisions. As with single-vehicle households, the one-time nature of C4C ensures that the household will follow a (potentially different) steady-state equilibrium path once the participation decision is made.

We simulate the impact of the C4C policy as implemented by Congress and find that 89.6% of trade-ins were trucks, quite close to the 85% reported by NHTSA. Transactions in which trucks are traded-in for cars account for 53.8% of all transactions in the model. This amount is almost equivalent to the 51.8% from the actual program. These results give us additional confidence that our model successfully captures the relevant incentives, particularly as we did not use any information from the C4C program to pin down the parameters of the model.

As our model allows us to simulate any tied-scrappage program, we present results below with respect to a number of variations on the C4C policy design in additional to the program as implemented. We focus on variations that change the incentives to switch between vehicle types, and allow the rebate from trading in a truck for a car to vary between \$2,500, \$4,500, and \$6,500. We consider only policies where switching to a car is at least as rewarding as the alternative.[8] Throughout our analyses, we maintain the \$3 billion subsidy budget appropriated by Congress.[9]

We begin our analysis by exploring the behavior of households under different program designs. Table 5 reports participation rates. The top number of each entry reports the percentage of U.S. households electing to participate, and the

---

[8]As we do not model the details of the new vehicle market, we are unable to explore alternative specifications of tiered subsidies based on MPG improvements.

[9]We do find that more consumers wish to participate in the program than can be accommodated by this budget. For example, under the C4C policy, we find that 4.8% of all households would have opted to participate in the program, which is an order of magnitude larger than the 0.58% of U.S. households who were able to participate. Given the speed as which the appropriation was spent, this result is reasonable. We allow households to participate with uniform probability.

bottom number reports the share of participating households which would not have purchased a vehicle in the year of the policy without the subsidy – in other words, the share of marginal consumers as opposed to inframarginal consumers which would have purchased a vehicle without the subsidy. Both the overall and marginal participation rates increase both in the subsidy amount and in the looseness of trade-in MPG requirement as more households become eligible. However, the highest marginal participation rates occur when the subsidy rate for truck-to-car transactions is higher than the the subsidy rate for other transactions.

TABLE 5.
Overall and Marginal Participation in Cash for Clunkers and Alternative Policies

Subsidy Amount		Trade-in MPG Requirement				
Truck → Car	All Others	16	17	18	19	20
\$2,500	\$2,500	0.06	0.08	0.11	0.14	0.18
		29.7	35.7	36.0	38.5	37.2
\$4,500	\$2,500	0.26	0.33	0.43	0.51	0.62
		51.4	53.0	53.4	54.5	52.6
\$4,500	\$4,500	0.55	0.74	**0.95**	1.14	1.38
		38.3	38.9	**40.9**	42.3	44.0
\$6,500	\$2,500	1.58	2.17	2.73	3.2	3.66
		62.5	62.0	62.3	62.7	62.4
\$6,500	\$4,500	1.87	2.57	3.24	3.82	4.41
		56.8	56.4	57.1	57.7	58.0
\$6,500	\$6,500	2.3	3.15	4.04	4.85	5.8
		55.0	54.4	55.1	55.3	55.6

Notes: For each row-cell combination, the top number is the percentage of households willing to trade-in a vehicle under the policy. The bottom number is the percentage of trade-ins which would not have happened without the policy – in other words the share of new vehicle purchases which are marginal. The Cash for Clunkers policy is in **bold**.

We detail this substitution behavior in Table 6, which reports the percentage of participants who trade-in a a truck for a car. When the subsidy rate is low, less than 20% of the transactions involve a change in vehicle type. When the subsidy for truck-to-car transactions is raised above other transaction types, the share of these transactions increases to as high as 97%.

TABLE 6.
Percentage of Participants Substituting from Trucks to Cars

Subsidy Amount		Trade-in MPG Requirement				
Truck $\rightarrow$ Car	All Others	16	17	18	19	20
\$2,500	\$2,500	12.5	14.3	15.3	18.5	18.9
\$4,500	\$2,500	78.2	78.2	78.2	78.4	76.4
\$4,500	\$4,500	36.3	34.8	**35.2**	35.1	34.2
\$6,500	\$2,500	96.5	96.7	96.6	96.6	96.0
\$6,500	\$4,500	81.4	81.4	81.0	80.8	79.6
\$6,500	\$6,500	65.1	65.5	64.2	62.7	59.8

Notes: The Cash for Clunkers policy is in **bold**.

Short-Term and Long-Term Stimulus

To investigate the extent to which C4C was effective at stimulating spending in the automotive sector, we compare spending on new vehicle purchases in 2009 in our steady-state equilibrium to spending in 2009 under the C4C program rules. The results are presented in Table 7. C4C generated \$8.70 billion in additional consumer expenditures over what would have occurred in 2009. Given the \$3 billion budget, we conclude that a marginal dollar spent on C4C generates \$2.9 in additional spending. The alternative designs explored in Table 7 reveal that increasing the size of the subsidy results in a smaller stimulus effect, even though the percentage of purchases which are marginal increases per Table 5, due to the decreased number of transactions which can be subsidized.

TABLE 7.
Additional Spending in 2009 Under Cash For Clunkers and Alternative Policies ($ billion)

Subsidy Amount		Trade-in MPG Requirement				
Truck → Car	All Others	16	17	18	19	20
$2,500	$2,500	10.1	*13.6*	12.83	11.42	10.4
$4,500	$2,500	7.29	7.38	7.51	7.65	7.84
$4,500	$4,500	8.59	8.78	**8.70**	8.40	8.12
$6,500	$2,500	5.13	5.08	5.10	5.18	5.25
$6,500	$4,500	5.17	5.14	5.17	5.23	5.28
$6,500	$6,500	5.36	5.28	5.24	5.20	5.14

Notes: The table above depicts expenditures in the new vehicle market in 2009 under the policy relative to expenditures without the policy. Expenditures are measured in billions of 2009 dollars. The Cash for Clunkers policy is in **bold**. The optimal policy by this metric is in *italics*.

While these short-term effects may seem considerable, many of the marginal households may have planned to replace their vehicle in the next year or two. In this case, C4C may have acted to simply 'pull-forward' future spending rather than inspire truly new consumption in the vehicle market. To examine this possibility, we take advantage of the dynamic nature of our model, which allows households which participate in C4C to re-optimize their future actions in the primary and secondary vehicle markets. As above, we compare spending in 2009 and each following year under the steady-state equilibrium to the time-series of spending under the C4C program. Figure 4 illustrates the results. In 2009, C4C generates nearly $14,000 per participating household. However, this is significantly offset by a reduction in 2010 of nearly $5,000 per household. Indeed, relative to a world in which C4C did not occur, participating households reduced their spending on new vehicles until 2017, when households begin to replace the vehicles they purchased under C4C *en masse*.

We summarize these offsetting effects in Table 8 by computing the discounted sum of additional spending under C4C assuming an annual discount rate of 5%.

FIGURE 4.
The Effect of Cash For Clunkers on Future Spending

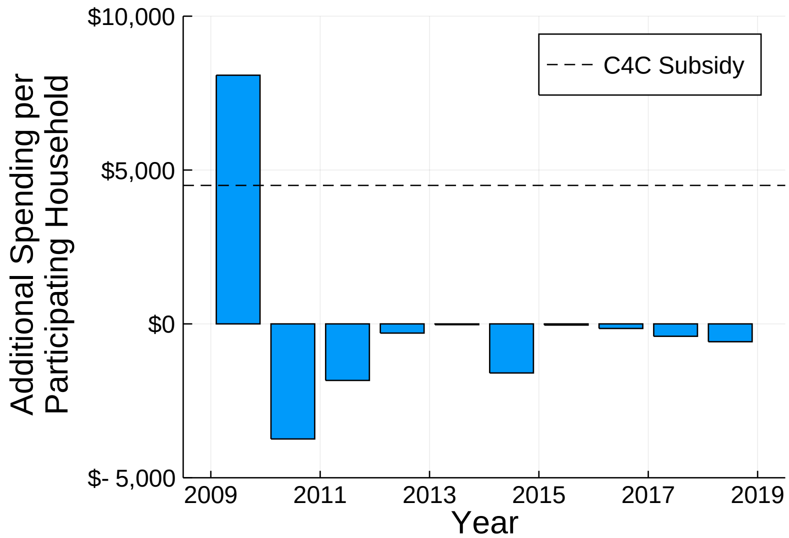

We find that, over the time series, C4C generated only $925 million in additional spending – a marginal stimulus effect of merely $0.31 per dollar spent. Increasing the size of the subsidy to $6,500 results in a negative stimulus effect, as the spending that is pulled-forward is never fully recovered.

Carbon emissions reductions

To calculate the environmental impact of the program, we compare the portfolio holdings of households across the steady-state equilibrium to the holdings after C4C occurs. We aggregate over households to obtain a measure of the net emissions savings, using the differences in the EPA estimated MPG by vehicle type and vintage. We adopt the EPA standard estimates that the average vehicle is driven 11,400 miles per year, and that consumption a gallon of gas leads to 8,887 grams of CO_2 emitted.[10] Finally, we use a EPA estimates of the social

[10]See EPA website

TABLE 8.
Discounted Long-Run Additional Spending Under Cash For Clunkers and Alternative Policies ($ million)

Subsidy Amount		Trade-in MPG Requirement				
Truck → Car	All Others	16	17	18	19	20
$2,500	$2,500	1,145	2,240	*2,361*	2,203	2,040
$4,500	$2,500	757	767	795	972	1,096
$4,500	$4,500	913	997	**925**	870	809
$6,500	$2,500	-786	-824	-790	-711	-654
$6,500	$4,500	-917	-938	-904	-838	-781
$6,500	$6,500	-810	-852	-841	-802	-764

Notes: The table above depicts expenditures in the new vehicle market over the long run under the policy relative to expenditures without the policy. Expenditures are measured in millions of 2009 dollars; the annual discount rate is 5%. The Cash for Clunkers policy is in **bold**. The optimal policy by this metric is in *italics*.

cost of carbon emissions to convert our CO_2 estimates to a present-dollar-valued environmental impact.[11]

The results are reported in Table 9. C4C results in savings of $317 million. Given the budget of $3 billion, our estimate implies that $1 spent on C4C resulted in $0.11 in environmental benefits. In general, making additional vehicles available for trade-in by raising the MPG requirement reduces the environmental impact as fewer emissions are reduced. Increasing the subsidy for truck-to-car trade-ins reduces the impact relative to C4C, even if the subsidy for car-to-car trade-ins is set to $0, because the increase in the subsidy amount reduces the number of households who are able to participate. Conversely, reducing the subsidy increases the benefits as additional consumers are able to participate and a sufficient number of willing consumers exist to exhaust the subsidy budget.

[11]See the 3% discount rate scenario of `https://www.epa.gov/sites/production/files/2016-12/documents/sc_co2_tsd_august_2016.pdf`

TABLE 9.
Environmental Benefits Under Cash For Clunkers and
Alternative Policies

Subsidy Amount		Trade-in MPG Requirement				
Truck → Car	All Others	16	17	18	19	20
\$2,500	\$2,500	*614*	565	497	438	381
\$4,500	\$2,500	418	406	400	391	382
\$4,500	\$4,500	357	342	**326**	307	281
\$6,500	\$2,500	293	286	282	280	277
\$6,500	\$4,500	289	282	278	273	267
\$6,500	\$6,500	265	259	249	241	228

Notes: Environmental benefits are in millions of 2009 dollars under a budget of \$3 billion. The Cash for Clunkers policy is in **bold**. The optimal policy by this metric is in *italics*.

Consumer Surplus and Purchasing Power Inequality

Finally, to calculate the effect of C4C on purchasing power inequality, we must first calculate the consumer surplus generated by the vehicle market under the steady-state equilibrium and under C4C. In our model, V_i represents the discounted sum of utility from income and vehicle use for household i. We transform this to a dollar value for each household by dividing by the marginal utility of net income $(y_i - C + R(\cdot))^{-1}$. We compute this value for each household in 2009 with and without the C4C program to calculate the consumer surplus generated by the policy.

The results are in Table 10. C4C generates a total of \$849 million in consumer surplus – the marginal pass-through of an addition dollar of government spending to consumer surplus is \$0.28. As in the other analyses, in general increasing the subsidy reduces the consumer surplus as fewer households can participate, though this effect is offset by the fact that the households that do participate receive a larger transfer. The highest consumer surplus increases come when the subsidy is reduced to \$2,500.

TABLE 10.
Consumer Surplus Under Cash For Clunkers and Alternative Policies

Subsidy Amount		Trade-in MPG Requirement				
Truck → Car	All Others	16	17	18	19	20
$2,500	$2,500	1,019	991	*1,447*	1,315	1,264
$4,500	$2,500	665	634	721	759	802
$4,500	$4,500	765	764	**849**	931	1,040
$6,500	$2,500	753	739	750	750	751
$6,500	$4,500	764	753	768	778	798
$6,500	$6,500	791	786	806	820	853

Surplus values are in millions of 2009 dollars. The Cash for Clunkers policy is in **bold**. The optimal policy in terms of consumer surplus increases is in *italics*.

These increases are not distributed uniformly for three reasons. First, households vary in the level of subsidy needed to incentivize a new vehicle purchase in the year of the program (and therefore the extent to which the fixed-size credit provides income above their purchase threshold). Second, households vary in their preferences for vehicles, which implies that the surplus from a new vehicle purchase will far as well. Finally, households vary in income and therefore in the marginal utility of income. Figure 5 illustrates the distribution of consumer surplus over participating households under C4C. Though the credit is $4,500, more than half of the consumers receive surplus of less than $2,000.

This uneven distribution of consumer surplus implies that, at least within the set of C4C participants, the program may increase inequality. We test this by calculating the proportion of the subsidy budget which is spent on households in the bottom half of the income distribution. The results are reported in Table 11. Under C4C, almost 99% of the subsidy dollars go to the top half of the income distribution. This is in large part due to the need for households to be able to purchase a new vehicle when given the subsidy; very few households below the median can afford to, even with a $4,500 subsidy. Increasing the maximum subsidy

FIGURE 5.
Distribution of Consumer Surplus Under Cash for Clunkers

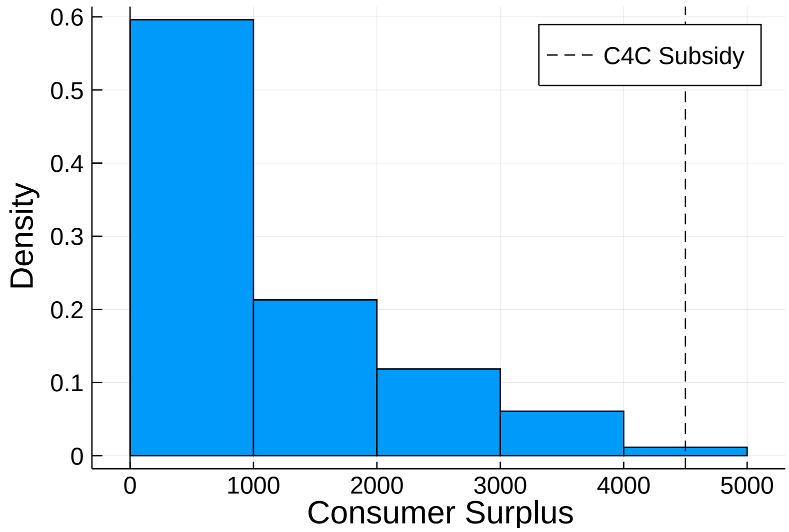

to $6,500 increases the share of subsidy dollars distributed to the bottom half of the income distribution to between 13% - 20% depending on the other implementation details. We thus conclude that these policies are not an effective mechanism for reducing income inequality.

