Title
Essays on Automotive Lending, Gasoline Prices, & Automotive Demand

Author
Schulz-Mahlendorf, Wilko Ziggy

Publication Date
2013

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Essays on Automotive Lending, Gasoline Prices, & Automotive Demand

A dissertation submitted in partial satisfaction of

the requirements for the degree

Doctor of Philosophy in Economics

by

Wilko Ziggy Schulz-Mahlendorf

2013

Abstract of the Dissertation

Essays on Automotive Lending, Gasoline Prices, & Automotive Demand

by

Wilko Ziggy Schulz-Mahlendorf

Doctor of Philosophy in Economics

University of California, Los Angeles, 2013

Professor Connan Andrew Snider, Chair

Two events in 2008-2009 had sweeping effects on the global economy and the US automotive sector in particular: the contraction of credit in 2008-2009 and the massive run-up in crude oil prices in late 2008. In 2008, as credit tightened, auto demand declined, and General Motors and Chrysler edged towards bankruptcy, begging the question: "What role does credit have in determining auto market outcomes?" In late 2008, crude oil prices rose sharply, leading many to ask "How will consumers respond to rising gasoline prices?" This dissertation addresses both of these questions.

In chapter one, I analyze car manufacturers and their lending subsidiaries, or "captive finance companies." I note that via their captive finance companies, car manufacturers compete against traditional banks and are majority lenders in the auto lending market. Government agencies like the Consumer Financial Protection Bureau and the Federal Deposit Insurance Corporation have recently expressed interest in captives and auto lending. They have asked for instance whether captives affect the market price for risk. This leads me to my research question: "How does the use of captive finance affect competition and risk in the auto lending and new vehicle markets?" I specify an equilibrium model of the new vehicle and auto lending markets. I estimate parameters from this model using the universe of loan originations and vehicle transactions from a 20% of all US vehicle dealerships. I find that consumers are more responsive to loan principal than to changes in interest rates. Using the estimated model to conduct a counterfactual in which captives cease to

exist, I also find that competition between captives and traditional banks leads banks to under-price risk.

In chapter two, I ask "How do consumers' vehicle choices respond to changes in their expectations for future gasoline prices?" The 2008 spike in gasoline prices resurfaced questions asked about consumer gasoline consumption and vehicle choice which were originally posed in the 1970's. Since that time, economists have modeled consumers' dynamic responses to gasoline prices. A necessary component of such models is the specification of consumers' expectations over future values of relevant variables like gasoline prices. As data on gasoline price expectations were unavailable, researchers had been forced to make assumptions on how consumers form expectations over future gasoline prices. To address this, I introduce entirely novel data on consumers' gasoline price expectations and stated vehicle preferences. Using these data, I find that consumers' vehicle choices do respond to their expectations for future gasoline prices. Consumers vehicle preferences are persistent, however, which attenuates the effects of gasoline prices on vehicle choices.

The dissertation of Wilko Ziggy Schulz-Mahlendorf is approved.

Moshe Buchinsky

Rosa Liliana Matzkin

Raphael Chaim Thomadsen

Connan Andrew Snider, Committee Chair

University of California, Los Angeles

2013

To my loving family who got me here &

to my leading lady who got me past the finish line

Contents

ACKNOWLEDGEMENTS

I am indebted my committee chair, Connan Snider, for his continued advice, tenacious encouragement, and indefatigable support. I am grateful to my committee members Moshe Buchinsky, Rosa Matzkin, and Raphael Thomadsen for their advice and guidance. I would also like to thank Daniel Ackerberg, Jinyong Hahn, Matthew Kahn, Kathleen McGarry, and Jorge Silva-Risso for their direction, encouragement, input, and dedication to my professional development. My advisors from Duke - Arie Beresteanu, Charles Becker, and Daniel Graham - deserve significant credit for inspiring me to pursue graduate studies in economics.

I gratefully acknowledge Rosa Matzkin for the opportunity to serve as her research assistant. I thoroughly appreciate financial support from UCLA. Above all, I appreciate support via the Chancellor's Prize, Graduate Mentorship Fellowship, and Dissertation Year Fellowship.

I would also like to acknowledge my first year study group - Matthew Hill, Mitsuru Igami, and Anthony Keats - without whom the first year would have been impossible to overcome and entirely devoid of humor. I benefited greatly from the Empirical IO reading group and discussions with its members - Mitsuru Igami, Andraz Kavalar, Jinyoung Lee, Jiyeon Seo, Yong Hyeon Yang, and Hisayuki Yoshimoto. My graduate studies were thoroughly enriched by interactions with Guillermo Beylis, Neil Lessem, Dan Ben Moshe, Tracy Orcholski, and Ryan Vaughn. I am also indebted to all of my instructors, professors, teaching assistants, fellow graduate students, as well as the many undergraduate students whom I had the pleasure of teaching.

VITA

B.S. Economics & B.A. Mathematics, ***Duke University*** **2006**

M.A. Economics, ***Duke University*** **2007**

M.A. Economics, ***UCLA*** **2009**

Chapter 1

Competition & Risk:

The Auto Finance Market & The New Vehicle Market

1.1 Introduction

The new car market is one of the most heavily studied markets in industrial organization. However, one critical component of the industry has been largely ignored – auto finance. In total, auto finance is a one trillion dollar sector. Its magnitude is driven by the fact that 80-90% of new car transactions occur with some form of vehicle financing. From an IO perspective, it is highly relevant that the majority of lending in this industry is originated by manufacturers themselves, via their lending subsidiaries or so-called captive finance companies. This paper is the first to empirically investigate these captive finance companies and their role in the market for new vehicles. In particular, in this paper I ask "How does the use of auto finance companies by car manufacturers affect competition and risk in the new vehicle and auto lending markets?"

The current topic is especially salient in the wake of the US auto bailout. In the run-up to the Chrysler & General Motors bankruptcies, the US Treasury Department weighed the option of letting the two companies' captive finance companies - Chrysler Financial and GMAC - go without government aid and implode. Ultimately, the Treasury Department deemed that the failure of GMAC would lead to the failure of GM and would pose a systemic risk to an already fragile financial sector. It bailed out the ailing captive, but severed

GM's ownership rights.[1] In contrast, Chrysler Financial was allowed to fail. Since then, industry analysts have asked whether car a manufacturer needs to be integrated with a captive finance company. General Motors and Chrysler have presented two alternative trajectories to industry watchers. GM has acquired a new captive, while Chrysler has survived without one. In spite of these examples, it remains an open question whether US car manufacturers are viable without captive finance companies.[2, 3]

Government regulators are similarly interested in car manufacturers' involvement in auto finance. The Consumer Financial Protection Bureau (CFPB) is pursuing regulation of dealer reserves - a markup that car dealers tack on to loans as a fee for arranging financing for their customers. The motivation for this is by analogy to mortgage brokers, who attach fees to the mortgages they arrange for their clients. Mortgage brokers differ fundamentally from car dealers - mortgage brokers cannot change house prices while car dealers can change prices on the vehicles they sell. Any policy directed at restricting dealers' capacity to markup loans must then consider the consequences on vehicle prices. Both captive finance companies and their parent carmakers will be impacted by this pending regulation.

The Federal Deposit Insurance Corporation (FDIC) also has an interest in this space. The FDIC has noted that competition between banks and captive finance companies has led to a weakening of underwriting standards at traditional banks. Competition between the two lender types may lead traditional banks to under-price risk. The intuition driving this concern is that captive finance companies may be able to offer interest rates below the marginal cost of lending provided that they can cross-subsidize this loss with a gain via their vehicle markup. To compete with captive finance companies, traditional banks may relax underwriting standards and weaken lending controls. The FDIC has noted that this may leave some banks' "...loan portfolios increasingly vulnerable to economic downturns."

In this paper, I assess how the use of captive finance companies by car manufacturers affects competition and risk in this market. My analysis allows me to comment on the pending CFPB regulation and the concerns of the FDIC. To accomplish these tasks, I analyze auto finance and vehicle transaction data from a major automotive marketing agency. I then specify and estimate an equilibrium model of the auto finance and new

[1] Warren, et. al (2010)
[2] Henry, Jim. "Flying without a captive - and with one." *AutoNews Online,* April 3rd, 2013. Retrieved, April 28, 2013.
[3] Henry, Jim. "No captive? No problem - so far." *AutoNews Online,* October 12, 2011. Retrieved April 28, 2013.

vehicle markets. I use parameters estimated from this model to conduct a counterfactual experiment in which captive finance companies cease to be controlled by car manufacturers. This counterfactual allows me to answer the key question of the paper: "How does the use of captive finance companies by car manufacturers affect competition and risk in this market?" My counterfactual allows me to speak to whether carmakers are viable without their captives; how regulating loan markups may affect vehicle prices; and whether competition with captive finance leads traditional banks to under-price risk.

All captive finance transactions occur through new car dealers. Therefore, to empirically measure captive finance's effects on competition and risk in auto finance and new vehicle markets, I employ data at the appropriate unit of analysis – new vehicle dealerships. The data comprise the universe of transactions from a 20% sample of all US vehicle dealerships in over thirty geographic markets, from 2007-2012. The data directly yield two key insights about captive finance: 1) car manufacturers use their captive finance companies to extend loans to subprime borrowers and 2) car manufacturers use their captive finance companies to subsidize interest rates on loans to consumers. These two insights are critical to note if one seeks to accurately measure outcomes in this space.

Using these insights, I specify an equilibrium model of the new vehicle and lending markets. I then estimate demand and supply primitives from the model using the dealership data. For demand, consumers solve a discrete choice problem over bundles consisting of a car and a loan. The supply model parsimoniously captures the essential features of competition among carmakers and banks. Each carmaker is comprised of a production and sales division, as well as a captive finance division. Carmakers choose prices and interest rates optimally. Captives compete against banks which are not restricted in the brands to which they can lend. The first order conditions to carmakers' optimization problems create an explicit link between production and lending decisions. In this sense, the model explicitly allows for interest rate subsidies in equilibrium: carmakers can adjust across their product and lending margins. To account for the fact that captives focus lending efforts on subprime borrowers, in the model I allow costs of lending to differ across captives and banks.

In practice, it is empirically challenging to identify the joint elasticity over prices and loan terms. To overcome this challenge, I follow Berry (1994) and estimate demand parameters using data on market level shares of car and loan bundles, as well as data on prices, loan terms, and vehicle characteristics. I face both the standard concern of endogenous prices as well as the non-standard problem of endogenous interest rates on car loans. I follow standard practice to identify demand elasticities in the face of endogenous prices – I use prices of the same car in other geographic markets. To identify demand elasticities in the face of endogenous interest rates, I employ two alternative identification strategies. First, noting that captives raise funds from credit markets to issue new loans, I use the 2008 credit crunch as a cost shifter. Second, noting that changes in expected repayment shift lending costs, I use 90+ day mortgage delinquency rates as an additional cost shifter. Both instrument strategies yield qualitatively similar results. Namely, consumers are more responsive to changes in loan principal than to changes in APRs. Following Nevo (2000 a, b), I combine my demand parameter estimates with the first order conditions from the supply-side model in order to back out firms' marginal costs - both marginal costs of lending and marginal costs of vehicle sales and production.

Using the model primitives, I re-solve for optimal prices and interest rates in counterfactual scenarios. I use a counterfactual experiment to address the core question of the paper – "How does the presence of captive finance affect competition and risk in the new vehicle and auto lending markets?" Assessing the impact of captive finance on competition is straightforward. I compare counterfactual prices and APRs to actual prices and APRs and assess changes in consumer welfare and profits across actual and counterfactual outcomes. Based on the counterfactual experiments, I find that on average APRs increase by roughly 15% and vehicle prices decrease by roughly 11%, relative to the levels observed in the data. This result is intuitive – with less competition, banks respond by raising interest rates on loans. Not only do banks no longer compete with another lender, they cease to compete with lenders who subsidize interest rates on loans. Car manufacturers can recover the profits sacrificed by losing their captives by lowering prices. On net, total surplus is higher in the counterfactual scenario. This is driven by the fact that in relative terms, consumers are more responsive to changes in loan principal than to changes in APRs. Car manufacturers can drop prices, banks can increase APRs, and sales can still increase, netting all firms greater profits in the counterfactual scenario. Since consumers value reductions in principal more than changes in interest rates, they are better off under the counterfactual scenario as well. I find that changes in APRs are inversely correlated with changes in prices. This result is relevant to the CFPB's intent to regulate dealer lending markups. It is conceivable that dealers will respond to caps in loan interest rates and restrictions on dealer loan markups by raising vehicle prices.

With regard to risk, I address an auxiliary question: "Does competition between captive finance companies and traditional banks lead banks to under-price risk?" To answer this question, I exploit the geographic variation inherent in my data. During the time span of my data (2007-2012), geographic markets vary in their observable risk measures (e.g. mortgage delinquency). I draw on this fact and examine changes in APRs across geographic markets. Using Minneapolis as an example of a lower risk market and Miami as an example of a higher risk market, I find that counterfactual interest rates increase by more on an average basis in Miami (15.8% increase) versus Minneapolis (13.8%). I interpret this as APRs increasing more in risky markets, relative to actual outcomes, or that competition between captive finance companies and traditional banks may lead traditional banks to under-price risk.

The paper proceeds as follows. In Section 2, I provide industry background and demonstrate that carmakers use their captives to subsidize interest rates on loans and to lend to subprime borrowers. In Section 3, I present the empirical model. In Section 4, I introduce the data. In Sections 5 and 6, I discuss details of the estimation procedure and provide the estimation results. In Section 7, I discuss the counterfactual and present the counterfactual results. I conclude in the final section.

1.2 Industry Background

1.2.1 Overview

General Motors (GM) established the first automotive finance company - or captive finance company - in 1919 to facilitate installment lending to consumers. In the wake of World War I, banks were unwilling to extend installment credit to car buyers.[4] Available lending came from operations resembling loan sharks - contracts were ambiguously worded and loans were loaded with both high interest rates and high fees. Noting that these loans effectively raised the price of its cars, and thus reduced sales, it served the interests

[4] Banner (1958), Tedlow (1998)

of GM to establish a lending subsidiary. In the years that followed, all major car manufacturers followed suit and established their own captive finance companies. In the early 1980s, captives cemented their roles as safety nets for their parent carmakers. In this period, prime rates hovered near 20%. Traditional banks had all but ceased to lend to car buyers.[5] Captive finance companies sustained car sales by continuing to lend to consumers through this period.

The auto finance industry has continued to grow since the 1980s. The scale of the industry - combining leases and loans on new cars - stood at $1.28 trillion dollars in 2004. In 2012, outstanding consumer car loans stood at $750 billion.[6] The chart below summarizes many of the key features of the contemporary consumer auto finance industry. The vast majority of consumers finance their vehicles. Over 80% obtain a lease or a loan.[7] There are two main channels through which consumers receive auto financing - 'direct financing' where consumers obtain loans directly from a bank, and 'indirect financing' where consumers obtain loans indirectly from a lender through a car dealership. This indirect channel is by far the larger of the two, comprising over 80% of originations and outstanding debt. Indirect lending is also dominated by captive lenders - either subsidiaries or contracted lending partners of auto manufacturers. Since the focus of the current study is captive finance, I concentrate attention on the indirect lending channel. As captive finance companies operate exclusively in dealerships, my dealership data are entirely appropriate for this analysis.