TABLE 11.
Percentage of Subsidy Budget Distributed to Households in the Bottom Half of the Income Distribution

Subsidy Amount		Trade-in MPG Requirement				
Truck → Car	All Others	16	17	18	19	20
$2,500	$2,500	0.0	0.0	0.0	0.0	0.0
$4,500	$2,500	2.0	1.6	2.3	3.1	2.5
$4,500	$4,500	0.9	0.7	**1.1**	1.4	1.1
$6,500	$2,500	19.2	*19.3*	18.5	18.8	18.6
$6,500	$4,500	15.9	16.2	15.6	15.8	15.5
$6,500	$6,500	13.4	13.7	13.9	14.2	14.9

Notes: The Cash for Clunkers policy is in **bold**. The optimal policy by this metric is in *italics*.

Conclusion

Different constituencies can have interests that are diametrically opposed. Competition between politicians for votes and campaign contributions can exacerbate these differences and result in socially sub-optimal policies (Grossman and Helpman, 1996). As a consequence, policymakers generally seek policies that allow them to form winning coalitions by appealing to the interests of different groups (Persson, 1998).

Cash for Clunkers is an example of one such policy. Touted as a way to stimulate flagging domestic industry and reduce carbon emissions and income equality all at the same time, it quickly gained support from multiple groups and easily passed the U.S. House and Senate by large margins (Bunkley, 2009). However, since consumers and vehicles are heterogeneous, the success of any tied-scrappage subsidy program is related to the distribution of consumers and vehicles which are affected.

We analyze Cash for Clunkers and variations on the policy with a dynamic model of vehicle portfolio choice in which consumers differ by income and preferences for different vehicle mixes and vehicles differ by type and age. We find that, per dollar of government spending, the program resulted in roughly \$0.11 in environmental benefits, \$0.28 in consumer surplus, and \$0.31 in net additional consumer spending on vehicles. As these benefits sum to less than \$1, we conclude that the policy did not achieve the trifecta and the government's aims could have been better served by simply allocating the budget of the program directly to households, environmental programs, and vehicle manufacturers according to each group's welfare weights. Insofar that our results indicate that auto manufacturers received the largest gains from the policy, C4C may represent an consequence of industry capture in the sense of Stigler (1971).

Our results are largely driven by the fact that much of the program's resources were spent on inframarginal consumers, and the marginal consumers were largely above the median income. Indeed, we find the program exacerbated income inequality in the sense that only one percent of the subsidy dollars went to individuals in the lower half of the income distribution. Variations of the program could improve each of the outcomes in turn. From the perspective of government spending, a positive rate of return could have been achieved simply by reducing the size of the subsidy from \$4,500 to \$2,500. Such a program would have generated \$0.17 in environmental benefits, \$0.48 in consumer surplus, and \$0.79 in net additional spending, for a total of \$1.44 per dollar of subsidy distributed. However, as such a program would also have increased income inequality, we conclude that it would not achieve the trifecta.

Our results speak to the need to consider both heterogeneity across vehicles and the dynamic behavior of consumers in order to capture the full effect of programs which seek to influence behavior in the primary or secondary vehicle markets. To the extent that policymakers seek to manage vehicle fleets within their jurisdictions, for example by implementing fuel economy standards (Portney et al., 2003; West and Williams, 2005; Anderson et al., 2011) or subsidizing the purchase of electric vehicles (Li et al., 2017; Seo and Shapiro, 2018), our approach may be used to analyze the extent to which such policies will be effective.

CHAPTER III

USING MACHINE LEARNING TO IDENTIFY HETEROGENEITY IN VEHICLE EXCHANGE PROGRAMS

Introduction

The United States is one of the largest contributors to CO_2 emission levels in the world, ranking number two in total emmisions, and number eleven in per capita terms per The World Bank.[1] Personal automobiles are significant contributors to these levels; in fact, the Environmental Protection Agency (EPA) has begun imposing stricter regulations on emission standards for automobiles. Many of the cars and trucks on the road today, however, were not subject to these standards when manufactured.

In 2009, the federal government introduced a temporary program, Cash for Clunkers. The program offered participants a credit of up to $4,500 for scrapping their old vehicle when buying a new vehicle. The goals of this program were twofold. As this program was introduced during the Great Recession, the first was to stimulate a hard-hit automobile industry in an ailing economy. The second was to get older vehicles, particularly trucks, off the street and replace them with newer, fuel-efficient ones. This goal was implemented by making it mandatory that all trade-ins must get a combined MPG rating of less than 18 MPG and that the purchase vehicle must increase fuel efficiency.

---

[1] China ranks higher the the U.S. in total emissions. In terms of per capita, Saudia Arabia is the only country in the top ten with a population above 10 million. The World Bank data can be found here: `https://data.worldbank.org/indicator/en.atm.co2e.pc?most_recent_value_desc=true`

Prior research into this topic includes Quan and Singer (2015) and Miller et al. (2018). Both developed a structural model to quantify the benefits of the program. The results of both papers were that the costs of the program greatly outweighed the benefits. However, none of the prior research into the Cash for Clunkers policy incorporated household heterogeneity when performing their analysis. Typically, structural estimation does not lend itself to this form of analysis due to computational limitations. In this project, I develop a new method to incorporate microdata from the Consumer Expenditure Survey to estimate correlations in observables among households. Using the CEX data, I can estimate the latent variables for household vehicle preference. This approach is an improvement over the previous literature because the estimates are dependent on household observation and their covariates, rather than a parameterization.

I achieve this by borrowing from the machine-learning literature. Machine learning (ML) is poised to be the next big thing within the field of economics. As it stands now, there are but a few papers that cross disciplines between the structural methods used in industrial organization and the algorithms of machine learning. The lack of research into that topic space leaves a relevant gap in the literature. In this paper, I develop and apply an approach to integrate these two methods to evaluate economic policy. I present the implementation of two methods that fall under machine learning that improve upon the current iterations of structural modeling.

These implementations aim to classify a better approach to calibrating structural parameters. In my analyses, I demonstrate that both principle-component analysis and the use of a neural network can be used to compact significantly more dimensions of the household into structural modeling—without increasing computation burdens—than with current methods.

In addition to the providing insights into the effectiveness of the Cash for Clunkers policy, the approach developed in this paper can be applied to a wide variety of structural problems. In the most general sense, the methods developed in this paper can be applied to structural models for which the researcher possess individual level data and unobserved preference parameters.

The remainder of the paper is presented as follows. Section 3.2 discusses in detail the previous economic research into both analysis of the Cash for Clunkers program, as well as the literature of machine learning and economics. Next, I give an overview of the Cash for Clunkers program, including its legislative history and eligibility requirements. In Section 3.4, I present the foundations for the model that I employ. To estimate the model, I incorporate multiple datasets from different sources—which are discussed in Section 3.5. Section 3.6 presents how I use the data to estimate and solve the model. Results are reported in Section 3.7, and Section 3.8 summarizes my findings.

Literature

This project is an intersection of ML techniques and economic policy analysis. In what follows, I discuss the relevant literature related to each of these two areas in turn. The ML literature is a relatively new to economics with most of the work being performed in the last decade. Athey (2018) states that "ML will have a dramatic impact on the field of economics within a short time frame." While there is a vast literature on ML, its impact on economic research is tempered. Much of the related research is yet to come; there are a few publications to reference.

In recent years, ML has been incorporated into economic research with greater frequency. Björkegren and Grissen (2017) uses mobile phone usage data to predict the likelihood of loan repayments. Their result outperforms models that incorporate credit bureau information. Goel et al. (2016) analyze the hotly debated stop-and-frisk policy in New York City and demonstrate that if the policy is limited to the 6% of most likely offenders, the officers could "recover the majority of weapons and mitigate racial disparities in who is stopped." Glaeser et al. (2016) run a tournament in conjunction with the city of Boston on restaurant hygiene prediction. The tournament provided algorithms that prove capable of predicting inspection outcomes on future inspections. These papers illustrate the wide range of possibilities in which ML algorithms can influence economics.

Unlike the previously discussed papers, this project is based on a structural model of consumer choice. Some papers in the literature that deal with this interaction of ML and consumer choice are Ruiz et al. (2017) and Athey et al. (2018). The first of these develops a probabilistic model of shopping data based on consumer choice. The second uses cellphone location data to estimate restaurant effectiveness. They do this by estimating latent variables about both the firm and consumers using a utility maximization framework. This latter paper more closely resembles the work outlined in this proposal.

The other aspect of this project is the Cash for Clunkers program itself. This program has been studied quite extensively in the literature. Many of these papers have investigated the environmental benefits of the program. These include, most notably, Quan and Singer (2015), which adopts a model of vertical product integration from Gavazza et al. (2014), henceforth GLR, to study the U.S. automobile industry. Miller et al. (2018) further develops this model by incorporating the horizontal differentiation between cars and trucks, a vital aspect of the government program. Both researchers find lackluster returns on investment from the Cash for Clunkers program.

However, despite the current research determining that the program was an apparent failure, similar programs have been introduced across the globe. The Romanian government has offered an automobile scrappage program since 2005, increasing the incentive to 6,500 Leu (about $1,600) in 2014, Auto (2015). Butcher (2018) highlights multiple proposals for a London-based program.

The Cash for Clunkers Program

The Cash for Clunkers program, officially the "Car Allowance Rebate System" (and henceforth referred to as C4C), was a nationwide vehicle exchange program implemented by Congress. The program allowed any household to trade in a vehicle and receive a credit toward the purchase of a new vehicle. The program was slated to run from July 1, 2009, to November 1, 2009; however, the initial $1 billion budget ran out within the first month. In response, Congress allocated $2 billion in extra funds. The total $3 billion allocation was spent within two months and the program ended on August 24, 2009.

The official goals of the program were two-fold. As durable goods, automobiles typically experience larger drops in sales during economic downturns. Therefore, the primary goal was to support the U.S. auto industry. The secondary goal, supported by the MPG requirements discussed later in this section, was the environmental benefit.

The size of the credit received by households was either $3,500 or $4,500. Trade-ins were grouped into multiple types: Passenger cars, Category 1 trucks, Category 2 trucks, and Category 3 trucks. Passenger cars typically included sedans (e.g. the Toyota Camry, Ford Focus). Category 1 trucks included minivans and SUVs. The most popular Category 1 trucks traded in were the Ford Explorer and Jeep Grand Cherokee. Category 2 trucks traded in were primarily standard trucks. The most popular trade-ins were the Ford F150 and Chevrolet C1500. Category 3 trucks are larger models, but accounted for only 1% of vehicles traded in under the program.

To be eligible for the credit, households not only needed to trade in a vehicle while purchasing a new vehicle, the new vehicle had to exhibit an increase in fuel efficiency. Traded-in vehicles were required to have a combined MPG of 18 or less, be in working condition and less than 25 years old. To be eligible for the $3,500 credit, households trading in a passenger car had to purchase a new vehicle with at least a 4 MPG improvement. If the new vehicle had a 10 MPG or more improvement, the household would be eligible for the $4,500 credit. If a household was trading in a truck, it only needed to improve the fuel efficiency by 2 MPG to receive the smaller $3,500 credit or only by 5 to receive the larger $4,500 credit.

Given that trucks typically have lower starting MPG and a lower improvement requirement in the context of the program, it was much easier to receive the credit if the household was trading in a truck. This is visible in the summary statistics of the program, as 85% of the vehicles traded in fell into one of the three truck categories, while only 41% of the new vehicles purchased as part of this program were trucks.

Under the Cash for Clunkers program, a variety of vehicles were traded in. Figure 6 provides two different histograms that demonstrate two dimensions for which there was substantial variation in trade ins. These plots emphasis why, when attempting to model the effect of the program, it is so important to focus intently on the heterogeneity in the households that led to such variation among the participants.

FIGURE 6.
Distributions of Vehicle Trade-Ins by Age and Value

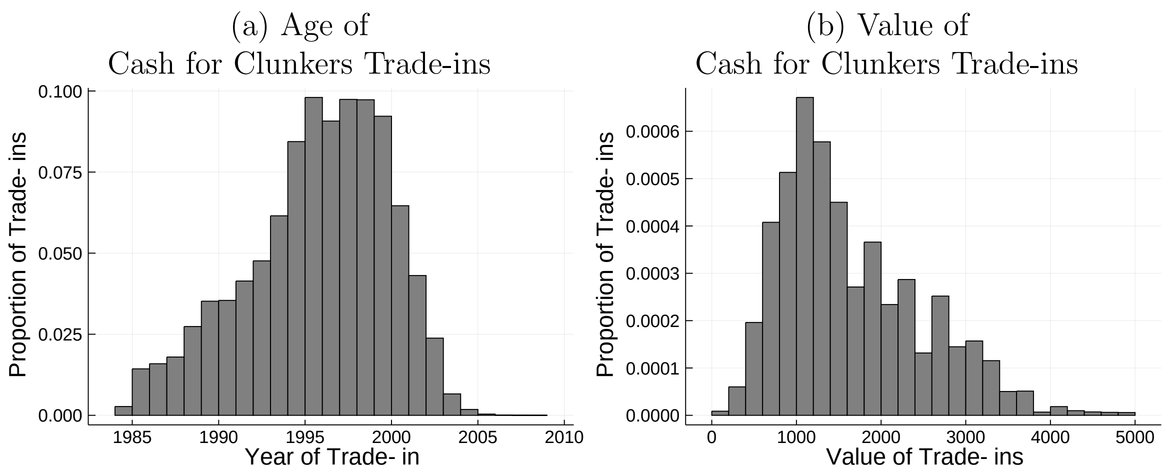

(a): Data from Consumer Expenditure Survey, 1996-2017 Q1 sample. (b): Data from National Highway and Saftey Administration. Age refers to years since model year.

While this study only considers the U.S. automobile market in the context of C4C, many similar programs have been implemented across the globe, particularly during the height of the financial crisis. The largest of these was the program in Germany. The German program had a budget of €5 billion Euros and ran for the entirety of 2009. Any car older than nine years was eligible as part of the German program. Spain, France and the U.K. also implemented vehicle exchange programs.

Model

In this section, I develop a model of discrete consumer choice. The model is developed similarly to GLR, Quan and Singer, and Miller et al., but allows for the introduction of many additional forms of heterogeneity in households. Presentation of the model is made in two parts. The first can be thought of as the "inner" or structural model. The second, or "outer" layer (grounded in ML techniques), determines how heterogeneity is incorporated into the model.

Structural Foundation

Vehicles are differentiated by vintage $a \in \{x \in \mathrm{Z} \mid 0 \leq x < A^\star\}$, where $A^\star$ is the first vintage for which no vehicles exist. The number of vehicles on the road is generally declining in vintage. Imposing $A^\star$ allows the non-consideration of vehicles that make up a disproportionately small fraction of the fleet. Imposing $A^\star$ has the effect of exogenously scrapping any vehicle that reaches this age. There is an exogenously determined stock of vehicles x_a. Vehicles have prices p_a and qualities q_a.

Households, indexed by i, are infinitely lived and optimally choose to own zero, one, or two vehicles. Households are heterogeneous in their preferences. Households are endowed with real income, y. Households have heterogeneous preference (γ_i) for vehicles, as well as a heterogeneous marginal utility (α_i) parameter for additional vehicles. The form that these parameters take is discussed further on in this section.

Households decide how many vehicles (up to two) and of what vintage to purchase those vehicles. Let a be the age of the first vehicle owned by the household, and b be the age of the second. The time-invariant utility that household i earns from their holding of vehicles is

$$u_i(a,b) = \begin{cases} 0 & \text{if } a = b = \emptyset \\ \gamma_i q_a & \text{if } b = \emptyset \\ \gamma_i(q_a + \alpha_i q_b) & \text{otherwise.} \end{cases} \tag{3.1}$$

The first case is a household that does not own any vehicles, while the second corresponds to the utility received by a household with one vehicle. The third is the utility received from a household with three vehicles.

While gaining utility from their holdings of vehicles, households can enter the market and either buy an additional vehicle, sell one of their fleet, or exchange vehicles (both buying and selling). Households incur transaction costs when they enter the market to sell a vehicle. There is a constant cost to all trade-ins (ψ_0), as well as a cost proportional to the price of the vehicle traded in (ψ_1). I use $\Psi(^\bullet)$ to denote the total transaction cost associated with trading in a given vehicle. Households are limited to one vehicle transaction per period. Now I definite the transition dynamics that govern how households move from their portfolio holdings in one period to the next. Let a be the age of the household's newest vehicle, and b the age of their second vehicle, if owned. Let $R(a', b'; a, b)$ be the cost associated with beginning a period with (a, b) and moving into the next period with (a', b'). There are three states for which a household can begin a period (zero, one, or two vehicles.) Households can either stay in their state with no action, transition up or down a state (either if they own one vehicle), or stay in their state by entering the market exchanging vehicles (not an option for zero vehicle households). This results in nine distinct values for for $R(^\bullet)$. However, there are 14 different relations between a, b, a', b' that determine which of the nine values that the households incur. I define $R(^\bullet)$ piecewise

$$
R(a', b'; a, b) = \begin{cases}
0 & \text{if } a = b = a' = b' = \emptyset \\
0 & \text{if } a' = a + 1, b' = b = \emptyset \\
0 & \text{if } a' = a + 1, b' = b + 1 \\
\Psi(p_a + 1) & \text{if } a \geq 0, a' = b' = b = \emptyset \\
\Psi(p_a + 1) & \text{if } a > b \geq 0, a' = b + 1, b' = \emptyset \\
\Psi(p_b + 1) & \text{if } a > b \geq 0, a' = a + 1, b' = \emptyset \\
-p_{a'} & \text{if } a' \geq 0, a = b = b' = \emptyset \\
-p_{b'} & \text{if } a' = a + 1, b' > a', b = \emptyset \\
-p_{a'} & \text{if } a' < b', b' = a + 1, b = \emptyset \\
\Psi(p_a + 1) - p_{a'} & \text{if } a' \neq a + 1, b = b' = \emptyset \\
\Psi(p_a + 1) - p_{a'} & \text{if } a' < b', b' = b + 1 \\
\Psi(p_b + 1) - p_{b'} & \text{if } a' = a + 1, b' > a' \\
\Psi(p_a + 1) - p_{b'} & \text{if } a' = b + 1, b' < a' \\
\Psi(p_b + 1) - p_{a'} & \text{if } a' < b', b' = a + 1.
\end{cases} \tag{3.2}
$$

The first three cases represent a household with zero, one, two vehicles electing not to enter the market and instead let their vehicles age. The fourth, fifth, and sixth case correspond to when a one-vehicle household sells their vehicle, a two-vehicle household sells their newer vehicle, or a two-vehicle household sells their older vehicle, respectively. The next three cases, the seventh through ninth, correspond to a household without any vehicles purchasing, or a one-vehicle household purchasing a vehicle older than their current vehicle, and in the ninth case, a vehicle newer than their current vehicle. The tenth is when a one-vehicle household enters the market and exchanges their vehicles. The final four cases involve a two-vehicle household entering the market and exchanging one of their vehicles. The eleventh is if they exchange their newest vehicle for a different model which is still their newest. The twelfth is if they enter the market with their older vehicle, while maintaining the ordering. The thirteenth is when they sell their newer vehicle and the vehicle that they purchase is old enough that their previously older vehicle is now their newest. The fourteenth and final case is the opposite of case thirteen. The household exchanged their older vehicle for one that becomes their newest, thereby transitioning the original youngest vehicle to the oldest.

In addition to earning utility from vehicles, and paying to enter the market, households pay a cost of ownership. Operating costs (C^1, C^2) depend on the amount of vehicles owned, but not the vintage. With all these components, I have the tools to write down the households' value function

$$V_i(a,b) = u_i(a,b) + \max_{a',b'} \left[\ln\left(y_i - C + R(a',b';a,b)\right) + \beta V_i(a',b')\right]. \tag{3.3}$$

When considering equilibrium in this model, I analyze it in the context of steady state and opt to evaluate the effect of a shock (C4C) at a later time. By considering steady-state behavior, it allows household behavior to be separately characterized by the decision of how many vehicles to purchase. This is because, in the absence of external shocks, no household that has previously decided to own a certain number of vehicles (resulting from optimal behavior), would suddenly elect to change their behavior. Note that this result does not apply to the vintage of those vehicles because, despite the removal of disrupting shocks, the vehicle portfolios of households still age along the steady-state equilibrium path. The value function for a one-vehicle household, suppressing the household notation, becomes

$$\begin{aligned}\tilde{V}(y,\gamma,a) &= \gamma q_a + \ln(y - C^1) \\ &\quad + \beta \max\left\{\tilde{V}(y,a+1),\ \max_{a'} \tilde{V}(y,a') + \ln\left(\frac{y - p_{a'} + \Psi(p_{a+1}) - C^1}{y - C^1}\right)\right\}.\end{aligned} \tag{3.4}$$

Here the household receives their utility from their vehicle and from their income endowment net operating costs. In the second line, they make the optimal decision between allowing their vehicle to age and entering the market to replace it.

Next, I consider the two-vehicle households. Quan and Singer argues that when transaction costs are sufficiently high, two-vehicle households will default into making within-household exchanges. This allows them to avoid a substantial amount of transaction costs by naturally allowing their newer vehicle to age into their older vehicle when they enter the market to replace an aging vehicle. This corresponds to the final case from the transaction cost dynamics. The value function of a two-vehicle household is represented by

$$\begin{aligned}\tilde{V}(y,\gamma,\alpha,a,b) = \gamma(q_a + \alpha q_b) + \ln(y - C^2) + \beta \max\Big\{&\tilde{V}(y,\gamma,\alpha,a+1,b+1), \\ &\max_{a'} \tilde{V}(y,\gamma,\alpha,a',a+1) + \ln\left(\frac{y - p_{a'} + \Psi(p_{b+1}) - C^2}{y - C^2}\right)\Big\}.\end{aligned} \tag{3.5}$$

Here, if the household opts to enter the market, they do so by selling their oldest vehicle and purchasing a vehicle which then becomes their newest, while allowing the vehicle that they keep to age. When the value functions are written in this form, they exhibit a recursive behavior which allows for an analytical solution in the household's decisions of the age of vehicles to purchase. Consider the household that owns a new vehicle. That household will let that vehicle age until they determine that the quality has dropped far enough to a point where they would like to replace it. When they enter the market, what age of vehicle would they replace it with? In the steady-state equilibrium path, they will replace it with a vehicle of the age that is equal to the purchase age of their original vehicle. So if I can determine the age of the vehicle the household will buy, which I represent by a, the only remaining unknown is how long households choose to keep their vehicles, represented by n. If both quantities are known, it is enough to completely characterize the steady-state behavior of the household. In practice, it is required to calculate every combination of a, n and identify the one with the highest value. Below is the closed form steady-state value of a one-vehicle household for a given combination of a, n[2]

$$V(y, \gamma, a, n) = \frac{\gamma q_a + \ln(y - p_a + \Psi(p_n) - C^1) + \sum_{j=1}^{n-a} \beta^j [\gamma q_{a+j} + \ln(y - C^1)]}{1 - \beta^{n-a+1}}. \quad (3.6)$$

From the closed form solutions of the value function, it is possible to compute decision rules, $\{a^\star(y_i, \gamma_i, \alpha_i)$, $n^\star(y_i, \gamma_i, \alpha_i)$, and $m^\star(y_i, \gamma_i, \alpha_i)$ that determine given household characteristics the optimally chosen a, n, and m, where, m is the spread in ownership of the two vehicles in a multi-vehicles household's fleet.

An equilibrium in this model is a collection of decision rules $\{a^\star(y_i, \gamma_i, \alpha_i)$, $n^\star(y_i, \gamma_i, \alpha_i)$, $m^\star(y_i, \gamma_i, \alpha_i)\}$ and prices p_a $\forall$ $a \in \{0, 1, \ldots, T^\star\}$ such that households are maximizing and markets clear

[2]Only the one vehicle case is presented here. Information on two-vehicle households can be found in Appendix A.

$$\int h_a^i(y, \gamma, \alpha) d(y, \gamma, \alpha) = x_a^i \; \forall a \in \{0, 1, \ldots, A^\star - 1\}, \tag{3.7}$$

where $h_a^i(y, \gamma, \alpha)$ is the steady-state distribution of owners of vehicle type i with vintage a. Simply put, the vehicle-ownership decisions for each vintage—of all households—aggregates to the observed exogenous stock of that vintage.

In the remainder of this section, I describe three different methods for implementing household heterogeneity.

Methods of Moments

Previous studies of this market have used method of moments to encapsulate consumer heterogeneity. See GLR, Quan and Singer (2015) and Miller et al. (2018). This requires assuming a parameter distribution for the heterogeneous parameters and then using optimization techniques to solve for the distributions of those parameters. In this section, that is the procedure that I implement. There are three forms of heterogeneity: income, y; the preference for vehicles, γ; and the marginal utility of an addition vehicle, α. I take y to be income drawn from my dataset on consumers.

The preference parameters γ is distributed Log-Normal with mean 1. This restriction is a normalization; as if it was not the case, it would be impossible to distinguish the mean of the γ distribution from the mean of the quality parameters q. The Log-Normal distribution is defined on the interval [0, ∞]. This property can be interpreted that the preference for vehicle ownership for any household cannot be negative.

The marginal utility parameter α, is assumed to be distributed according to a Beta distribution. Notably, the Beta distribution is defined on the interval [0, 1], meaning that the marginal utilities are within this interval. This decision is made based on economic theory because we expect that the utility gained from additional units of the same good should be positive, but less than the first. Both parameters of the Beta distribution are free to be set by the optimization algorithm.

The distribution of the moments method for implementing heterogeneity can be summarized by

$$\gamma \sim \text{Lognormal}(1/2 - \sigma_\gamma^2/2, \sigma_\gamma), \ \alpha \sim \text{Beta}(\beta_1, \beta_2) \tag{3.8}$$

The shortcoming of the above method is that the draws from income, preference, and marginal utility are completely independent from each another. This is not a how we expect the world to operate, nor is it a restriction made from lack of data. The reason that household heterogeneity has a limited scope is due to computational limitations. In practice, optimizing over as few as the three parameters described above can take hours or days. In what follows, I implement two distinct ML techniques to reduce the computation burden from including additional dimensions of heterogeneity.

Principle Components Analysis

If computational power was not a restriction, one might hope to run a regression of the form

$$\gamma = X'\beta_\gamma + \epsilon \tag{3.9}$$

to determine the γ values. I focus on γ rather than α but they are analogous. Where the X matrix was a collection of data that the econometrician believed to depend on γ; however, there are issues with this proposed approach. The first is that the γ values are not observed. This means that solving for the values of β would require optimizing over a $\beta_\gamma + \beta_\alpha$ size state-space, where the second set comes from the analogous α regression. This is not computationally feasible if you were to consider all possible observables that influence γ. Even if it were, it would require the existence of at least $\beta_\gamma + \beta_\alpha$ moments with which to optimize.

I use principle components analysis (PCA) to alleviate this issue. Principle component analysis reduces the data into ordered orthogonal components. Components are ordered such that the first component is that which encapsulates the largest amount of variance in the dataset. The second principle component would contain the second largest variance; the intuition being, that using only the first few principle components would capture a substantial majority of the variation in X.

To compute the principle components, one must start with a data matrix X, in which each row is an observation and each column is mean zero. Next, compute the covariance matrix $C = \frac{X'X}{n-1}$, where n is the number of observations. With the covariance matrix, find the eigenvalues of C, $\Lambda = V^{-1}CV$. Λ should be a diagonal matrix where the non-zero elements are the eigenvalues of C.

If the original dataset had K columns, let $g_K = \sum_{k=1}^{K} \Lambda_{k,k}$. Sort Λ such that $\Lambda_{1,1}$ is the largest eigenvalue, and the eigenvalues descend further down the matrix. Finally, for a given threshold level T find the smallest L such that $\frac{g_l}{g_K} = \frac{\sum_{k=1}^{L} \Lambda_{k,k}}{\sum_{k=1}^{K} \Lambda_{k,k}} > T$. The first L eigenvectors of V are the principle components of X.

With the principle components in hand, we can define how they relate to γ, α. Let σ be a 3×1 matrix. Consider the first two principle components of X. I define

$$\gamma = \sigma_1 \frac{X'PCA_1}{\overline{X'PCA_1}} + (1 - \sigma_1) \frac{X'PCA_2}{\overline{X'PCA_2}} \tag{3.10}$$

$$\alpha = \sigma_3 \left[\sigma_2 \frac{X'PCA_1}{\overline{X'PCA_1}} + (1 - \sigma_2) \frac{X'PCA_2}{\overline{X'PCA_2}} \right] \tag{3.11}$$

Here, γ is once again set to be mean 1. σ_1 dictates how much of the first principle component goes into γ while the rest come from the second principle component. For α, the term inside the brackets is mean 1 and σ_2 determine how much of the first principle component goes into α. σ_3 determine the mean of α, which is confined to be between 0 and 1.

Notice that this has a similar form to the distributions of γ and α from the methods of moments. In fact, to solve for σ requires using methods of moments. However, unlike the regression presented at the start of this subsection, this is characterized in full by just three unknown parameters. In contrast to the original method of moment proposal, this definition incorporates any number of demographic variables to which the researcher has access. The restriction of this instance is that the ratio $\frac{\sum_{k=1}^{2} \Lambda_{k,k}}{\sum_{k=1}^{K} \Lambda_{k,k}}$ might not be as large as desired to incapsulate enough of the variation in the data. In what follows, I provide one final alternative that removes this concern.

Neural Network

Previously, I discussed the possibility of running the following regression

$$\gamma = X'\beta_\gamma + \epsilon \tag{3.12}$$

while noting that to do so would require having already obtained γ or require running the structural model within the regression. Here I explore an alternative to obtain values of γ. While γ is unobserved, what if there were some observed data that could act as a proxy for γ? Call this data Y. Mathematically, if the researcher had access to $Y = f(X)$ and a $g(\bullet)$ such that $g(Y) = \gamma$, then they could obtain the values of γ. Note that it is necessary to find a similar set of proxy data and transformation function for α.

It turns out that such as dataset Y does exist. While γ and α are not directly observed, the vehicle ownership decisions are. Specifically, it can be observed if a household purchases zero, one, or two vehicles. It can reasonably be interpreted that a household that owns more vehicles has a higher preference for vehicles. Similarly, a household that falls into the two-vehicle category likely has a higher preference for their second vehicle than those that only own one vehicle.

If vehicle ownership decisions are acting as Y, then what plays the role of $f(\bullet)$? I use a neural network to predict which household characteristics X are more likely to result in different vehicle ownership decisions Y. A neural net is typically used to predict a finite set of outcome variables from a finite set of input variables subject to some loss function. Many different options exist for the complexity of the network, learning rate, and the loss function selected. Without providing a whole textbook worth of information, I present the specifics of the basic neural net that I employ.

The input variables are made up of any amount of household characteristics, while the output is the decision of how many vehicles to own. A standard feature of a neural net is the inclusion of "hidden" layers that are placed between the input and output layers. These allow the model to identify patterns in the data that a one-stage setup might otherwise overlook. I use one of these layers with a classic sigmoid activation function. Because households can only opt into one (and must opt into one) of the three choices of vehicle ownership, I transform the final activations values into probabilities.

I use a cross-entropy loss function, which is used for estimating the efficiency of an estimated probability distribution from its true distribution. The loss function is

$$\sum_{h \in \text{“0”},\text{“1”},\text{“2”}} - P(h)log(\widehat{P(h)}) \qquad (3.13)$$
$$= -P(\text{“0”})log(\widehat{P(\text{“0”})}) - P(\text{“1”})log(\widehat{P(\text{“1”})}) - P(\text{“2”})log(\widehat{P(\text{“2”})})$$

Since the household will have chosen one of the three outcomes, one of $P(h)$ will be equal to 1, while the other two will be equal to 0. The loss for a given observation is just $log(\widehat{P(h)})$ for whichever group (h) to which the household belongs.