Figure I : Outstanding Auto Debt Receivables, 2004

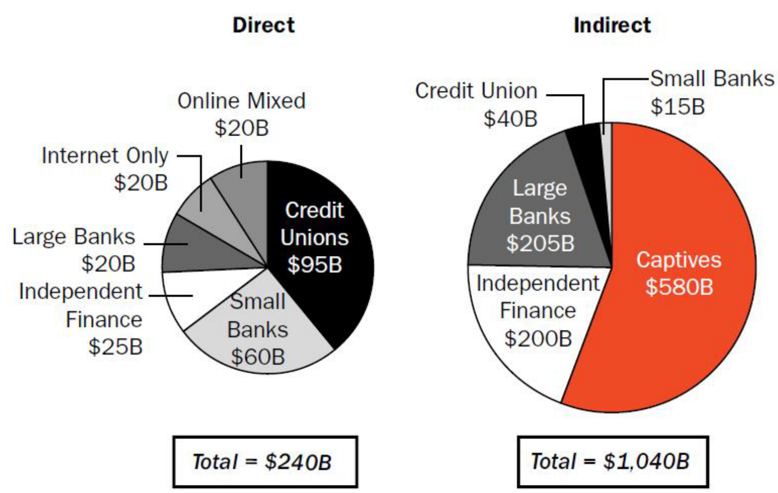

Source : Hersh, et al (2004)

[5] Crain, Keith. "Cars run on money and credit, not gas." *AutoNews Online,* September 22, 2008. Retrieved November 8, 2012.
[6] Lunby, Tami. "Consumers cut up credit cards, but buy cars." *CNN Money,* August 29, 2012. Retrieved November 21, 2012.
[7] Warren et al (2010).

To raise funds to issue new loans, lenders rely on a variety of sources. Auto asset backed securities and commercial paper provide sources for all lenders. Banks can also draw upon deposits. Overall, between 10-15% of US auto loans originated are bundled into asset backed securities. There are stark differences across lender type in the issuance of auto loan asset backed securities. Captives rely predominantly on asset backed securities, with up to two-thirds of auto securitizations being originated by captive finance companies. [8]

The rest of this section draws on the data to underscore features of captive finance which are critical to include in order to accurately model competition in this market. Specifically, I demonstrate that manufacturers use their captive finance companies to subsidize interest rates on loans. I also give evidence of the importance of subprime loan originations to captive originations and manufacturer sales. I use these data findings to inform the model. In the model, optimal price and interest rate setting by carmakers allows for the possibility of interest rate subsidies in equilibrium. Additionally, I allow for differences in cost structure between captives and banks allows for the differential contribution of subprime loan originations to each lender type's business.

1.2.2 Captive Interest Rate Subsidies

Within the auto finance industry, the practice of subsidizing interest rates is referred to as subvention. In practice, carmakers issue payments to their captives for the net present value of the difference between the market rate for a loan and the subvented rate. Carmakers can employ this strategy because they can use their product margin to buffer below marginal cost pricing on loans. Traditional banks do not have this second margin available to them; any pricing below marginal costs of lending would incur negative profit.

Figure III displays histograms of annual percentage rates (APRs) on loans originated through dealerships to consumers receiving financing from either captive finance companies or traditional banks. The data are pooled over markets, time periods and vehicles spanned by the data. Similar figures can be generated on a manufacturer-by-manufacturer basis. This figure serves to demonstrate that captive finance companies are much more willing to charge APRs below the prime rate. Over 20% of the captive finance observations occur with APRs less than 2.5%. In contrast, only 0.1% of traditional bank lending occurs with APRs of less than

[8] Ashcraft, Malz, Poszar (2012), Cetorelli & Peristiani (2012).

2.5%. The prime rate peaked in July 2006 at 8.25% and has declined to 3.25% in January 2009, and persists at that level to this day. Isolating attention to observations from January 2009 onward, similar results obtain - 28% of captive transactions occur at APRs less than 2.5% and 0.2% of traditional bank transactions occur with APRs of less than 2.5%. From January 2007 through December 2008, 14% of captive transactions occur with APRs less than 2.5%, while only 0.005% of bank transactions occur below that rate. Using 2.5% as an upper bound for subvented APR is conservative - as indicated, prime rates exceeded 2.5% for the entire period and some captive finance companies are known to subvent on loans to consumers with higher credit risk, leading to subvented APRs of 7%.

Figure II : Histograms of APRs by Lender Type

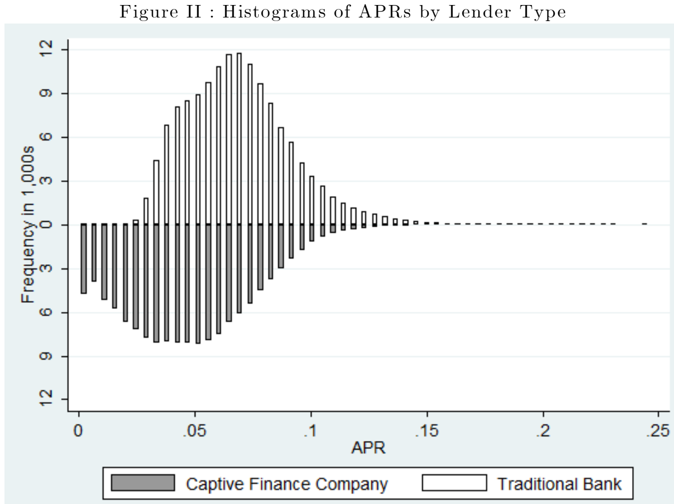

Any conclusions drawn from the above histograms would be spurious were there to be selection on unobserved risk. This could occur for instance, if captive finance companies were to retain all of the lowest credit risk consumers for their own loan portfolio, leaving riskier borrowers for traditional banks. The data provide means by which to ensure against this possibility. The transaction and loan origination data are linked to zip level Census estimates for consumer demographics, including median income. Examining median income by APR range and across lender types yields Table I.

Table I : Interest Rate Charged vs. Consumer Income by Lender Type

APR	Bank	Captive	95%
0-2.5%	$79,656	$76,418	Y
2.5-5%	$73,825	$72,030	Y
5-7.5%	$70,932	$68,851	Y
7.5-10%	$66,007	$65,568	Y
10-12.5%	$61,011	$62,738	Y
12.5-15%	$58,029	$59,589	N
>15%	$56,129	$58,288	N

Treating income as a proxy for credit risk, it is reasonable to assume that risk is decreasing in income. For a given column - ‘Bank’ or ‘Captive,’ APRs and risk move in the expected direction; as income decreases, risk increases as does APR. Comparing the ‘Bank’ and ‘Captive’ columns, especially for the first for rows, a more interesting trend emerges. Captive finance companies are willing to charge lower interest rates to riskier consumers, relative to their traditional bank competitors. To ensure that differences across lenders are statistically significant, I run unpaired two-sample t-tests assuming unequal variances. The null hypothesis is that differences between row entries are equal to zero. We reject the null for the top five rows. Note, only the direction and significance of the top four rows is meaningful for our purposes. Taken as a group, carmakers tend not to subsidize interest rates to the highest credit risk consumers (with APRs greater than 10%).

To present the results in another form, I regress observed APRs on my income measure, an indicator for whether the lender is captive, and a term interacting income with the captive indicator. Across specifications, I layer in quarter and market fixed effects, as well as my non-income demographic measures. Without attaching a causal interpretation to the regression, the precisely estimated coefficient on the captive dummy indicates that loans originated by captive lenders tend to have 2% lower APR, conditional on income/risk.

Table II : APR Regression Results

	Model 1		Model 2		Model 3	
	Coefficient	SE	Coefficient	SE	Coefficient	SE
Covariate						
Constant	0.100	0.0003	0.115	0.0003	0.109	0.0020
Captive	-0.022	0.0004	-0.020	0.0004	-0.019	0.0003
Income	-4.56e-07	4.3e-09	-4.8e-07	4.2e-09	-1.2e-07	6.1e-09
Captive*Income	1.59e-08	6.1e-09	-5.7e-09	5.3e-09	-2.4e-08	5.2e-09
Fixed Effects & Demographics						
Market Fixed Effects	N		Y		Y	
Quarter Fixed Effects	N		Y		Y	
Demographics	N		N		Y	
R^2	0.22		0.41		0.45	
N	278364		278364		276855	

Couching the result in a practical example, assuming a vehicle were to cost \$25,000 net down payment with a loan term of 60 months, and a bank APR of 6% (thus a captive APR of 4%), the total bank loan would be \$29,000 and the total captive loan would be \$27,625. While obtaining the loan through the captive would represent total savings to the consumer, the carmaker treats the \$1,375 as a cost to be subtracted from the vehicle markup. In this light, we see subvention as a costly strategy to move inventory. Subvention is infeasible if captives cannot issue loans (if borrowing costs increase dramatically) or if carmakers have insufficient cash on hand (to pay the spread between the market APR and subvented APR).

1.2.3 Captive Subprime Originations

One manner in which a carmaker's captive may generate considerable profit and boost sales is via extending loans to subprime borrowers. GM's former captive, GMAC, was known to originate a substantial amount of subprime auto loans. In the years leading up to 2008, of GMAC's total auto loan originations over 50% were to sub or near prime borrowers.[9]

GMAC's lending response to the 2008 credit crunch is particularly enlightening. GMAC - like other captive lenders - leverages itself to then extend loans to consumers. The primary sources of financing for captives are the auto asset backed securities market and the market for commercial paper. Both seized up and essentially froze in the fourth quarter of 2008.

[9]Nussel, Philip & Donna Harris. "GMAC auto finance business loses \$1.31 billion." *Automotive News,* February 3, 2009. Retrieved November 21, 2012.

In response to tight credit market conditions, GMAC tightened underwriting standards and ceased lending to anyone with a FICO score below 700. FICO scores range from 0-999, but most fall within 300-800. A FICO score of 640-660 is typically considered 'subprime.' As we see in point A in Figure III, cutting off 700's and below had a significant impact on GMAC's loan originations to consumers of GM vehicles. The drop from third quarter 2008 to fourth quarter 2008 represents a 92% decline in originations. Comparing year on year dollar amounts, this translates into the difference between $13.4 billion in fourth quarter 2007 originations versus $2.7 billion in fourth quarter 2008 loan originations.

Figure III : GMAC APR with Subprime Events

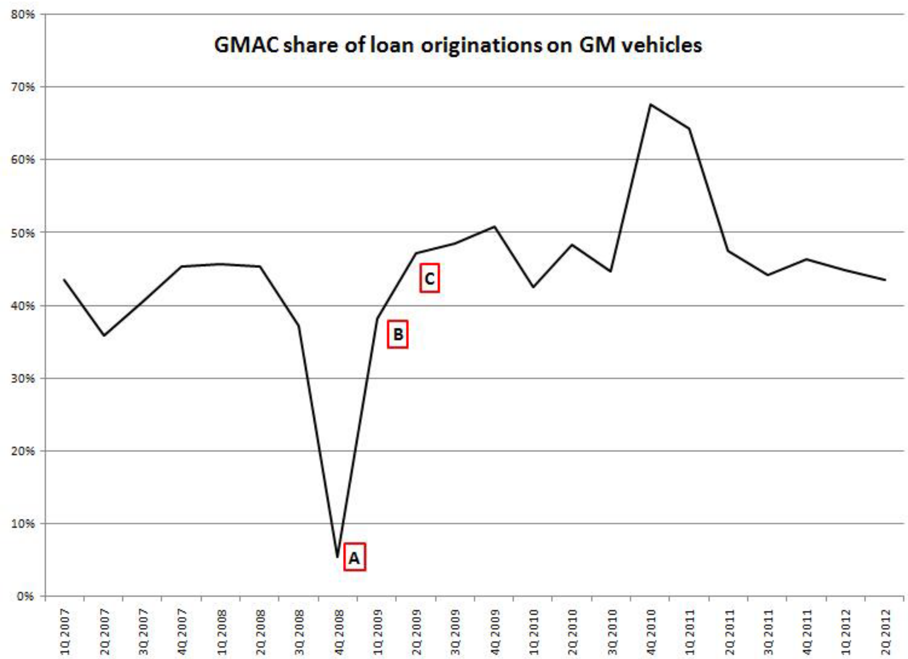

Next, after receiving a $6 billion loan from the TARP fund in December 2008 (before first quarter 2009), GMAC decided to allow FICOs of 620 or above. In other words, GMAC was now extending credit to subprime borrowers. Finally, in April 2009, GMAC received another $7.5 billion to support lending operations.[10] At this point in time, GMAC lifted its previous restrictions on FICO scores.

This decline in originations which coincided with the credit crunch and with announcements of tighter underwriting standards was not specific to GM and GMAC. From the third quarter to the fourth quarter of 2008, there were declines in originations from captives across the major manufacturers; Chrysler (-80%), Ford

[10] Warren et al (2010).

(-37%), GM (-92%), Honda (-25%), Nissan (-41%), and Toyota (-18%). These declines serve to highlight the importance of subprime originations to captive loan portfolios and to manufacturer sales. Furthermore, by virtue of the fact that we do not see declines of a similar scale in the level of bank originations provides evidence for the fact that captives and banks face different cost structures.

1.3 Model

I follow the standard modelling approach for differentiated demand systems delineated in Berry (1994), Berry, Levinsohn and Pakes (1995) - henceforth BLP - and Nevo (2000a & b) and subsequent literature.[11] Demand is given by a logit discrete choice model. Consumers make static purchase decisions over bundles of loans and cars. The decisions of simulated individuals are aggregated and matched to data on market shares in order to recover demand primitives. There are two types of firms - multiple carmakers and a single bank. Carmakers are comprised of two divisions - a captive finance division and a production & sales division. The bank's sole purpose is to lend to consumers. Carmakers make optimal choices over both vehicle prices and interest rates on loans originated by their captives. This joint optimization explicitly captures carmakers' incentive to subvent interest rates on loans. The bank optimizes over the interest rates that it charges on its loans. The first order conditions of carmakers and the bank can be inverted to yield marginal costs in the spirit of Nevo (2000 b). Projecting these marginal costs onto control variables yields supply relations for both firm types. Combining demand primitives and the estimated supply relation, I re-solve for new sets of prices and interest rates under alternative counterfactual scenarios.

1.3.1 Demand

Consumers make discrete choices over vehicle-loan combinations to maximize their utility. Goods are projected onto characteristics space. Each good is characterized by a vehicle price (p_{jmt}) ; a vector of physical characteristics of the vehicle - horsepower divided by weight, size, and miles per gallon (X_{jt}); loan contracts - monthly interest (z_{jlmt}) , loan term (T_{jlmt}) , and down payment (D_{jlmt}); as well as an unobserved product

[11] In ongoing work, I estimate the demand model with a random coefficient logit discrete choice model, as in BLP. In the current analysis, I employ a no heterogeneity version of the model.

characteristic ξ_{jlmt}. Consumers are characterized by their idiosyncratic preferences for bundles - their type one extreme value draw for each car-loan combination (ϵ_{ijlmt}).