After training the model on a hold-out sample, I calculate the estimated probabilities to which group the household belongs. Let p_0 be the probability that the neural net assigns to the household owning zero vehicle, p_1 the probability it assigned to owning one vehicle, and p_2 to two. I then calculate the expected number of vehicles owned

$$E(V) = p_1 + 2 * p_2 \qquad (3.14)$$

and the ratio of the probabilities of owning two vehicles to owning one vehicle

$$\text{Ratio} = \frac{p_2}{p_1}. \qquad (3.15)$$

A higher number of expected vehicles owned implies a higher preference for vehicles, so this expectation relates to γ. A higher value of the ratio implies a higher marginal utility from the second vehicle, relating to α. Keep in mind that the expected number of vehicles owned, while similar, is not the same thing as the preference for vehicles. Instead, the only conclusion is that households with a higher expected number of vehicles should have higher γ. To get the values of γ requires the $g(\bullet)$ that transforms the value of the expectation to a preference parameter.

The values of γ should come from a distribution confined to the positive values, with mean 1 (so that the qualities can be identified). One such distribution is the Log-Normal distribution. What remains is to map from the expectation on number of vehicles to own, into the Log-Normal distribution. To do this for the whole range of possible households would require an analytical form for the distribution of expectations; however, it is not necessary to map the whole distribution, just the household in the model. To do this, assume there at N households, each with a calculated expected value of vehicle ownership, sort so that the highest expectation is first. To proceed, I take N draws from a Log-Normal distribution, sorted in descending order; then map, one-to-one, the household to the draws.

A similar procedure takes place to obtain the values for α. In this case, draw from a Beta distribution and assign the highest draws to the households with the highest value for the ration calculated prior. To determine the best parameters from distributions from which to draw α and γ, the familiar method of moments is implemented. Alternatively, computational requirements can be substantially reduced if the user is comfortable specifying reasonable parameters for the Log-Normal and Beta distributions.

Data

The datasets used to evaluate the model presented in this paper are from a variety of sources. As previously discussed in Section 3.3, data from the

Cash for Clunkers program is provided from the National Highway and Safety Administration. The primary data source used in this chapter is the Consumer Expenditure Survey (CEX). Additionally, I use data from the Kelley Blue Book for analysis of transaction costs.

Consumer Expenditure Survey

The CEX provides data on American's spending habits, containing published data from 1996. The data used in this study is from 1996-2017. The data is gathered by interviewing selected participants in four consecutive quarters. This study uses the first quarter subsample from each report. This procedure ensures the capture of each interviewee once, but not more.

The dataset reports the purchase decision of all item made by participating households. In addition, it creates a record of all vehicle holdings of the households. The vehicle holdings are used in the calibration of the model. This is elaborated upon in Section 3.7. The CEX dataset also includes a wide range of demographic variables for participating households. The demographic observations are used to discern unobserved household preferences for the structural parameters in the second-stage analysis. These primarily consist of dummy variables. The summary statistics for the demographic variables are displayed in Table 12.

National Highway and Traffic Saftey Administration

The National Highway and Traffic Safety Administration (NHTSA) data covers all trade ins and new-vehicle purchases related to the Cash for Clunkers program. A component of this data is vehicle model and type, with data provided on the number of times the vehicle was traded in and purchased (if it was new in 2009) as

TABLE 12.
CEX Demographic Information

	Mean	Std
Income	60046	58690
Adults	1.94	0.91
Children	0.58	1.02
Age[a]	49.6	17.4
Urban	0.94	-
Male[a]	0.47	-
Vehicle Holdings[b]		
No Vehicles	0.14	-
One Vehicle	0.34	-
Two Vehicles or more	0.53	-
Size of Locality		
125K or less	0.13	-
125-330K	0.23	-
.33-1.2M	0.06	-
1.2-4M	0.24	-
4M or more	0.34	-
Region		
Northeast	0.19	-
Midwest	0.23	-
South	0.35	-
West	0.22	-
Eductation[a]		
High-school or less	0.39	-
Bachelor Degree	0.50	-
Graduate Degree	0.11	-
Race[a]		
White	0.83	-
Black	0.11	-
Asian	0.04	-
Hispanic[a]	0.13	-

Data from 2009Q1 CEX. $N = 5,727$.
[a]Refers to the interviewee, rather than the household.
[b]Historical average from 1996-2017.

part of the program, as well as other descriptors of the vehicle. Summary statistics for the program are in Table 1 in Chapter II. This data was used to plot the age of vehicles that were traded in as part of the program in Figure 6.

Kelley Blue Book

The July-September 2009 edition of the Kelley Blue Book (KBB) provides data on the 2009 value of the most popular used vehicles produced from 1994-2008. I use the KBB to compute transaction costs for the model. The KBB reports four different values for each vehicle: the trade-in value for vehicles in fair condition; the trade-in value for vehicles in good condition; the private-property value of vehicles in good condition; and the retail value of vehicles in excellent condition. Summary statistics for the trade-in value for vehicles in good condition and the private property of value of vehicles in good condition are presented in Table 13.

TABLE 13.
Kelley Blue Book Summary Statistics

	Mean	Std
Trade-in Value	2957	2122
Private Property Value	4276	2406

Notes: Data from Kelley Blue Book. N = 1301.

Estimation

In this section, I discuss solving the model. This includes how I calibrate the necessary parameters and what methods I use to solve for the variables that remain. Unless stated otherwise, information pertains to all versions of the model.

Calibration

The main dataset used for calibration is the CEX data, which is used to calibrate the steady-state level of vehicle supply and price in the model. This is done for the stock of vehicles. To do this, first obtain a record of all vintages owned as a percentage of total vehicles on the road for each year in the data set. Then, take the year-over-year percentage change in the occurrence of a given vintage. After this, average the percentage changes with respect to age. Noting that in the steady state, age is the same as vintage, apply this average change of age to the relative stock of each vintage. For use in the simulation, a 5th-order polynomial is fit to the data, this order is used due to nature of the curvature. Figure 7 illustrates the relative stock obtained by applying this process to both cars and trucks.

FIGURE 7.
Distributions of Vehicle Stocks and Prices by Age

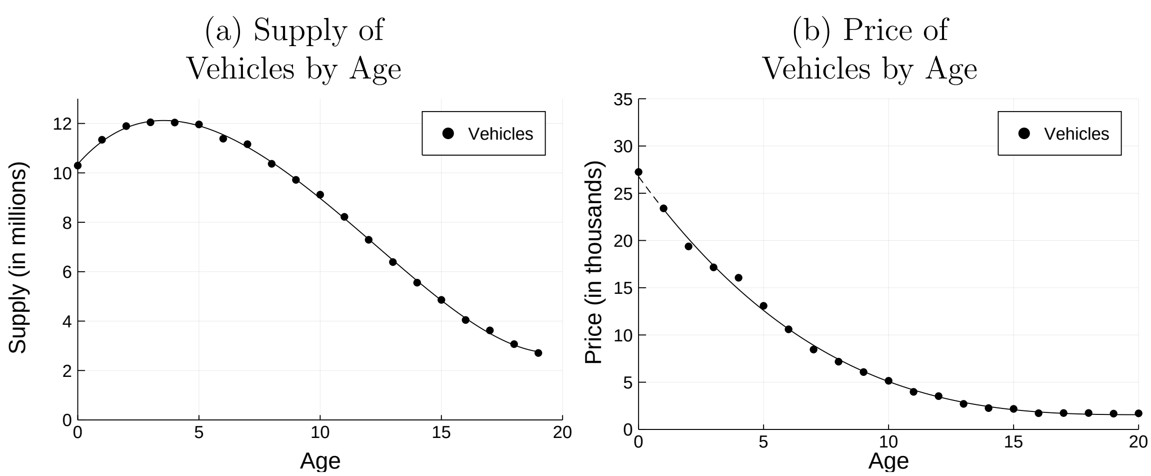

Notes: Data from Consumer Expenditure Survey, 1996-2017 Q1 sample. Points denote sample averages, lines are best-fit 5th-order polynomials. Age refers to years since model year.

With prices, take a similar approach. Parsing the records of vehicles purchased in the previous 12 months reported in the Q1 sample from 1996-2017 of the CEX, obtain the median-price level of used vehicles of every vintage available on the market in each year. New prices for each year are obtained by looking at the median purchase price of new vehicles in that year of that vintage. Next, obtain the relative price of used-vehicles to the new-vehicle price for each year. Using the median of these estimates by age, obtain the average price of a used vehicle of any age relative to a new vehicle. Once again, the steady state allows us to interpret ages as vintages; the price of each vintage is relative to the new-vehicle price. New-vehicle price is matched to the median new-vehicle price in the Q1 sample of the 2009 CEX. As evidenced in Figure 7, there are rarely transactions of 15+ year-old trucks observed; as such, the estimates are non-monotonic. I fit a 5th-order polynomial to used vehicles with the restriction of a negative derivative to the data. The price of new vehicles are set to the point estimate from the CEX data.

Transaction costs in the model are calibrated using data from the KBB. As in Miller et al., I used the difference between the private-property value and the value of vehicles in good condition to model transaction costs. To get a smooth relation, OLS estimates are fit to these data. The results are displayed in Figure 2 in Chapter II. The results show that a linear fit is a good approximation to the difference between the value of a vehicle and what a consumer can expect to receive when trading that vehicle in. This relationship extends to the upper end of the distribution without breaking down.[3]

I set the yearly household discount rate β equal to .95 and the maximum age ($A^{\star}$) of vehicles to 20 years. Operating costs differ by the number of vehicles owned. Operating costs are the sum of maintenance, registration, and fuel costs, and are set at four times the median operating cost from the CEX Q1 2009. The

[3]I experimented with varying transaction costs by vehicle type, but the estimates are essentially identical.

frequencies of all vehicle combinations in the data and associated operating costs are shown in Table 14.[4]

TABLE 14.
Distribution of Households and Ownership Costs Across Vehicle Portfolio Types

Vehicles	Household Frequency	Operating Cost
No Vehicles	13.9%	$0
One Vehicle	33.6%	$1360
Two Vehicles	52.5%	$2800

Notes: Data from Q1 CEX 2009. Operating costs are in 2009 dollars and are annualized from median quarterly measure.

The distribution for the income parameter, y, is drawn from reported income from the 2009 CEX. The stock of cars and trucks, x_a, are from figure 7. The values are scaled so that the total stock is equivalent to the observed stock per household.

$$\sum_{a=0}^{A^\star} x_a = 1.386 = .336 + 2 * .525$$

The frequency of households into certain types provides the primary moment that the data aims to match. However, in all the implementations I described, there were three free parameters to be estimated. This requires at least three moments for them to match. Currently, there is only one because, while there are three household types, by nature they must sum to one, removing one moment. Likewise, the frequency of one vehicle household plus two times the frequency of two vehicle households must sum to the stock of available vehicles. This leave only one moment to match with the free parameters to be estimated.

There are other moments in that data that could be matched. GLR propose the use of three addition moments. These three are described below and the value

[4] I experimented with allowing costs to vary by vintage, but it did not have a significant effect.

from the data is provided for context.

$$\frac{\text{Households with at least one vehicle}}{\text{Households that acquired a vehicle in the last 12 months}} = 2.89$$

$$\frac{\text{Total stock of vehicles}}{\text{Vehicles acquired in the last 12 months}} = 5.30$$

$$\frac{\text{Vehicles acquired in the last 12 months}}{\text{New vehicles acquired in the last 12 months}} = 3.26$$

These three provide the necessary additional information to complete the calibration of the model.

Simulation

The remainder of the variables require solving the model rather than calibration. Estimation of the model requires simulation. The final goal is to match the data on:

1. the optimal decisions of the households result in the correct stock of vehicles on the road, as determined by the calibration of x_a; and
2. the frequency of household profiles are distributed according to Table 14.

I match the data with two optimization loops, referred to as an inner and an outer loop. The inner loop solves for the qualities for a given set of consumers, while the outer loop determines the process that generates the distribution of consumers. Upon completion, the routine solves for qualities, q_a, $\forall\ a$, and for the parameters that govern the distributions of γ and α. Below the procedure is outlined.

1. Guess values for the parameters that govern the distributions of γ and α.
2. Generate consumers according to the distributions of γ and α.
3. Guess qualities $q_a\ \forall\ a \in \{0, A^\star - 1\}$.
4. Update qualities as follows.

a. Calculate optimal a, n, and m for each consumer given qualities.

b. Aggregate consumer vehicle decisions by vintage and type.

c. Increase qualities that do not have adequate representation and decrease others.

5. Repeat step 3 until convergence.
6. Aggregate consumer decisions by household type.
7. Compare the frequency of each type of household to the values in Table 14.
8. Update guess of the parameters that govern the distributions of γ and α.
9. Repeat steps 2-8 until convergence.

In Section 3.4, I presented three different means of evaluating heterogeneity in this model. Referring to the procedure above, the aspect that changes between these different approaches is steps 1 and 2. In the next section I present my results from evaluation of the model.

Results

There are three implementations of the model to consider: The method that uses purely moment analysis; the method that incorporates principle component analysis; and the method that uses a neural net. I begin with a comparison of the results from the methods.

Comparison of Models

Table 15 provides context for how the distributions of γ and α behave in the context of the three models. Notice that the correlation between income and the two distributions is equal to zero in the baseline case. This is a result of the fact that that specification draws γ and α independently.

Each model has its own set of quality estimates. These qualities are not necessarily on an equal scale, due to any number of inputs that might change the

TABLE 15.
Comparison of Machine Learning Models

	Baseline	PCA	Neural Net	Data
Vehicle Preference (γ)				
σ_γ	0.990	-	1.007	
σ_1	-	-0.296	-	
Std(γ)	1.830	0.529	1.870	
Cor(γ, Y)	0	-0.562	0.197	
Marginal Utility (α)				
β_1	0.654	-	0.448	
β_2	1.024	-	1.025	
σ_2	-	0.210	-	
Mean(α)	0.387	0.692	0.300	
Std(α)	0.294	0.241	0.292	
Cor(α, Y)	0	0.676	0.237	
Moments				
Moment 1	3.42	3.56	3.40	2.89
Moment 2	5.54	6.21	5.48	5.30
Moment 3	2.90	2.57	2.89	3.26

Moment 1 = $\frac{\text{Households with at least one vehicle}}{\text{Households that acquired a vehicle in the last 12 months}}$
Moment 2 = $\frac{\text{Total stock of vehicles}}{\text{Vehicles acquired in the last 12 months}}$
Moment 3 = $\frac{\text{Vehicles acquired in the last 12 months}}{\text{New vehicles acquired in the last 12 months}}$
The fraction of variance captured in the PCA model by the first two principle components is 99.9%.
$N = 10,000$

level of the unobserved quality. Instead of presenting the raw numbers, deviations from the mean are presented. The standardized qualities are presented in Figure 8.

All three models exhibit similar behavior in how the qualities of each vehicle differ from the mean. Since the model has been designed to closely fit the distribution of vehicle types by households, I cannot evaluate the effectiveness of the model in that regard; however, there are other ways to determine if the model is capturing the relevant aspects of the data. One way is to examine the rate at which households with different income levels select different types of vehicle ownership. The results from this exercise can be found in Table 16. The model is solved to closely fit the distribution of vehicle types by households, but it also performs quite well in representing the households' decision by income brackets.

TABLE 16.
Household Vehicle Type Portfolios by Income Quartile, CEX
Averages and Predicted Model Values

Income Quartile	Zero Vehicles	One Vehicle	Two Vehicles
	30.5%	46.7%	22.8%
Lowest	45.5%	34.8%	19.8%
	56.0%	35.4%	8.5%
	51.8%	34.2%	14.0%
	10.3%	47.0%	41.7%
Mid-Low	9.7%	41.1%	49.2%
	4.3%	54.2%	41.5%
	5.5%	50.1%	44.4%
	3.7%	27.8%	68.3%
Mid-High	2.6%	31.4%	66.0%
	0.0%	34.6%	65.4%
	1.4%	31.3%	67.3%
	2.8%	14.4%	88.2%
Highest	0.4%	22.7%	76.9%
	0.0%	5.4%	94.6%
	0.2%	16.8%	83.0%

Values from the data are reported first. The second row are from the baseline model. The third row is from the PCA model, and the final row if from the neural net.

FIGURE 8.
Standardized Qualities Estimates

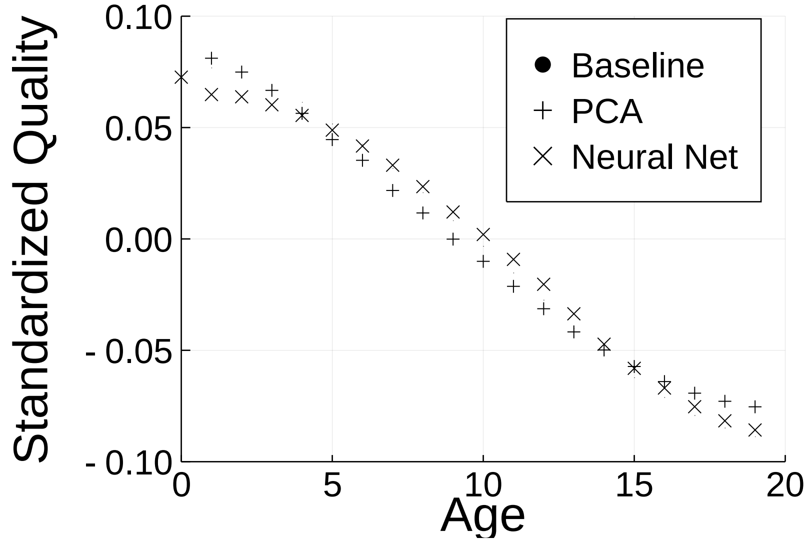

Counterfactuals

While the expanded models do not perform measureably better in the analysis of income brackets, the counterfactuals are the primary reason that they are included. In this subsection, I put the models through a series of counterfactual exercises like that of C4C. First, I provide an oversight of how effective these counterfactual programs can be on aggregate indicators. Due to the improvements incorporating heterogeneity, it is possible to backtrack the distributional effects that these counterfactuals have on different groups of households.

For all policies presented, four different credit levels \$3,500, \$4,500, \$6,000, and \$9,000, are considered. The first two are the two credits offered under the C4C program, and in addition, there are two larger credit amounts. For each credit level, I also impose an MPG restriction on trade-in vehicles. Under C4C, this was 18MPG for most vehicles. The effects if this requirement was removed are presented for comparison. All policies assume the same \$3 billion budget, and that that there is sufficient demand to exhaust it. This is not a baseless assumption, as the actual C4C program had significantly excess demand and the policies that considered consist of either larger credits or less stringent eligibility requirements. Either of these changes should only result in increased demand.

To calculate which households participate in the counterfactual programs, I model the program as a one-time shock to the model. Note that C4C, and all policies considered, are tied subsidies, meaning that it required both the purchase of a new vehicle and the trade in of an old vehicle. This means that any household participating in the program would leave with the same number of vehicles as with which they began. As such, I only need to consider how the policies effect households within a set amount of vehicle ownership.

Households without any vehicles are not eligible because they do not have access to a vehicle to trade-in. The policy shock for single household vehicles is modeled as follows. First, calculate the value to the household of maintaining its current path through the steady-state equilibrium. As an alternative to this, present each eligible household (with a single vehicle meeting the MPG requirement) the option to exchange that vehicle, with a new vehicle discounted by the credit amount. If they opt to do this, calculate the value function associated with keeping their new vehicle until the time their decision rule states that they would sell it in equilibrium. All future vehicles they purchase would be consistent with their steady-state equilibrium path. To determine if they opt-in to the program, compare these two value function calculations.

For multiple vehicle households, first check for eligibility contingent on whether their oldest vehicle meets the MPG restriction. Then, like single-vehicle households, calculate the value function associated with remaining on their equilibrium path. To determine the value of participation, calculate the value function contingent on them trading in their oldest vehicle for a new vehicle, with cost discounted by the size of the credit. While I do assume that households keep their new vehicle purchased under the program until their equilibrium decision rule dictates, I allow households to keep what becomes their older vehicle for a different number of periods than they would have without the program. By allowing them to retain their vehicle for a longer (or shorter) period of time, it allows households to maintain the gap between the ages of the two vehicles in their fleet, while still returning to their equilibrium path in the long run. In order to determine how long households keep this vehicle, the households are allow them to optimally choose the best option given their characteristics. Finally, compare the value of participating with optimal behavior, with that of maintaining their equilibrium path.

This first statistics illustrated below in Table 17 show what portion of eligible households choose to participate under the different schemes. Unsurprisingly, as the size of the credit increases, the number of willing participants increases. It appears that as MPG requirement is waived, a lower percentage of eligible households elect to participate. The reason for this is that the program only appeals to households with older vehicles which typically have a lower MPG. Recall that the traded-in vehicles are scrapped, meaning any household that would be trading in a vehicle worth more than the size of the credit would never be rational. Thus, opening the policy to households with all ranges of MPG requirements would reduce the fraction of willing participants.

TABLE 17.
Household Participation Rates from Alternative Policies

Credit	Restiction	Baseline	PCA	Neural Net
$3,500	18 MPG	2.50	1.39	2.69
$3,500	No Limit	2.22	2.07	2.37
$4,500	18 MPG	4.63	8.10	4.66
$4,500	No Limit	4.21	7.45	4.34
$6,000	18 MPG	9.74	20.91	10.33
$6,000	No Limit	9.52	19.60	9.58
$9,000	18 MPG	37.37	55.58	36.06
$9,000	No Limit	35.37	53.92	33.29

Notes: The table above depicts the fraction of eligible households that participate in the proposed programs.

The primary goal of C4C was economic stimulus, so the next three tables presented are indended to summarize how effective the counterfactuals would have been at promoting it. The first metric is what fraction of the participants would not have purchased a vehicle in 2009. This can be interpreted as how many purchases would not have occurred in the absence of the C4C program. These values can be seen in Table 18. As the size of the credit increases, a larger faction of participants would not have purchased a vehicle in 2009 in the absence of the program. This meaning that the program would have been more successful in inducing new purchases with a larger credit.

TABLE 18.
Fraction of Purchases New to 2009 from Alternative Policies

Credit	Restriction	Baseline	PCA	Neural Net
$3,500	18 MPG	48.3	33.3	41.5
$3,500	No Limit	45.8	26.7	41.2
$4,500	18 MPG	44.4	48.6	46.9
$4,500	No Limit	51.9	44.4	48.8
$6,000	18 MPG	48.9	58.7	48.2
$6,000	No Limit	52.7	58.2	55.4
$9,000	18 MPG	62.2	66.3	60.3
$9,000	No Limit	61.8	66.5	62.3

Notes: The table above depicts the fraction participating households that would not have purchase a vehicle in 2009 in the absence of the proposed programs.

The number of incremental purchases does not tell the whole story concerning the effectiveness of these proposals toward economic stimulus. As the credit size increases, there are fewer households that can participate for the fixed $3 billion budget. Table 19 illustrates how spending increased in 2009 as a result of the counterfactuals. The total effect is diminished as the size of the credit increases, due to reduced participation. The removal of the MPG requirement seems to have no significant effect on the 2009 stimulus.

TABLE 19.
Increase in 2009 Output from Alternative Policies

Credit	Restriction	Baseline	PCA	Neural Net
$3,500	18 MPG	13,697	14,840	12,460
$3,500	No Limit	13,225	13,159	12,536
$4,500	18 MPG	10,908	12,850	10,863
$4,500	No Limit	11,727	12,292	11,190
$6,000	18 MPG	8,896	10,713	8,681
$6,000	No Limit	9,408	10,426	9,434
$9,000	18 MPG	7,617	7,348	7,092
$9,000	No Limit	7,588	7,451	7,394

Notes: The table above depicts the change in 2009 output as a result of the proposed programs. Values are in millions of 2009 dollars.

While the previous two tables aimed to summarize the stimulus effects in 2009, the model developed in this paper can evaluate market effects beyond only the year in which the policy is implemented. This is done by looking at the residual effect of the decisions by the households that opted to participate. The overall finding in Table 20 is that the long-run benefit of the program is significantly diminished from the initial effect exhibited in 2009. This is because, while not all participating households would have otherwise purchased in 2009, they likely would have purchased within the next few years. In long run the only effect this program has is to accelerate purchases by a few years.

The other stated goal of the program was to increase environmental quality. In order to calculate the environmental benefits of the program, I considered CO_2 offsets. To calculate the amount of CO_2 removed from the atmosphere, I calculated the difference in emissions from the household's fleet of vehicles if they were to participate in the program, compared to if they did not participate. To do this, I assume that all vehicles are driven an average number of miles; therefore, the results are solely the difference in fuel efficiency of the new cars compared to their replacements. To convert the reduction in emission into a dollar value, I use the social cost of carbon estimates from the EPA during the Obama Administration.

TABLE 20.
Discounted Increase in Long-Run Output from Alternative Policies

Credit	Restriction	Baseline	PCA	Neural Net
$3,500	18 MPG	3,268	2,576	3,131
$3,500	No Limit	3,044	2,439	3,054
$4,500	18 MPG	2,328	2,476	2,622
$4,500	No Limit	2,711	2,982	2,751
$6,000	18 MPG	2,204	3,447	2,012
$6,000	No Limit	2,560	3,490	2,450
$9,000	18 MPG	3,201	3,843	2,399
$9,000	No Limit	3,133	2,999	2,790

Notes: The table above depicts the discounted change in long-run output as a result of the proposed programs. Values are in millions of 2009 dollars.

$C0_2$ offsets can be seen in Table 21. Similar to the results from the stimulus, larger credits result in smaller benefits. The removal of the MPG restriction also reduces the environmental benefit of the program.

TABLE 21.
CO_2 Offsets from Alternative Policies

Credit	Restricton	Baseline	PCA	Neural Net
$3,500	18 MPG	262	235	287
$3,500	No Limit	175	137	185
$4,500	18 MPG	187	186	209
$4,500	No Limit	132	131	139
$6,000	18 MPG	134	142	141
$6,000	No Limit	89	100	99
$9,000	18 MPG	78	85	80
$9,000	No Limit	52	56	54

Notes: The table above depicts the total CO_2 offsets of the proposed programs. Values are in millions of 2009 dollars.

By modifying the baseline model to include household demographics, it allows me to analyze the effect that these counterfactual exercises have on varied demographic groups. For the remainder of the paper I present just a few of the analyses that are possible with this modification. The first is what percent of participants are male, seen in Table 22. As the size of the credit increased, fewer males participate in the program. One reason for this may be that, with smaller credits, only the oldest vehicles are traded-in as part of the program—which are more likely to be owned by males. A similar explaination would explain why, when the MPG restriction is removed, fewer males participate.

TABLE 22.
Percent of Male Participants from Alternative Policies

Credit	Restriction	PCA	Neural Net
$3,500	18 MPG	66.67	58.46
$3,500	No Limit	53.33	56.15
$4,500	18 MPG	51.43	56.64
$4,500	No Limit	47.22	54.07
$6,000	18 MPG	53.26	52.57
$6,000	No Limit	48.43	52.87
$9,000	18 MPG	54.22	49.22
$9,000	No Limit	49.18	49.00

Notes: The table above depicts the percent of male participants in the proposed programs.

The final proof of concept that I demonstrate is how individual regions of the country are impacted by the proposed policies. For this analysis, I present the results from the neural net model in Table 23. The conclusion is that as the size of the credit increases, the participants shift from the Northeast (the richest region of the country) to the South (the poorest) and, to a lesser degree, to the West. This type of analysis could be incredibly useful if policy makers wished to target specific areas of the country, or aimed policy toward specific demographic groups.

Conclusion

Recently, automobile scrappage subsidies—such as Cash for Clunkers—have fallen out of favor. This is most likely due to these programs' relative ineffectiveness, which is the consensus of the literature. Despite the noble goal of acting as both an economic stimulus and a method to promote a lower carbon

TABLE 23.
Region of the Country of Program Participants (Neural Net) from Alternative Policies

		Region			
Credit	Restriction	Northeast	Midwest	South	West
$3,500	18 MPG	36.92	20.00	27.69	15.38
$3,500	No Limit	31.55	14.44	34.76	19.25
$4,500	18 MPG	33.63	16.81	30.97	18.58
$4,500	No Limit	30.23	14.53	35.17	20.06
$6,000	18 MPG	30.43	13.04	36.36	20.16
$6,000	No Limit	27.55	14.36	36.95	21.15
$9,000	18 MPG	22.33	18.78	37.78	21.11
$9,000	No Limit	21.72	19.36	37.05	21.87

Notes: The table above depicts the percent of participants that live in each region of the country in the proposed programs.

footprint, the application tends to fall short on both fronts when compared to the associated high cost.

In this paper, I develop a heterogeneous consumer model in which households optimize over multiple vertically differentiated products. The model is applied to the U.S. car market using a variety of data sources. Importantly, data from the Cash for Clunkers program is not used in estimating the model. With the results, I perform a counterfactual meant to replicate the effects of the Cash for Clunkers program.

I expand on the current literature by developing a new method for modeling large amounts of household heterogeneity. Comparing my innovation on modeling heterogeneity to the previous works evaluating the program at an aggregate level, I can closely match their qualitative results. By expanding on the current literature, I not only estimate aggregate effects but also distributional effects. I find that as the conditions of the program changes, it causes large shifts in what regions of the country and which demographic groups are affected.

CHAPTER IV

NEW PRODUCTS AND INNOVATION IN THE SMARTPHONE INDUSTRY

Introduction

Since the introduction of the iPhone by Apple on June 29th, 2007, the smartphone market as experienced astronomical growth. Since the original iPhone launched, the smartphone industry has expanded by over four times in terms of annual units sold, growing to an industry with over $50 billion in yearly sales across many competing firms.

These devices have become a daily part of our lives. Three of every four American adults own a smartphone, of which over 94% reported to carry their smartphone with them frequently, according to the PEW Research Center. Over this timespan, the phones in our pockets have gotten more and more advanced. For example, consider that AI assistance in the form of Apple's Siri and Google's assistant, is a standard feature of almost every phone sold today after being introduced only seven years ago with the iPhone 4s. Consumers have access to professional level camera replacements in their pockets, and of course the myriad of apps free available for download.

In a market so large and ubiquitous, surplus is contested for by both consumers and firms. This constant struggle is exemplified by the rapid pace of technological improvement exhibited in the market. In this paper, I model the initial growth in the smartphone market and separately identify the value of the rise in accessibility, and the value of technological innovation. My findings determine that considering these two effects in tandem, the consumers have benefited by $3.5 billion per year or $11.50 per individual per year.

The makeup of the smartphone industry has evolved over years. The market originated out of the exclusive two-year contract agreements, a remnant of the "feature phone" era. During this time there were every few options of phones to pick from, with the market confined to physical space limitations of the dealership. The term feature phone originated as a concept of the plans themselves. As the plan itself was the only differentiated aspect of the bundled purchase, a phone would be included as part the plan as an additional feature. The modern smartphones are anything but homogeneous. With extensive coverage chronicling every imagine adaption.[1]

Today, consumers have many more options for obtaining access to a smartphone and the number of phones on the market as exploded. While there are still dominant brands (Apple, Samsung) the fringe firms provide and an increasing amount of value to the market. Brynjolfsson et al. (2003) show that capturing the effect of these fringe firms is vital to obtaining welfare estimates. Greenstein and Ramey (1998) and Weiss (2003) examine different industrial makeups in vertically integrated markets to determine when firms have incentive to innovate.

The introduction of new products has a long a recent history. In his study of minivans, Petrin (2002) finds that the introduction of the vehicle added \$500 million in yearly welfare. The analysis presented in this paper expands on the work done by Petrin. A paper by Crawford et al. (2015) has looked at welfare effects in a market of substantial size, the cable television market. The smartphone, unlike the minivan or cable television, has undergone vast technological change since its introduction. The technological advance of the smartphone has been on a broad scale, ranging across many attributes (camera, processing power, resolution, etc.) In studying technological growth, Goettler and Gordon (2011) allow firms to make strategic decisions about when to release new models when studying the microprocessor market. However, the authors of this paper boil down microprocessors to a quality variable in one dimension, and only must consider the strategic nature of a two-firm market. This approach does not lend itself to the makeup of the smartphone market. One reason is the top firms release their

[1] In 2010Q1 there were 74 options while in 2014Q3 consumers had 163 models to choose from.

flagship devices at the same time every year, implying that that strategic decision in the smartphone market is degenerate. Another is that the smartphones are differentiated on more attributes than just speed. Attempting to include all the different attribute dimensions in a Goettler and Gordon (2011) style model would soon become untractable.