Formally, consumer i in geographic market m in time t receives the following utility from a combination of vehicle j from lender l,

$$u_{ijlmt} = X_{jt}\beta + \alpha L_{jlmt} + \xi_{jlmt} + \epsilon_{ijlmt}$$

Total loan (L_{jlmt}) is the sum of monthly payments over the life of the loan.

$$L_{jlmt} = T_{jlmt}\left(\frac{z_{jlmt}\left(p_{jlmt} - D_{jlmt}\right)}{1-\left(1+z_{jlmt}\right)^{-T_{jlmt}}}\right)$$

Normalizing the value of the outside good to zero, with some abuse of notation, we obtain the following expression for bundle shares for market m in time t,

$$s_{jlmt} = \frac{\exp\left(X_{jt}\beta + \alpha L_{jlmt} + \xi_{jlmt}\right)}{1+\sum_{j,l}\exp\left(X_{jt}\beta + \alpha L_{jlmt} + \xi_{jlmt}\right)}$$

1.3.2 Supply

Carmakers and the bank compete à la Bertrand, in prices and in interest rates. Carmakers are indexed by $f = 1..F$ and the bank is indexed by B. Carmakers are comprised of a manufacturing & sales division as well as a captive finance division. Each carmaker's manufacturing & sales division produces and sells vehicles from a a subset J_f of the $j = 1...J$ different vehicles. In addition, carmaker f's captive originates loans only on products it also produces. The bank B originates loans on all vehicles $j = 1...J$. Carmaker $f's$ total variable profits are given by the following expression,

$$\Pi_f = \underbrace{\sum_{j\in J_f}\{Ms_j(p_j - mc_j^f)\}}_{\text{manufacturing \& sales profits}} + \underbrace{\sum_{j\in J_f}\{Ms_{jf}\left(L\left(T_{jf}, z_{jf}, p_j\right) - p_j + D_{jf}\right) - vc_f\left(Ms_{jf}\right)\}}_{\text{captive finance profits}} \tag{1.1}$$

Market size is given by M. Profits are maximized independently over markets and time, so I drop the market and time subscripts. The first term in Π_f captures carmaker f's manufacturing & sales profits.

These profits are a function of vehicle price p_j, constant marginal cost of production mc_j^f, and s_j. [12] The latter is the sum of the share of j with loans originated by carmaker f's captive plus the share of j with loans originated by bank B. The second term in Π_f represents carmaker f's captive finance profits. The captive effectively buys cars from the manufacturing & sales division at price p_j, receives down payments D_j and then generates revenues from loans $L(T_{jf}, z_{jf}, p_j)$. The share of originations s_{jft} are necessarily smaller than total product shares s_j. In addition, the captive finance division faces non-constant marginal costs which will be discussed in detail after introducing the bank's profit function. Note first that the specification of carmakers' profits explicitly accounts for the carmaker's incentive to subvent interest rates.

The bank's total variable profits are given by,

$$\Pi_B = \sum_j M s_{jB} \left(L\left(T_{jB}, z_{jB}, p_j\right) - p_j + D_{jB} \right) - vc_B \left(M s_{jB} \right) \tag{1.2}$$

The bank's profit function is nearly identical to the captive finance component of carmaker f's profit function. There are two important differences to note. First, banks cannot profitably lend at below marginal cost, as they do not have the second margin available to carmarkers. Second, the bank considers lending decisions over all vehicles in the choice set.

Returning to marginal costs of lending, I assume that both captives and the bank face non-constant marginal costs. This assumption is motivated by two empirical regularities; we observe subvention and we observe both captives and banks lending to consumers of a given vehicle. Assume for the moment that marginal costs of lending are constant. If a captive's marginal cost is weakly less than the bank's marginal cost, then the captive will obtain the entire market for originations of that carmaker's vehicles. If a captive's marginal cost exceeds that of the banks, it may or may not subvent, depending on the relation between product markups, demand elasticities, and the magnitude of the difference in marginal costs across lenders. In any event, with a captive's marginal costs exceeding those of the bank, either the captive will subvent and originate all loans or the bank will originate all loans. The assumption of constant marginal costs is empirically invalidated.

[12] Constant marginal costs of auto production are assumed by Bresnahan (1987) and Goldberg (1995). Berry, Levinsohn & Pakes (1995) examine alternative specifications for marginal cost, including non-constant specifcations. Beresteanu & Li (2011) test and cannot reject a constant marginal cost specification.

Profits are maximized independently over markets and time. Carmakers optimize over interest rates and vehicle prices; the bank over interest rates. Full derivations are reserved for the appendix. It is however enlightening to compare the captives' first order conditions with respect to interest rates with those of the bank. First, define the following markups; vehicle j product markup $\left(\Delta_j^V\right)$, captive lending on vehicle j by firm f markup $\left(\Delta_{jf}^L\left(Ms_{jf}\right)\right)$, and bank lending on vehicle j by the bank, B $\left(\Delta_{jB}^L\left(Ms_{jB}\right)\right)$.

$$
\begin{aligned}
\Delta_j^V &\equiv p_j - mc_j \\
\Delta_{jf}^L\left(Ms_{jf}\right) &\equiv L\left(T_{jf}, z_{jf}, p_j\right) - p_j + D_{jf} - mc_f\left(Ms_{jf}\right) \\
\Delta_{jB}^L\left(Ms_{jB}\right) &\equiv L\left(T_{jB}, z_{jB}, p_j\right) - p_j + D_{jB} - mc_B\left(Ms_{jB}\right)
\end{aligned}
$$

Define the square matrices $\Omega_{B,}\left(z;p\right)$ for the bank and $\Omega_C\left(z;p\right)$ for all captives, as in Nevo (2000 b).These matrices consist of partial derivatives of shares of bank or captive originations, with respect to interest rates. Since the bank lends to consumers of all vehicles, the j^{th} row and r^{th} column entry of $\Omega_B\left(z;p\right)$ equals $\frac{ds_{rB}}{dz_{jB}}$. Captives on the other hand only set interest rates for the vehicles their parent produces. The j^{th} row and r^{th} column entry in $\Omega_C\left(z;p\right)$ is given by,

$$
\Omega_{C,r,j}\left(z;p\right) = \begin{cases} \frac{ds_{rf}}{dz_{jf}} & \text{if } r, j \in J_f \\ 0 & \text{otherwise} \end{cases}
$$

Finally, define $\frac{dL_C}{dz_C}, \frac{dL_B}{dz_B}$ as diagonal matrices with j^{th} row, j^{th} column entries equal to derivatives of total loan with respect to the interest rate, for captives and the bank respectively. Let s_C and s_B vectors of all shares of captive and bank originations, respectively. Then, after stacking captive first order conditions for interest rates and manipulating, we arrive at the following expression,

$$
\Delta^V + \Delta_C^L\left(Ms_C\right) = -\Omega_C\left(z;p\right)^{-1}\frac{dL_C}{dz_C}s_C
$$

Were we to divide element by element by the vector of interest rates, we would have a multiproduct firm version of Lerner Index for captives. In contrast, the expression for the bank is,

$$\Delta_B^L (Ms_B) = -\Omega_{B,} (z;p)^{-1} \frac{dL_B}{dz_B} s_B$$

The first order conditions of the captive finance company and the bank capture the fundamental empirical difference between the two lenders. The bank's only considerations in setting interest rates are its marginal costs of lending and consumer responsiveness to changes in interest rates. Captives' interest rate setting decisions differ in that their interest rate setting also factors in their vehicle markups. Carmakers and banks differ; carmakers *total* markup relates to consumers' interest rate elasticities. As they only lend, the bank's lending markup relates to consumers' interest rate elasticities.

The first order conditions are stacked into a system of equations relating prices, shares, interest rates, marginal costs of lending, marginal costs of production and estimated demand parameters. In a manner similar to Nevo (2000 b), this system of equations is used to invert out marginal costs. The results of the inversion are also provided in the appendix. These marginal costs are projected onto control variables to obtain supply relations for vehicles and loans. With the demand estimates and estimated supply relations in hand, I conduct counterfactual experiments.

1.4 Data

The complete dataset for analysis is comprised of four individual sources of data. I combine data from a large subset of vehicle dealers in the US, Ward's automotive data, demographic data from the Census, and mortgage delinquency data from TransUnion.

The dealer data are from a major automotive consulting company and comprise a 20% sample of all new vehicle dealerships in the US. All new car transactions from a given dealer in the sample are recorded in the data. The data available for use in this study are disaggregated but not individual transaction level data. The final dataset to be used for demand analysis is comprised of 213,192 new vehicle - lender type - geographic market - calendar time observations.