Instead, I use a modified version of the model from Berry et al. (1995). I can track technological innovation over the product cycle and obtain welfare estimates in the smartphone market. Additionally, I can diagnose the value to consumers of firm competition and of brand loyalty to firms. Using the later to preform welfare analysis of potential mergers.

What makes this market an important topic of study is primarily its size and dramatic growth rate. Today, almost every American adult owns a smartphone (and uses it daily), a device barely a decade old. New devices are repurchased, on average every 33 months and can cost up to $1000. Over time the prices of these devices have continued to rise. Even when new products are introduced, the previous model only see a small, if any, drop in price. This makes the smartphone market is one of the largest in world. Understanding how this market operates is not only necessary to accurately regulate the market but can useful to guide a understanding of the increasingly large and complex digital markets at the center of today's economy.

In the next section, I lay out the specifics of the utility specification. In section three, I discuss the datasets used in estimation. In section four, I present results.

Discrete Choice

I adopt a random coefficients model from Berry et al. (1995). In models such as these, individuals receive utility from their purchase decisions. This specification can be written in the fashion presented below

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \epsilon_{ijt}. \tag{4.1}$$

Here, i indexes individuals and j indexes products on the market, and t time. δ has been separated out as it contains the terms that do not depend on i, that means that this term (δ) captures aspects of the product that affect all consumers equally. The term μ is individual specific, capturing individuals' heterogeneous preferences. ϵ_{ij} is an individual specific error term which is assumed to be i.i.d. across consumers. The term δ, common to all consumers, can be further disaggregated

$$\delta_{jt} = \alpha p_{jt} + \sum_{k=1}^{K} \beta_k x_{jtk} + \xi_{jt} \tag{4.2}$$

Here k indexes product attributes, the product attributes (p, x) are assumed to enter linearly. α is the marginal utility of income, and β, the marginal utility of one addition unit of associated attribute. ξ is utility earned from aspects of the product that are unobserved by the econometrician (but observed by the consumer). Individuals select the product that gives them the most utility. Assume that the error term ϵ follow a type 1 extreme value distribution. This assumption allows market shares (s) to be calculated from the model

$$s_{jt} = \frac{1}{I} \sum_{I} \frac{e^{\delta_{jt}+\mu_{ijt}}}{1 + \sum_{j} e^{\delta_{jt}+\mu_{ijt}}}, \tag{4.3}$$

where I is the number of consumers. To disentangle the values of α, β from ξ requires use of appropriate instruments for price. This is because the unobserved component is expected to be positively correlated with price. In other words, the econometrician recognizes that there are characteristics that are unobserved to him but not the consumer which are expected to be positively correlated with price. In the specific case of no heterogeneity ($\mu_{ijt} = 0$), the model can be solved analytically.

If appropriate instruments (Z) are available, the estimates for α, β become,

$$\alpha, \beta = \left(X' \Delta X\right)^{-1} X' \Delta \delta, \tag{4.4}$$

where, $\Delta = Z(Z'Z)^{-1}Z'$.

If $\mu_{ijt} \neq 0$, the model can be solved with a Berry et al. (1995) contraction map. In practice, I use a modified version proposed by Nevo (2000b). Let μ_{ijt} have the form

$$\mu_{ijt} = \sum_{m=1}^{M} v_{i,m} \sigma_m x_{j,m}, \tag{4.5}$$

where, $v_{i,m} \overset{i.i.d.}{\sim} N(0,1)$, and σ_m is the standard deviation of consumer preferences for product attribute m. This specification allows consumers' preferences to vary across product attributes but does not allow interactions between attributes. The Nevo (2000b) contraction map solves for δ_{jt} by perfectly matching predicted shares $s_{ijt}(x, p, \delta, \sigma, \xi)$ to shares from the data $S_{ijt}, \ \forall i, j, t$ for a given set of guesses σ. Nevo (2000b) suggests updating δ_{ijt} as follows,

$$\delta'_{ijt} = \delta_{ijt} + \ln(S_{ijt}) - \ln(s_{ijt}(p_t, x_t, \delta_t, \sigma)), \tag{4.6}$$

until convergence.

With this in place, the same first-order condition from before can be applied, specifically, $\alpha, \beta = \left(X' \Delta X\right)^{-1} X' \Delta \delta$. From this I can back out the unobserved product attributes ξ.

$$\xi = \delta - \alpha p - X\beta. \tag{4.7}$$

All that is left is to solve for σ. The solution to sigma, is to find the set of value that minimize the sum off the errors from the above process. This amounts to solving the minimization problem,

$$\min_{\sigma} \xi' \Psi \xi \tag{4.8}$$

where Ψ is an appropriate weighting matrix.

It is possible to simultaneously estimate the supply side of the market. This requires solving the first-order condition resulting from the firm's profit maximization problem.

$$\max \Pi_f = \sum_{j \in J_f} (p_j - mc_j) s_j(p, X, \xi) \tag{4.9}$$

where, f indexes firms, and J_f the suite of products offered by firm f. If we assume a linear system for marginal cost, i.e,

$$mc_j = w_j \gamma + \omega_j \tag{4.10}$$

where w_j are product observables that impact marginal cost, γ are parameters to be estimates, and ω is an error term (which is possibly correlated with price.) Solving the first-order condition gives the following relation,

$$p - \Gamma(p, X, \xi)^{-1} s(p, X, \xi) = w\gamma + \omega \tag{4.11}$$

where $\Gamma_{lj} = \frac{-\partial s_l}{\partial p_j}$ if product l and j are produced by the same firm, zero otherwise. The benefits of estimating the supply side are that it allows the potential to estimate efficiency gains due to economies of scales when considering potential mergers. Additionally, there can be efficiency gains in the estimation regarding the standard errors of demand side instruments.

Welfare

With estimates in hand, the model lends itself quite nicely to welfare calculations. From Zhao et al. (2008), the welfare given a specific option set is

$$E(Max(U_{ijt},\ \forall j)) = \ln \sum_j e^{V_{ijt}} \tag{4.12}$$

This can be used to calculate the improvement in welfare when the set of options expands

$$E\left(Max(U_{ijt}) - Max(\tilde{U}_{ijt})\right) = \ln \sum_j e^{V_{ijt}} - \ln \sum_j e^{\tilde{V}_{ijt}} \tag{4.13}$$

Finally, the compensating variation can be obtained by scaling the previous result by the marginal utility of income. This provides a monetary figure that is equivalent to the loss in utility the individual would sustain, were the choice set varied from $V \rightarrow \tilde{V}$.

$$E(CV_{it}) = -\frac{1}{\alpha + \sigma_\alpha v_{i\alpha}} \left(\ln \sum_j e^{V_{ijt}} - \ln \sum_j e^{\tilde{V}_{ijt}} \right) \tag{4.14}$$

Mergers

This model can be used to examine welfare applications from proposed mergers. There are a few common methods for evaluating mergers in the literature. One such method is used by Nevo (2000a) and Petrin (2002). The simulation uses demand estimates from the original estimation to calculate post-merger outcomes. Another, developed by Morrow and Skerlos (2011) iterates over the markup equation to solve for post-merger shares. Either method allows me to back out on estimate of the effect on price. Using these estimates, I can calculate the effect on the Herfindahl-Hirschman Index (HHI), consumer surplus and producer surplus.

Data

The data in this analysis is provided by the Internation Data Corporation (IDC). The data is the US subset of their Worldwide Quarterly Mobile Phone Tracker. The sample ranges from the first quarter of 2007 to the third quarter of 2014. Over this time-period I have market level data on the price and amount of sales of all phones sold in the US at the quarterly frequency. In addition to market values, the dataset has a select number of product attributes on the phones, such as screen size, network connectivity, etc. However, the attribute data is only available for smartphones sold in or after 2010Q1. Due to the lack of data on all phones pre-2010, and feature phones thereafter, I have restricted my analysis to just the smartphone market for the years which I have detailed data. The data from 2007Q1-2009Q4 account for 40% of the observations, 39% of the sales, and 16% of the revenue from overall dataset. The feature phones from 2010Q1-2014Q3 account for 25% of observations, 22% of sales, and 5% of total revenue. In addition to the product data provided by IDC, I acquired additional data on smartphone models from `gsmarena.com`. The data I collected is on battery life and physical specifications.

Selection Criteria

An important descriptor of phones in the US market is their availability at the major carriers. The four primary carriers are Verizon, AT&T, Sprint and T-Mobile. The observation of carrier availability is an important factor in determining the appropriate market for the phone. However, the dataset used in this study only has aggregate sales. On top of that, carrier information is not provided. There are a few possible solutions to this problem, on which I elaborate later. In this section I propose a basic solution.

Most phones are requisitioned for a specific carrier (exclusive) or have general availability and are stocked by all major carriers. Since I cannot determine the specific phones or market sizes for exclusive phones, I instead focus on the generally available phones. By implementing a selection criterion, I aim to isolate the sample of phones that are universally available. Phones that are stocked by all carriers have a larger reach and should ship in larger quantities. The selection criteria I use is based off that idea. I select a subsample of phones that satisfy the following criteria:

$$\sum_{t} \frac{U_{jt} N_t}{U_t} > 4 \text{ where } U_t = \sum_{j=1}^{J} U_{jt}. \tag{4.15}$$

In the above, U_{jt} is the number of units sold by product j in quarter t and N_t is the number of products on the market in quarter t. I selected this criteria over a purely unit sold measure to make the criteria invariant to the growing market. In other words, the aggregate market sales for a phone must be above the expected average over the time-period for which the phone was on the market. The cutoff of 4 comes from Table 24 below, which shows that the average phone spends 366 days (4 quarters) on the market. The high cutoff is due to the preference to reduce the likelihood of a type II error (failing to omit an exclusive phone) over a type I error (rejecting a universal phone.) The downfall is that the remaining sample will only be of popular universal phones. When presenting results, I perform sensitivity analysis to the selection criteria. Below are some summary statistics for the dataset containing all smartphone and the subsample based off of the selection criteria.[2]

Summary Statistics

Table 24 provides the summary statistics for both the entire smartphone market, and the subsample generated by the selection criteria.

The data for this analysis is specific to the phones and the amount of sales from the perspective of the firm. The firm in this case being the company that

[2]This is common practice in the sabermetric community to control for player performance across era. Source: `https://www.fangraphs.com/library/offense/wrc/`

TABLE 24.
IDC Data Summary Statistics

	Complete		Selection Criteria	
Variable	Mean	Std	Mean	Std
Price ($)	339	152	427	167
Units Sold (Thousands)	241	629	814	1263
Time on Market (Days)	366	204	587	229
Camera (Megapixels)	5.74	3.05	6.00	3.10
Memory (Gigabytes)	7.64	13.6	12.7	15.7
Thickness (Millimeters)	11.8	3.79	11.4	2.35
Weight (Grams)	137	25.0	136	21.8
Battery Life (Hours)	9.19	5.62	8.56	4.55
Battery Capacity (Milliampere Hours)	1729	526	1677	495
Screen Coverage (% of Face)	54.9	11.7	55.1	13.1
Processor Speed (GHz)	1.19	0.37	1.16	0.38
Number of Cores	1.81	1.17	1.65	0.97
N	2244	-	403	-
Number of Models	528	-	51	-
Number of Firms	30	-	8	-

$T = 19$. The average price of a purchased phone from the entire dataset was $456 while it was $526 in the subsample.

produces the phone (e.g. Apple, Samsung, etc.) In the United States, phones are generally purchased congruently with an access plan from a major provider (e.g. Verizon Wireless, Sprint, etc.)

Market History

While the smartphone market has 30 different firms contributing products, the market is dominated by only two firms. Apple accounts for 45.4% of the total revenue of the market, while Samsung accounts for 24%, or 43.9% of the remaining market. However, this was not always the case, in Figure 9 below, we see that at the beginning of the sample, BlackBerry began hemorrhaging its market share. The Samsung we know today, Apple's prime competitor, was not always a foregone conclusion. Throughout 2010 and most of 2011, Samsung was in a fierce fight with HTC for their market position. Only with the success of the Galaxy S line was

Samsung able to position itself as the number two player in the market. Since 2012, there has not been any substantial movement in market shares of the top firms.

FIGURE 9.
Revenue Weighted Market Shares

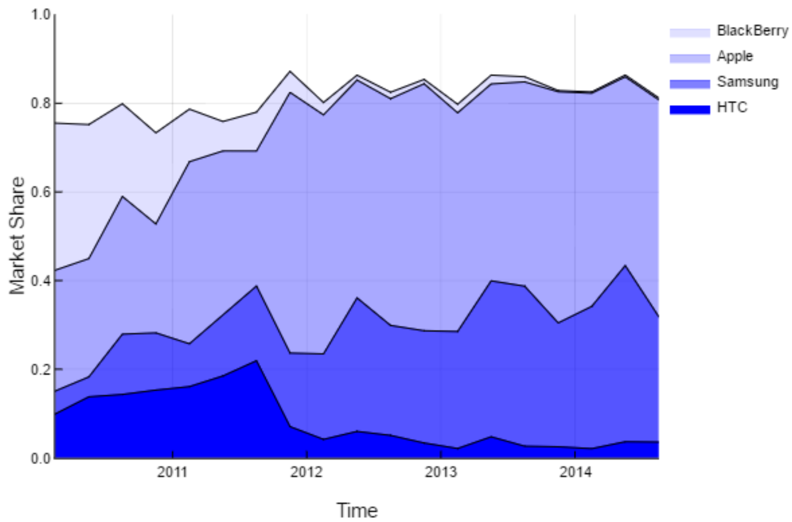

Most sales for both Apple and Samsung come from their flagship line, the iPhone and Galaxy S, respectively. The two product lines are remarkably similar. Both products have a distinct product cycle lasting about three years. Both see immediate and sustained high sales upon introduction, only to experience a drop-off with the introduction of the next model. The only exception to this is the iPhone 4s, which sold considerably well after the introduction of the iPhone 5. One possible explanation for this is that the iPhone 4s was the first model to come with Siri voice assistant. Since their introduction, both lines have been refreshed about once per year. As the dominant firm, Apple chooses to introduce their new iPhone in the fall, during the height of shopping season, while Samsung releases their new Galaxy S phone separated by almost exactly 6 months. Figure 10 visualizes the product life cycle of the iPhone and Galaxy S lines.

With every new product, the producers iterate and improve on their design. This does not always happen in every dimension, for example, the iPhone 4s increased the megapixels in the camera over the previous iteration, however it

FIGURE 10.
Product Life Cycle of Popular Products

(a) iPhone
Product Cycle

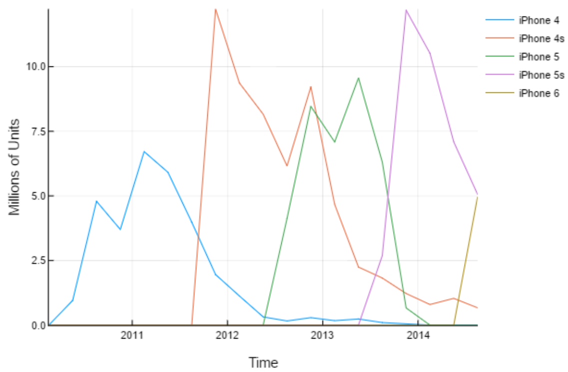

(b) Galaxy S
Product Cycle

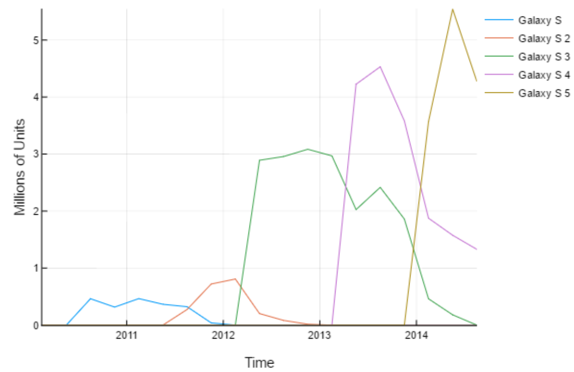

Notes: Data from the Internation Data Corporation, 2010Q1-2014Q3.

made no improvements on the thickness of the device. The general trend among characteristics is always towards better components. I present the trend for camera quality and thickness (Figure 11). In these diagrams, the value for any singe quarter is the average of phones sold specifically in the quarter, rather than the average of all devices in the hands of consumers at that point in time. Recall from Figure 9, we saw the market shares of the leading firms had remained constant over the past three years. Despite constant market shares, the market has continued to evolve in the form of technological advance. For welfare estimates to be precise, the model must be able to accurately measure the value that these continued improvements bring to consumers.