Table III : Geographic Markets Represented in the Data

Atlanta, GA	Miami, FL
Baltimore, MD / Washington DC	Milwaukee, WI
Boston, MA	Minneapolis, MN
Northern CA	NV
Southern CA	New York City, NY
Charlotte, NC	Norfolk/VA Beach, VA
Chicago, IL	OK
Cincinnati, OH	Orlando, FL
Cleveland, OH	Philadelphia, PA
Columbus, OH	Phoenix, AZ
Dallas/Ft. Worth, TX	Pittsburgh, PA
Denver, CO	San Antonio, TX
Detroit, MI	Seattle, WA / Portland, OR
Houston, TX	St. Louis, MO
Indianapolis, IN	Tampa, FL
Kansas City, MO	TN

Table III indicates the thirty-two geographic markets represented in the data, on a quarterly basis from first quarter 2007 through third quarter 2012. Variation over geographic markets and calendar time is used in three distinct ways. First, variation across markets and time induces choice set variation, which identifies key demand parameters. There are a total number of 310 new vehicle models present in the full sample, yet at the market level the range of models offered is from 145 to 298 models and at the quarterly level the range spans 168 to 217 models. Second, since I observe a measure of negotiated final vehicle price - not sticker price - the data yield significant price variation over models, markets and time. This price variation will be used to identify the price elasticity of demand for vehicles. Third, lenders view different geographies as representing different average levels of risk . It is not uncommon for two consumers with identical credit scores to obtain different loan terms across geographies. This fact will manifest itself in different annual percent rates (APRs), loan term (in months) and down payments across markets and time. This variation aids in identifying the interest rate elasticity of demand for credit.

Table IV : Summary Statistics for Key Demand Model Variables

Var	Mean	Std. Dev.	Min	Max
APR	0.059	0.026	0	0.223
Term	64.5	5.96	22	89
Total Down	6239.5	3893.75	-7179	54537
Percent Down	0.21	0.09	-0.21	0.94
Price	28674.4	9609.9	10254	108141
Share	0.0035	0.0059	6e-05	0.1192

In the demand analysis, new vehicles are projected onto their physical characteristics, from Ward's Automotive Handbook. Following standard practice in the analysis of differentiated demand for cars (e.g., Berry, Levinsohn & Pakes (1995), Beresteanu & Li (2011)), I use horsepower divided by weight, miles per gallon and vehicle size (length x width) as physical characteristics. There are a total of 310 new vehicle models represented, representing 24 brands produced by 13 manufacturers.

The dealer data are linked to Census estimates for zip level demographics. A unit of observation in my data is an average of variables over transactions in this cell. For example, an observation in my data may be the total number of Ford F-150 transactions which occurred in the third quarter of 2010, received financing from Ford's captive (Ford Credit), and occurred in Houston, TX. In point of fact, each of the individual transactions which contribute to this total number of transactions are linked to consumers who reside in zip codes in and around Houston. The Census estimates provided for this observation are then weighted averages of the zip code-level Census estimates for each of these zip codes. The Census variables that I employ are education (percentage in education levels ranging from High School through Graduate School); ethnicity (percentage African-American, Asian, Caucasian, Latino, Other); marital status (percentage single, married); median household income and occupation (percentage Blue Collar, Service, White Collar). In addition, the transaction data include measures for consumer age and gender. At the aggregated level that I use in my analysis, I observe average age and percentage of females.

To instrument for endogenous interest rates, I use 90+ day mortgage delinquency data from TransUnion. These data are provided via the New York Federal Reserve. The data are for county level delinquencies, for the fourth quarter of each year, 2006-2011. I link these county level data to the appropriate level of geographic aggregation implied by dealer transaction data.

1.5 Estimation

I estimate multiple specifications of the demand model. These models are tractable given the size of the data and capture the salient feature of demand in this market - that the demand for cars and the demand for loans are determined jointly. The following form for consumer utility yields the preferred specification for estimation.

$$u_{ijlmt} = X_{jt}\beta + \alpha L_{jlmt} + \xi_{jlmt} + \epsilon_{ijlmt}$$

Physical characteristics of the vehicle (X_{jt}) are observed at the model year level for each vehicle. Total loan is the sum of monthly payments over the life of the loan.

$$L_{jlmt} = T_{jlmt}\left(\frac{z_{jlmt}\left(p_{jlmt} - D_{jlmt}\right)}{1 - \left(1 + z_{jlmt}\right)^{-T_{jlmt}}}\right)$$

The unit of observation for loan characteristics (monthly interest rate, loan term, vehicle price and down payment) is at the level of vehicle- lender type - geographic market - calendar time. Following Berry (1994), I assume that the ϵ_{ijlmt} are iid type one extreme value errors. Integrating over these errors and normalizing the value of the outside option to equal zero, I obtain the following estimating equation which is linear in the parameters. This specification provides the basis for OLS and IV regressions. When estimating the model, I include geographic market and time fixed effects.

$$\ln\left(\frac{s_{jlmt}}{s_{0mt}}\right) = X_{jt}\beta + \alpha L_{jlmt} + \xi_{jlmt} \tag{1.3}$$

Estimating equation (1.3) suffers from endogeneity due to simultaneity. Interest rates and total demand for loans are determined simultaneously. Vehicle prices and the total demand for vehicles are similarly simultaneously determined. To control for the latter, I use standard instruments as outlined in Nevo (2000a), Hausman, Leonard & Zona (1994) and Hausman (1996). To control for the enogeneity of price of vehicle j in geographic market m at calendar time t, I use prices for vehicle j in other markets in the same time period. The data are sufficiently rich to allow the use of prices for the same good across nine geographic markets.

Prices for the same good in other markets are valid instruments under the assumption that demand shocks for the same good are uncorrelated across markets. There are two clear challenges to this assumption over the period that the data span - the Great Recession of 2008-2009 and the period of elevated gas prices in 2008. Unemployment rates and housing prices - proxies for the extent of the recession - were differentially impacted across markets and time. Across markets, gasoline prices share a trend but differ in fixed proportions. Thus, conditional on time and market fixed effects, it seems plausible to assume that demand shocks will be uncorrelated across markets and time even in the face of the Great Recession of 2008-2009 and the gas price shock of 2008.

To instrument for endogenous interest rates, I employ two alternative cost shifter arguments. For the first, I note that lenders choose interest rate offers for consumers based in large part on consumers' FICO scores. Officially, the two dominant factors in determining consumer FICO scores are payment history and debt outstanding, which respectively contribute 35% and 30% to FICO score determination.[13] Variation in these two FICO determinants serve as valid shifters to the cost of lending as they vary the risk associated with lending. As a proxy for payment history and debt outstanding, I use lagged 90+ day mortgage delinquency data from TransUnion across markets and time.

For the second cost shifter argument, I note that captives raise funds to issue new loans by selling commercial paper and by bundling existing loans into asset backed securities. The 2008 credit crunch all but shut down the commercial paper and ABS markets. Prior to the credit crisis, the non-mortgage ABS market operated smoothly. It did not experience the same boom in new issuance exhibited in mortgage backed securitization.[14] In words, the 2008 credit crunch is plausibly exogenous to the credit markets in which captive finance companies engage.

The assumption that the ϵ_{ijlmt} follow a type one extreme value distribution is arguably strong. Implications of this assumption on implied substitution patterns are discussed in the results section. This assumption does however considerably simplify the identification of model parameters. The identification argument is equivalent to that for the identification of means of random coefficient distributions in the random coefficient model, as noted in BLP and Gowrisinkaran & Rysman (2012). To identify the the coefficients

[13] FICO, *What's in my FICO Score.*
[14] Ashcraft, Malz, & Poszar (2012).

(α, β), it is necessary to observe considerable variation in shares in response to changes in levels of characteristics associated with the parameters (α, β). Given that I observe variation in negotiated vehicle price, monthly interest, loan term and down payment across markets and time I am reasonably confident that the α parameter is identified. Further, I observe choice set variation across markets and time, and changes in vehicle characteristics across time. This variation should suffice to identify the vector β.

Given the estimated demand parameters and the marginal costs of lending and production can be obtained by solving the system of first order conditions implied by the supply side model. Marginal costs of production are assumed to be constant, and marginal costs of lending are assumed non-constant. I do not project constant marginal costs of production onto control variables, in keeping with Nevo (2000b).