Results

In this section I present my findings. Prior to the final results, I discuss the different instruments available and their relative benefits. Upon arriving at the optimal combination of controls and instruments, I present the results in the context of a model with homogeneous consumers, and one of random coefficient. I

FIGURE 11.
Evolution of Camera Quality and Thickness

(a) Improvement in Camera Quality

(b) Reduction in Phone Thickness

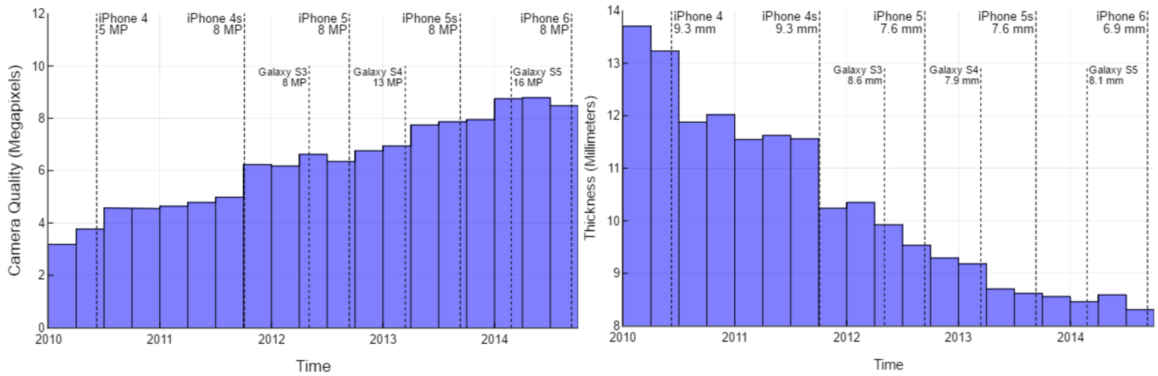

Notes: Data from the Internation Data Corporation, 2010Q1-2014Q3.

find that there is quite a large different in welfare effects, dependent on if consumer heterogeneity is included in the model.

As I have previously discussed, the general form of the model is

$$u_{ijt} = (\alpha + \sigma_\alpha v_{i\alpha})p_{jt} + x'_{jt}\bar{\beta} + \sum_k \sigma_k v_{ik} x_{jkt} + \xi_{jt} + \epsilon_{ijt} \tag{4.16}$$

when the heterogenous terms are ignored, this can be solved analytically for shares.

$$\ln(\text{Share}_{jt}) - \ln(\text{Outside Share}_t) = \alpha p_{jt} + \sum_{j=1}^{J} \beta_j x_{jt} + e_{jt} \tag{4.17}$$

Note that the right-hand side is a linear system. The two variables on the left side are reported in the data, this can be solved analytically using the relevant explanatory variables and appropriate instruments for price. I focus on a set of demand side variables in the analysis

- Time on Market: The amount of time, in quarters, the phone has been on market. A just released phone takes a value of zero. The longer a phone is on the market the lower its sales are expected to be.

- ln(Storage Space): The amount of files and/or media that can be stored on the device. It is expected that larger capacities are sought after.

- ln(Camera Quality): A portable camera is one of the most sought after features of the smartphone. As such, better qualities (measures in megapixels) should have a larger coefficient estimate.

For the supply side of the model I include a set of variables that might have an impact on the marginal cost of the individual units.

- ln(Storage Space): It is expected that larger capacities increase the production cost.

- ln(Camera Quality): Shipping phones with better cameras should increase the cost of production.

- Time Trend: It is reasonable to expect the production process to be improved upon over time. The time trend should have a negative estimate.

The above are not the only explanatory variables used in the model. The relevant control variables are included as well. In my analysis, I control for time with quarterly dummies on demand and firm dummy variables on both supply and demand. A more detailed analysis of these controls can be found in Appendix B. While other explanatory variables were considered, the degree of collinearity between alternative attributes and those listed above made included them unrealistic. The following is a discussion of the potential instruments used in my analysis.

Instruments

In this section, I discuss the options available to instrument for price and the choice ultimately adopted. Commonly in this line of literature, authors use instruments with the purpose of capturing how “far” a product is from its nearest competitors as defined by the product space. Examples of these would be the original instruments from Berry et al. (1995) or differentiation IVs developed by Gandhi and Houde (2016). Berry et al. (1995) accomplished this by aggregating values of product characteristics differentiated by offering from the same firm, then by the competitions' products. Gandhi and Houde (2016) proposes two sets of instruments. Like Berry et al. (1995), the differentiation IVs considers each product attribute in isolation. The first is a measure of the number of other products on the market within one standard deviation in the specific attribute. The second is a summation of all second-order polynomials between the product and others on the market. If desired, this second set of instruments can extended beyond a single attribute by considering the second-order interaction among attributes. In my analysis, I use the instruments proposed by BLP. More details on this choice can be found in Appendix B.

Findings

It been shown that even in the absence of average value for an attribute it is possible to obtain a measure of the spread of preference for that attribute. The basis for this is by extracting the information contained in a products' sales relative to that of product space over time. Putting together all of the pieces I present the estimates for the full model in Table 25.

Here the coefficient on price (α) is appropriately negative although still insignificant at conventional levels. Many of the other variable have the expected sign as well. The coefficient on the time since release is substantially negative, implying the individuals strongly prefer newer phones. The storage capacity does seem to favor larger capacities, if only slightly. This is an improvement over the previous models. Finally, the preference for camera quality is also stronger under the full specification.

TABLE 25.
Comparison of Estimates from Various Specifications

Demand Estimates	OLS	IV	Full
Price	0.213	-6.07	-11.6
	(0.845)	(0.337)	(0.184)
Time on Market (Quarters)	**-0.230**	**-0.278**	**-0.319**
	(0.000)	(0.000)	(0.000)
Log(Gigabytes of Memory)	-0.096	0.065	0.331
	(0.332)	(0.762)	(0.127)
Log(Megapixels on Camera)	**0.971**	**1.82**	**2.59**
	(0.016)	(0.030)	(0.031)
Supply Estimates			
Log(Gigabytes of Memory)			**0.027**
			(0.000)
Log(Megapixels on Camera)			**0.167**
			(0.000)
Trend			**-0.016**
			(0.000)
Quarter (Demand)	✓	✓	✓
Firm (Demand and Supply)	✓	✓	✓
R^2	0.415	0.315	—

p-values are reported in parentheses. p-values are generated using heteroskedastic robust standard errors and are clustered at the quarterly level. Estimates significant at the 5% level are in **bold**. BLP instruments are used. N = 403.

Welfare

The goal of this exercise is to determine a welfare value of all phones in the market by time-period. These estimates provide context into the scale of the market and from that the importance of relevant policy measures. Using these estimates, I can calculate the level of compensating variation. Figure 12 gives a time-series representation of the welfare added each period.

FIGURE 12.
Estimates for Compensating Variation

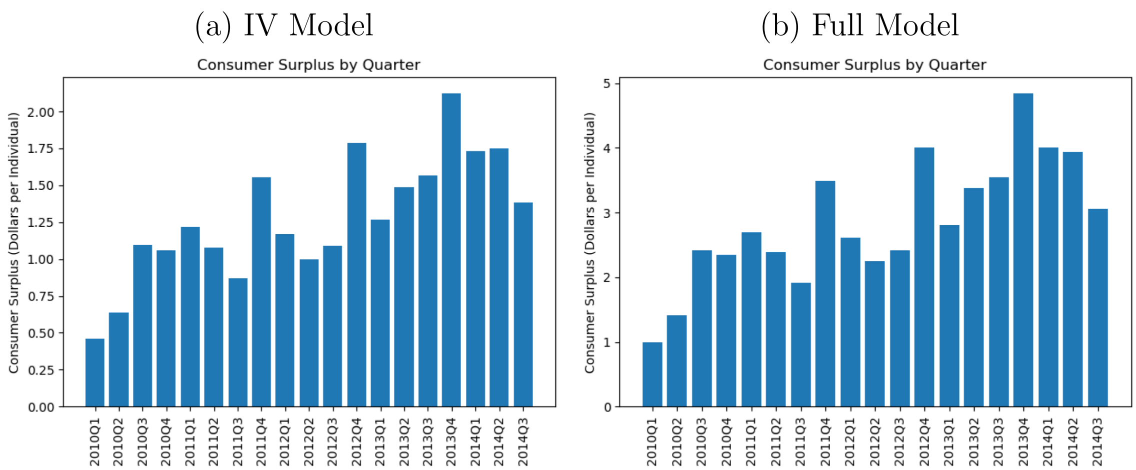

Notice the stark differences between the two models. While the shape is the same, the full model estimates consumer surplus over 100% larger. The returns to the consumer in the model where only demand was estimated were \$5 per consumer per year. In the model with supply side estimates, consumers benefited by \$11.50 per consumer per year. Aggregated over the population, the benefit exceeds \$3 billion annually.

Merger Analysis

With a fully solved random coefficients model, I can perform counterfactuals such as proposed mergers. I propose two potential mergers. The first is an acquisition of Motorola by Samsung. The second is a merging of Motorola and LG Electronics, the third and fourth largest firms of the time, respectively. Both analyses assume the merger takes place at the end the second quarter of 2013. The first result of note is the effect on market concentration. In the quarters after the mergers, the HHI increased by over 500 on average in the case of merger 1. The result of the second merger is more tempered, with an increase of just over 100. These effects can be seen in Figure 13.

FIGURE 13.
Estimates for Changes in HHI and Consumer Surplus from Mergers

(a) Change in HHI (b) Change in CS

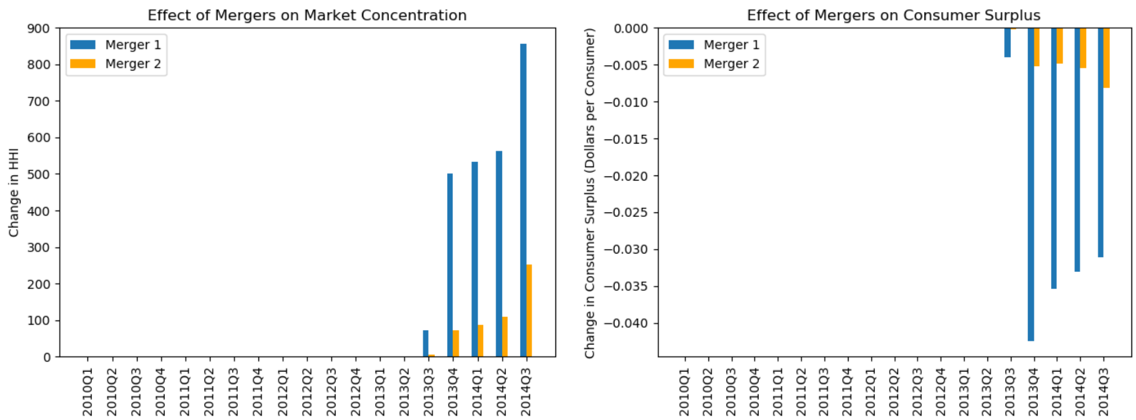

Similarly, the decline in consumer surplus is substantially larger when considering merger 1. Evidenced in Figure 13, the first merger hurts the average consumer by about \$0.15 over the remaining sample period. When factoring in gains from the producers, the overall welfare of the market would have decreased by \$38 million because of an acquisition of Motorola via Samsung in the summer of 2013. The merger in question would have consolidated power in a market that is already dominated by two firms. It should not be surprising that this proposed merger would have drastically decreased welfare. The second of the two proposals would have created a potential potent third player. Due to neither Motorola nor LG being prominent players, the negative effects of the proposal are significantly lower. The second proposed merger would have reduced total surplus by \$6 million.

Conclusion

While the previous literature has long had the tools available to perform welfare analysis, such projects are rarely undertaken in markets of substantial size. Of which the smartphone market is one of the largest. In this paper, I apply a random coefficient model of heterogeneous demand to data provided by the IDC. One of the most problematic features of this market is an incredible strong preference for expensive phones, to the extent where consumers are seemingly ambivalent to price. In my analysis, I instrument price and make good use of controls to break through the strong upward trend on prices.

I find that the market has provided over $860 billion in surplus to the consumers, over half of that coming from the two largest firms Apple, and Samsung. While there have not been any significant shifts in market share over the later parts of the sample period, it is not to say that there have not been attempts. Some of the largest players in the tech industry have tried to break into the smartphone market. Recently Google has made another push into the market with the acquisition of HTC after their failed bid to revitalize Motorola.[3][4] The market could be further limited by policies that have the potential to reduce foreign involvement, such as the Commerce Departments ban on ZTE, or the Intelligence Department Vendetta on Huawei. [5][6] In a market what such a large chunk goes to only two firms, any move towards consolidation should be met with a skeptical eye. Using the model presented, it is possible to conduct welfare considerations of any proposed merger.

[3] Epstein (2014)

[4] Russell (2018)

[5] Freifeld (2018)

[6] Cheng (2018)

CHAPTER V

CONCLUSION

I use a variety of structural modeling techniques to provide insight on differing markets with an eye toward the role that consumer heterogeneity plays in those markets.

In Chapter II, we consider the U.S. automobile industry. We write down a dynamic model of vertical quality differentiation. In our model, consumers optimally choose what types of vehicles to purchase, and how long to own those vehicles. This innovation is substantial, because previous studies that use dynamic models to study the industry do not allow for differentiation in the type of vehicles. We can more precisely estimate the effect of the Cash for Clunkers program. We do this by noting that the program incentivized households to trade-in fuel-inefficient trucks for newer fuel-efficient cars. By exploiting the differentiation between cars are trucks, we estimate that the benefit of CO_2 offsets achieved by the program are larger than the conclusion of the rest of the literature. Additionally, we find that if the level of the maximum credit offered to households had been set at $2,500 instead of $4,500, the environmental benefit associated with $C0_2$ offsets would have been 88% larger. While the distinction between car and trucks led to increased environmental benefits, by forcing this substitution, the program reduced its desired effect on economic stimulus. This is because the trucks that households substituted away from are typically more expensive than their sedan counterparts.

In Chapter III, I continue to examine the Cash for Clunkers program and the U.S. automobile industry. The analysis in Chapter II did not lend itself to incorporating the microlevel data that is available from the CEX. In Chapter III, I propose new methods to analyze the data. The proposed methods use ML techniques to compact the size of the data-space so that estimates can be obtained with standard computational methods. I focus on two different techniques, principle component analysis, and the use of a neural network. Using the methods presented in Chapter III I can trace the effects of the Cash for Clunkers program to which demographic groups it affected. Likewise, I can follow the same procedure to determine which groups were affected by numerous counterfactual exercises. The primary benefit of this research is that it would allow policy makers to implement a program nationwide while targeting specific groups of individuals or regions of the country.