The marginal costs of lending are projected onto control variables and levels of sales originations (Q_{jlmt}), noting that $Q_{jlmt} = Ms_{jlmt}$. The preferred specification for marginal costs of lending is,

$$mc_l\left(Q_{jlmt}\right) = \gamma_0 + \gamma_1 1\left(l \neq B\right) + \gamma_2 1\left(l \neq B\right) Q_{jlmt} + \gamma_3 Q_{jlmt} + \gamma_4 MD_{mt} + \omega_{jlmt}$$

where l indexes whether the firm in question is the bank or a captive lender, $1\left(l \neq B\right)$ is a dummy indicating whether the firm is captive or the bank, and MD_{mt} is defined as the 90+ day mortgage delinquency for a given market-time unit. The preferred specification includes both market and time fixed effects.

Since marginal costs are assumed non-constant, I face enogeneity due to simultaneity. I instrument Q_{jlmt} with demand shifters - average Census demographics at the market-time unit of analysis.[15] I omitted market and time fixed effects from the main marginal cost specification and used these instead as instruments - interpreting the fixed effects as demand shifters. The results are qualitatively similar to the current specification. Having time and market fixed effects in the main equation for marginal costs is however preferred for the purposes of counterfactual simulation.

[15] While the preferred demand specification omits demographics, this inconsistency will be reconciled with the estimation of the random coefficient logit model (in progress).

1.6 Estimation Results

The table below provides results for the OLS-Logit and IV-Logit demand models. There are two IV models - one uses mortgage delinquency rates as instruments for interest rates, the other uses the 2008 credit crunch to instrument for interest rates. Across both IV specifications, I use the same instruments for price - prices for a given good in nine different markets. First, note that the estimate on 'Total Loan' is precisely estimated, and is comparable across specifications. This is reassuring as this parameter is a key ingredient of my implied demand elasticities.

Table V : OLS-Logit, IV-Logit Demand Estimation Results

	Coefficient	SE	Coefficient	SE	Coefficient	SE
	OLS		IV - Delinquency		IV - Credit Crunch	
Variable						
Total Loan	-3.8E-05	4.4E-07	-4.8E-05	6.2E-07	-3.5E-05	2.2E-06
Horsepower/Weight	-1.35	0.26	0.74	0.27	-2.49	1.04
Miles/Gallon	1.4E-02	6.0E-04	1.4E-02	6.2E-04	2.0E-02	2.4E-03
Size (in^2)	9.6E-05	1.9E-06	1.2E-04	2.1E-06	1.0E-04	7.8E-06
Fixed Effects						
Market Fixed Effects	Y		Y		Y	
Quarter Fixed Effects	Y		Y		N	
R^2	0.25		0.25		0.24	
N	213192		212226		16573	

Second, note that instrumenting ensures that the sign of all estimated parameters moves in the anticipated direction for the delinquency specification. Intuitively, we expect utility from a vehicle to be increasing in its physical characteristics and decreasing in its effective price (total loan). The credit crunch specification yields a negative coefficient on horsepower / weight and a larger coefficient on miles per gallon. This is consistent with the fact that for this specification only the third and fourth quarters of 2008 are used as data. National average retail gasoline prices peaked at over \$4/gallon in the third quarter of 2008 and only began to abate late in the fourth quarter. Vehicles with high horsepower/weight tend to be sports cars and SUVs and trucks with powerful engines. Since there is less of a premium for such vehicles when gas prices are high, we should not be surprised by the negative coefficient in spite of instrumenting.

As a robustness check for my demand specification, I estimate the model allowing for separate coefficients on vehicle price, monthly interest rate, and term. Across specifications, I obtain estimated parameters that yield comparable price and interest rate elasticities. I omit these results from the current discussion, but present the findings in Appendix C. In what follows, the IV-Delinquency model is the preferred specification.

The primary focus of the demand estimation exercise is to obtain estimates for parameters which will characterize substitution patterns in the markets for loans and for vehicles. Logit demand systems are known to suffer from potentially unrealistic substitution patterns due to the fact that there is no correlation in consumer preferences across goods. This shortcoming of logit demand systems served as a motivating factor in introducing random coefficients into consumer preferences in BLP and the literature that followed. The elasticities implied by my logit estimates are still informative and will capture features of the average response of consumers to price and interest rate changes in these markets.

In the table below, I indicate summary statistics of the own price elasticity of demand for vehicles (formulas in Appendix B). An observation uses data for a given vehicle-lender-market-time period. The data in question are APR, down payment, price and term. The final column of the table indicates the percent of these implied elasticities which are greater than one in absolute value. The rows represent results from the OLS estimation (OLS) and the IV estimation with mortgage delinquency as an instrument (IV).

Table VI : Price Elasticities of Vehicle Demand Implied by Demand Estimates

	Mean	Std. Dev.	Min	Max	% Elastic
OLS	-1.27	0.4	-4.78	-0.42	72%
IV	-1.61	0.51	-6.06	-0.54	93%

When moving from OLS to IV, we increase the proportion of price elasticities which exceed one in absolute value. We would expect that consumers are price-responsive in their demand for vehicles, so it is reassuring that on average consumers are price elastic and that nearly all observations are associated with elasticities greater than one in absolute value.

Turning to interest rate elasticities of loan demand, it would appear that consumers are on average inelastic with respect to interest rates. In this setting, I interpret interest rate inelasticity as consumers being willing to accept interest rates over the support of APRs in the data. If banks and captives are lending, consumers are borrowing.

Table VII : Interest Rate Elasticities of Loan Demand Implied by Demand Estimates

	Mean	Cap-Mean	Bank-Mean	Std. Dev.	Min	Max
OLS	-0.15	-0.12	-0.18	0.08	-1.03	-6e-05
IV	-0.29	-0.15	-0.23	0.11	-1.3	-7e-05

In point of fact, we know that consumers face considerably less choice over interest rates than implied in the IV-logit version of demand. Consumers' individual-specific credit risk determines the subset of interest rates for which they are eligible. In this setting, there is a subset and not a single set of interest rates given that subvention implies that the bank and the set of captives will charge different APRs. Within the set of captive APRs, there may also exist variation at the consumer level, as captives do not perfectly coordinate timing of special finance offers.

Finally, I provide the results for the estimation of marginal costs. Constant marginal costs of production and non-constant marginal costs of lending are obtained via inverting the system of first order conditions implied by the supply model. Only the non-constant marginal costs of lending are projected onto control variates. I use the constant marginal costs of production obtained via the inversion of first order conditions as data, as suggested in Nevo (2000 b). As the marginal costs of lending vary with total originations, they suffer from simultaneity bias. I instrument with demand shifters - consumer demographics as described in the data section. Instrumenting changes the direction and magnitude of the coefficient on total originations. Instrumenting also changes the sign on mortgage delinquency to the expected direction. With mortgage delinquency representing risk, we would expect an increase in the delinquency rate to increase marginal costs of lending.

Table VIII : Non-Constant Marginal Costs of Lending Estimation Results

	Coefficient	SE	Coefficient	SE
	OLS		IV	
Variable				
Constant	-15265.96	589.96	-20110.19	196.20
Quantity	-1.94	17.74	100.54	3.77
Captive Indicator	15754.99	1696.62	20317.37	167.64
Captive*Quantity	-0.45	3.64	-130.76	4.66
Mortgage Delinquency	-1971.47	170.95	2097.56	451.77
Fixed Effects				
Market Fixed Effects	Y		Y	
Quarter Fixed Effects	Y		Y	
R^2	0.95		0.66	
N	213192		213192	

The results also imply that the bank faces increasing marginal costs. Noting that the slope on marginal costs for captive lenders is the sum of the coefficient on quantity and the coefficient on the interaction of quantity with a captive indicator, captives face decreasing marginal costs. A two-tailed t-test rejects the null of constant marginal costs of lending for captives at the 99% level. A single-tail t-test rejects the null of increasing marginal costs of lending for captives at the 99% level.