Finally, in Chapter IV, I consider the smartphone industry, which has seen unprecedented growth in recent years. This industry is already very concentrated, with very few firms controlling a significant portion of the market. This is worrisome because consumers can be negatively impacted because of severe market concentration. With large tech firms continuing to acquire start-ups at a furious pace, it become necessary to understand these industries to appropriately analyze anti-trust concerns. In my analysis, I estimate a random coefficients model to capture the effect of technological increases as well as characterize the state of the industry. I find that, on average, the American consumer benefits by \$11.50 per year from technological advancements in the industry from 2010-2014. Regarding anti-trust concerns, I preform hypothetical merger analysis. I find that even when considering the merging of non-dominant firms, the costs of consumers far outweighs the gains from economies of scale. Like the other chapters of my dissertation, I show that if the researcher does not include consumer heterogeneity in the analysis, that the estimates obtained are substantially ill-informed.

APPENDIX A

DECISION RULES FOR TWO-VEHICLE HOUSEHOLDS

Presented below is the closed form steady-state value of a two-vehicles households ofr a given combination of a, n, and m.

$$
\begin{aligned}
V_2(y, \gamma, \alpha, a, n, m) = \frac{1}{1-\beta^{n-a+1}} \Bigg[& \left[\gamma(q_a + \alpha q_{a+m}) + \ln(y - p_a + \Psi(p_n) - C^2)\right] \\
& + \sum_{j=1}^{n-m} \beta^j \left[\gamma(q_{a+j} + \alpha q_{m+j}) + \ln(y - C^2)\right] \\
& + \beta^{n-m+1} \left[\gamma(q_a + \alpha q_{a+n-m+1}) + \ln(y - p_a + \Psi(p_n) - C^2)\right] \\
& + \sum_{j=n-m+2}^{n-a} \beta^j \left[\gamma(q_{a+1+j-n+m-2} + \alpha q_{a+j}) + \ln(y - C^2)\right] \Bigg].
\end{aligned}
$$

APPENDIX B

CONTROLS AND INSTRUMENTS

Controls

An important part of any economic analysis is careful consideration of the control variables that go into the model. Table A1 below documents an extensive list of potential control variables. The results in this table are from running a BLP style regression on shares where price is instrumented for with the BLP instruments. A detailed discussion on the choice of instrument can be found later in the appendix.

The first column shows what the estimate would like with a complete lack of fixed effects. The second column only controls for the changing landscape of the market over time. Moving down the list increases the scale of the model's fixed effects. Firm fixed effects are added to control for differences in popularity across firms. This is necessitated by the perception that consumer have about different brands. A phone offered by ZTE would not be expected to sell as well as a phone with the same internal and external attributes offered by Apple. Some of this effect is due to advertising while some are unobserved qualities of the phone, such as customer service. The unobserved effects constant across all phones from a single provider are incorporated into the firm fixed effects.

When all this is put together, the final model had the only negative price coefficient among the specifications. While the coefficient is not significant, this is the only the results from the homogeneous version of the model. In the main body of the paper, I discuss the role that heterogeneity plays in the estimation of this model.

TABLE A1.
Comparison of Controls

Demand Estimates	None	Quarter	Firm
Price	0.918	0.768	-11.6
	(0.208)	(0.552)	(0.184)
Time on Market (Quarters)	**-0.150**	**-0.174**	**-0.319**
	(0.000)	(0.000)	(0.000)
Log(Gigabytes of Memory)	-0.109	-0.051	0.331
	(0.067)	(0.539)	(0.127)
Log(Megapixels on Camera)	-0.027	-0.176	**2.59**
	(0.889)	(0.360)	(0.031)
Supply Estimates			
Log(Gigabytes of Memory)	**0.028**	**0.032**	**0.027**
	(0.000)	(0.000)	(0.000)
Log(Megapixels on Camera)	**-0.033**	**-0.039**	**0.167**
	(0.006)	(0.028)	(0.000)
Trend	0.002	0.002	**-0.016**
	(0.387)	(0.457)	(0.000)
Quarter (Demand)		✓	✓
Firm (Demand and Supply)			✓

p-values are reported in parentheses. p-values are generated using heteroskedastic robust standard errors and are clustered at the quarterly level. Estimates significant at the 5% level are in **bold**. BLP instruments are used. N = 403.

Instruments

In this section, I discuss the options for instruments available and the choice ultimately adopted. Commonly in this line of literature, authors use instruments with the purpose of capturing how "far" a product is from its nearest competitors as defined by the product space. Examples of these would be the original instruments from Berry et al. (1995) or differentiation IVs developed by Gandhi and Houde (2016). Berry et al. (1995) accomplished this by aggregating values of product characteristics differentiated by offering from the same firm, then by the competitions' products. Gandhi and Houde (2016) proposes two sets of instruments. Like Berry et al. (1995), the differentiation IVs considers each product attribute in isolation. The first is a measure of the number of other products on the market within one standard deviation in the specific attribute. The second is a summation of all second-order polynomials between the product and others on the market. If desired, this second set of instruments can extended beyond a single attribute by considering the second-order interaction among attributes.

Table A2 compares the results of an instrumental variable regression using the different instruments discussed in this section. As you can see, the differentiation instruments moved the estimate on price into the positive region. The standard BLP instruments were the only instruments with enough power to move the coefficient on price into an economically sensible region.

Investigation of the Selection Criteria

In Table A3 below I present the results of varying the selection criteria. I inquire as to the effects of lowering the threshold from 4 to 2, 1 and 0 (the full sample), as well as an alternative rule for sampling, based off units sold.

TABLE A2.
Investigation of Potential Instruments

Demand Estimates	1^{st}	2^{nd}	BLP
Price	5.48	-8.21	-11.6
	(0.224)	(0.716)	(0.184)
Time on Market (Quarters)	**-0.167**	**-0.224**	**-0.319**
	(0.000)	(0.000)	(0.000)
Log(Gigabytes of Memory)	-0.187	-0.048	0.331
	(0.125)	(0.696)	(0.127)
Log(Megapixels on Camera)	0.296	**1.55**	**2.59**
	(0.583)	(0.028)	(0.031)
Supply Estimates			
Log(Gigabytes of Memory)	**0.030**	1.32	**0.027**
	(0.012)	(0.974)	(0.000)
Log(Megapixels on Camera)	**-0.204**	-5.13	**0.167**
	(0.002)	(0.978)	(0.000)
Trend	**-0.021**	0.079	**-0.016**
	(0.015)	(0.989)	(0.000)

p-values are reported in parentheses. p-values are generated using heteroskedastic robust standard errors and are clustered at the quarterly level. Estimates significant at the 5% level are in **bold**. Time and firm level fixed effects are used. N = 403.

TABLE A3.
Alternative Selection Criteria

		Invariant Criteria (Cutoff)			Sales Criteria (Units)	
	Full Sample	(1)	(2)	(4)	(250000)	(1000000)
Price	7.74	0.549	3.20	-7.32	-1.48	-1.94
	(0.220)	(0.911)	(0.298)	(0.281)	(0.676)	(0.533)
Time on Market	**-0.207**	**-0.246**	**-0.241**	**-0.282**	**-0.248**	**-0.250**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log(Memory)	-0.228	-0.022	-0.109	0.128	0.054	0.033
	(0.244)	(0.900)	(0.407)	(0.597)	(0.671)	(0.802)
Log(Megapixels)	-0.585	0.295	0.140	**2.07**	0.493	**1.16**
	(0.451)	(0.591)	(0.627)	(0.021)	(0.218)	(0.000)
R^2	0.363	0.396	0.4072	0.288	0.383	0.426
N	2214	1568	961	403	1614	668

p-values are reported in parentheses. p-values are generated using heteroskedastic robust standard errors and are clustered at the quarterly level. Estimates significant at the 5% level are in **bold**. Quarter and firm fixed effects are used. Price is instrumented with BLP instruments.

REFERENCES CITED

Abrams, B. A. and Parsons, G. R. (2009). Is CARS a clunker? *The Economists Voice*, 6(8):1–4.

Aidt, T. S. (1998). Political internalization of economic externalities and environmental policy. *Journal of Public Economics*, 69(1):1–16.

Anderson, S. T., Parry, I. W. H., Sallee, J. M., and Fischer, C. (2011). Automobile Fuel Economy Standards: Impacts, Efficiency, and Alternatives. *Review of Environmental Economics and Policy*, 5(1):89–108.

Athey, S. (2018). The impact of machine learning on economics. In *The Economics of Artificial Intelligence: An Agenda*. National Bureau of Economic Research, Inc.

Athey, S., Blei, D., Donnelly, R., Ruiz, F., and Schmidt, T. (2018). Estimating heterogeneous consumer preferences for restaurants and travel time using mobile location data. *AEA Papers and Proceedings*, 108:64–67.

Auto, Z. (2015). Under 3,000 tickets of the 5,000, allocated to legal entities; apply one more week.

Ayres, R. U. and Kneese, A. V. (1969). Production, consumption, and externalities. *The American Economic Review*, 59(3):282–297.

Barahona, H., Gallego, F., and Montero, J.-P. (2015). Adopting a cleaner technology: The effect of driving restrictions on fleet turnover.

Becker, G. S. (1983). A theory of competition among pressure groups for political influence. *The quarterly journal of economics*, 98(3):371–400.

Berry, S., Kortum, S., and Pakes, A. (1996). Environmental change and hedonic cost functions for automobiles. *Proceedings of the National Academy of Sciences*, 93(23):12731–12738.

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

Björkegren, D. and Grissen, D. (2017). Behavior revealed in mobile phone usage predicts loan repayment. *CoRR*, abs/1712.05840.

Blinder, A. S. (2008). A modest proposal: Eco-friendly stimulus. *The New York Times*.

Böckers, V., Heimeshoff, U., and Müller, A. (2012). Pull-forward effects in the german car scrappage scheme: A time series approach.

Brynjolfsson, E., Hu, Y., and Smith, M. D. (2003). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science*, 49(11):1580–1596.

Bunkley, N. (2009). Government will end clunker program early. *The New York Times*.

Butcher, L. (2018). Vehicle scrappage schemes. Briefing Paper CBP 8091, House of Commons.

Cheng, R. (2018). The us government's pressure on huawei takes its toll.

Copeland, A. and Kahn, J. (2011). The production impact of "cash-for-clunkers": Implications for stabilization policy.

Crawford, G. S., Lee, R. S., Whinston, M. D., and Yurukoglu, A. (2015). The welfare effects of vertical integration in multichannel television markets. Working Paper 21832, National Bureau of Economic Research.

Epstein, Z. (2014). Google bought motorola for $12.5b, sold it for $2.9b, and called the deal 'a success'.

Freifeld, K. (2018). Exclusive: China's zte signed preliminary agreement to lift u.s. ban: sources.

Gandhi, A. and Houde, J.-F. (2016). Measuring substitution patterns in differentiated product industries.

Gavazza, A. and Lanteri, A. (2018). Credit shocks and equilibrium dynamics in consumer durable goods markets.

Gavazza, A., Lizzeri, A., and Roketskiy, N. (2014). A quantitative analysis of the used-car market. *American Economic Review*, 104(11):3668–3700.

Glaeser, E. L., Hillis, A., Kominers, S. D., and Luca, M. (2016). Crowdsourcing city government: Using tournaments to improve inspection accuracy. Working Paper 22124, National Bureau of Economic Research.

Goel, S., Rao, J. M., and Shroff, R. (2016). Precinct or prejudice? understanding racial disparities in new york city's stop-and-frisk policy. *Ann. Appl. Stat.*, 10(1):365–394.

Goettler, R. L. and Gordon, B. R. (2011). Does amd spur intel to innovate more? *Journal of Political Economy*, 119(6):1141–1200.

Greenstein, S. and Ramey, G. (1998). Market structure, innovation and vertical product differentiation. *International Journal of Industrial Organization*, 16(3):285–311.

Grossman, G. M. and Helpman, E. (1996). Electoral Competition and Special Interest Politics. *The Review of Economic Studies*, 63(2):265–286.

Gürtler, M., Gutknecht, S., and Hibbeln, M. T. (2016). The price effect of supply and demand shocks on secondary markets - evidence from the 'cash-for-clunkers' program in germany.

Helfand, G. E., Berck, P., and Maull, T. (2003). Chapter 6 - The Theory of Pollution Policy. In Mäler, K.-G. and Vincent, J. R., editors, *Environmental Degradation and Institutional Responses*, volume 1 of *Handbook of Environmental Economics*, pages 249–303. Elsevier.

Hoekstra, M., Puller, S. L., and West, J. (Forthcoming). Cash for corollas: When stimulus reduces spending. *American Economic Journal: Applied Economics*.

Knittel, C. R., Miller, D. L., and Sanders, N. J. (2016). Caution, Drivers! Children Present: Traffic, Pollution, and Infant Health. *The Review of Economics and Statistics*, 98(2):350–366.

Lenski, S. M., Keoleian, G. A., and Bolon, K. M. (2010). The impact of 'cash for clunkers' on greenhouse gas emissions: a life cycle perspective.

Li, S., Linn, J., and Spiller, E. (2013). Evaluating "cash-for-clunkers": Program effects on auto sales and the environment. *Journal of Environmental Economics and Management*, 65(2):175–193.

Li, S., Tong, L., Xing, J., and Zhou, Y. (2017). The Market for Electric Vehicles: Indirect Network Effects and Policy Design. *Journal of the Association of Environmental and Resource Economists*, 4(1):89–133.

Li, S. and Wei, C. D. (2013). Green stimulus: A dynamic discrete analysis of vehicle scrappage programs. *Available at SSRN 2254375*.

Mäler, K.-G. (1974). *Environmental economics: a theoretical inquiry*. RFF Press.

Mian, A. and Sufi, A. (2012). The effects of fiscal stimulus: Evidence from the 2009 cash for clunkers program. *The Quarterly Journal of Economics*, 127(3):1107–1142.

Miller, K., Wilson, W., and Wood, N. (2018). Stimulus, environmentalism, and inequality reduction through industrial policy: Did cash for clunkers achieve the trifecta?

Morrow, W. R. and Skerlos, S. J. (2011). Fixed-point approaches to computing bertrand-nash equilibrium prices under mixed-logit demand. *Operations Research*, 59(2):328–345.

Munk-Nielsen, A. (2015). Diesel cars and environmental policy.

Nelder, J. A. and Mead, R. (1965). A Simplex Method for Function Minimization. *The Computer Journal*, 7(4):308–313.

Nevo, A. (2000a). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics*, 31(3):395–421.

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

NIADA (2017). 2017 used car industry report.

Persson, T. (1998). Economic policy and special interest politics. *The Economic Journal*, 108(447):310–327.

Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110(4):705–729.

Portney, P. R., Parry, I. W., Gruenspecht, H. K., and Harrington, W. (2003). Policy watch: The economics of fuel economy standards. *Journal of Economic Perspectives*, 17(4):203–217.

PTI (2018). Vehicle scrap policy to see another round of consultations with stakeholders: Nitin Gadkari. *The Economic Times*.

Quan, T. and Singer, E. (2015). Consumer heterogeneity and tied subsidies: Participation in automobile scrappage programs.

Ruiz, F. J. R., Athey, S., and Blei, D. M. (2017). Shopper: A probabilistic model of consumer choice with substitutes and complements. Papers 1711.03560, arXiv.org.

Russell, J. (2018). Google completes its $1.1b deal to buy a chunk of htc's smartphone division.

Rust, J. (1985). Stationary equilibrium in a market for durable assets. *Econometrica*, 53(4):783–805.

Seo, B. and Shapiro, M. H. (2018). Hetergeneous effects of subsidy and infrastructure investment in electric vehicles adoption.

Stigler, G. J. (1971). The theory of economic regulation. *The Bell Journal of Economics and Management Science*, pages 3–21.

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

The Economist (2009). Jump-starting the car industry.

Weiss, P. (2003). Adoption of product and process innovations in differentiated markets: The impact of competition. *Review of Industrial Organization*, 23(3):301–314.

West, S. E. and Williams, Roberton C., I. (2005). The cost of reducing gasoline consumption. *American Economic Review*, 95(2):294–299.

Zhao, Y., Kockelman, K. M., and Karlstrom, A. (2008). Welfare calculations in discrete choice settings: Role of error term correlation. In *Proceedings of the transportation research board 87th annual meeting*. Citeseer.