Decreasing marginal costs of lending for captives may be attributable to the fact that the marginal cost of lending incorporates the possibility that borrowers will default on their car loan and lenders will subsequently repossess vehicles. Captive finance companies have a competitive advantage in the remarketing of these repossessed vehicles, effectively facing a lower marginal cost of lending. Alternatively, decreasing marginal costs may capture the fact that captives tend to rely more heavily than banks on asset backed securities to raise funds to issue new loans.

1.7 Counterfactual

To gain insight into how captive finance affects market outcomes, I simulate and evaluate market outcomes in the absence of captive finance. I re-solve for prices and APRs assuming that car manufacturers no longer control their captive finance companies. I then evaluate the difference in prices and APRs between actual and counterfactual scenarios. I also quantify how consumer and producer surpluses differ between actual and counterfactual outcomes. This exercise allows me to examine how captive finance affects competition among manufacturers and banks. The results also contribute to the discussion surrounding new Consumer Financial Protection Bureau regulation of dealer lending markups. Re-solving for prices and interest rates across different geographies – which proxy for different risk categories – I can also speak to how captive finance impacts risk in the auto lending market.

In order to conduct my counterfactual experiment, I need to address a technical detail regarding the appropriate marginal costs of vehicle sales and production to be employed. The profit function specification for carmakers in (1.1) explicitly assumes that product pricing and interest rates are jointly determined by the carmaker. This assumption allows for the possibility of subvention to occur in equilibrium. Carmaker marginal costs of lending and production obtained from the FOCs for profit-maximization will embed the implicit trade-off to carmakers of shifting profits across their product and lending divisions. To evaluate the counterfactual in which carmakers no longer retain their captives, I re-specify the model - the marginal costs are incorrect given that all supply-side actors are now independent. In the alternative model, production and sales divisions set prices, now-independent captives set interest rates, and the bank's problem is unchanged. Inverting the FOCs implied by this alternate model yields a second set of marginal costs.

I also need to make additional assumptions on the model. I assume that the now independent captive finance companies have the same lending cost function as traditional banks. Consistent with this assumption, I assume that the now independent captives offer loans with the same loan term and down payments as traditional banks. Additionally, I fix the choice set to be the same as in the actual data - car manufacturers do not respond to the loss of their captives by altering their physical product selection. I solve for prices and interest rates for a given market separately, for each time period. As firms are static decision-makers, there is no loss in generality associated with this choice. I re-solve the model for the Miami and Minneapolis markets to provide some contrast in terms of lending risk.

Table IX reports the primary findings from the counterfactual. The first two columns represent the average percent change in APRs and prices in the counterfactual relative to their actual levels in the data. I average over all simulated time periods. The third column presents the correlation between the percentage change in APRs and prices.

I find that vehicle prices decrease and APRs increase relative to actual outcomes. This statement is consistent across the markets which I investigate. While averages over time periods are reported in table IX, the result is consistent across time periods as well.

Table IX : Percent Change in APRs & Prices in Counterfactual Relative to Actual

	Change in APR (%)	Change in Price (%)	Correlation
Miami	15.8	-11.6	-0.35
Minneapolis	13.8	-10.1	-0.4

The increase in APRs is in accordance with intuition. Banks now face less competition, in particular, less competition from captives who previously subsidized interest rates on loans, further depressing APRs. The results imply that to recover profits lost by the spin-off of captive finance companies, car manufacturers will drop prices, stoking sales volume. It is reasonable to see both increasing APRs and decreasing prices, as is evident in the negative correlation between prices and APRs in table IX. This inverse correlation between prices and interest rates plays out in an interesting way at the brand level (see Appendix D). Brands often associated with higher credit risk consumers see the biggest increases in APRs and relatively large declines in price. The result is driven by the estimated demand parameters. Consumers are more responsive to changes in loan principal than to changes in interest rates.

Using the new optimal APRs and prices, I calculate how consumer and producer surplus changes between the actual and counterfactual outcomes. The measure of consumer surplus that I employ is the compensating variation which is defined as,

$$CV = \frac{M}{\alpha}\left(\ln\sum_{j,l}\exp\left(\delta^{C}_{j,l}\right) - \ln\sum_{j,l}\exp\left(\delta^{A}_{j,l}\right)\right)$$

Where $\delta^{C}_{j,l}$ is the mean utility from vehicle j with a loan from lender l in the counterfactual scenario. The term $\delta^{A}_{j,l}$ is similarly defined for the actual data. M represents the relevant market size, and α is the estimated disutility of expenditure parameter. This measure compares utility in the counterfactual scenario to the actual data and asks what additional income is necessary to maintain consumer utility at the pre-change level, given the new prices.

Table X presents compensating variation figures for Miami and Minneapolis. The table also presents the change in manufacturer, captive and bank profits attributable to the counterfactual. Profit figures are presented as differences. Estimates of actual profits are subtracted from profits at counterfactual prices and APRs. All reported figures are in millions of US dollars. Reported figures are averages over simulated time periods.

The interaction of higher APRs and lower prices in the counterfactuals leads to lower loan sizes. Thus, consumers would need to be compensated to return to the equilibrium evident in the actual data. In the case of consumers in Miami, they would be compensated by just over $2 billion dollars to remain indifferent between the actual and counterfactual outcomes. The comparable figure is $11 billion for Minneapolis. These figures seem rather large. I discuss potential drivers for these magnitudes at the end of this section.

Turning now to changes in producer surplus, the results for captives and banks are straightforward and intuitive. Mechanically, captive profits drop in the counterfactual scenario relative to the outcomes actually observed in the data: captives cease to operate in the counterfactual scenario. It would be surprising if bank profits did not at least weakly increase - banks assume all lending operations in the industry where they previously represented fringe lenders. The growth in bank profits is attenuated by their increasing marginal costs of lending. Manufacturer profits increase in the counterfactual scenario. This is driven by an interaction among the following: prices decrease, shares increase, and marginal costs of production have decreased after the link between finance and production is severed.

Table X : Change in Consumer and Producer Surplus Between Counterfactual and Actual ($Million)

	CV	Manufacturer Profits	Captive Profits	Bank Profits
Miami	2,200	73.8	-0.56	6.6
Minneapolis	11,400	74.3	-0.37	4.4

The results of the counterfactual shed some light on issues pertinent to new Consumer Financial Protection Bureau (CFPB) regulation of auto lenders and vehicle dealers. The CFPB seeks to regulate markups that dealers charge consumers for securing loans. The motivation for this regulation is by analogy to mortgage brokers who tack fees on to mortgages. Mortgage brokers differ fundamentally from car dealers - brokers cannot change house prices while dealers can. Therefore, any policy directed at restricting dealers' capacity to markup loans must then consider consequences on vehicle prices. In the counterfactual, I consider a slightly different yet informative scenario. There we see that APRs and prices are inversely correlated.

Dealers respond to higher APRs by lowering prices. A reasonable take-away is that dealers would respond to lower APRs - as part of CFPB regulation - by rasing vehicle prices. Given the results of demand estimation - that consumers value changes in loan principal more so than changes in APRs - there are potentially negative welfare consequences associated with regulating dealer lending markups.

I also use the counterfactual interest rates to evaluate whether competition between captives and banks leads to looser underwriting standards, especially at banks. The FDIC suspects looser underwriting standards as captive subvention may induce banks to accept riskier consumers or charge these consumers less in a bid to remain competitive. Here, I interpret underwriting standards as pricing conditional on credit risk. My mortgage delinquency and APR variables allow me to categorize regional markets at different points in time in terms of risk relative to each other. By examining changes in interest rates when I remove captives entirely, relative to the base case, I assess whether the bank's interest rates increase in risky areas and by how much. I isolate attention to one high-risk market (Miami) and one low-risk market (Minneapolis). I use my mortgage delinquency measure as a proxy for risk. With 90+ day mortgage delinquency rates averaging 16% of mortgaged single family homes, Miami is in the top 95th percentile of mortgage delinquency rates during my sample period (2007-2012). In contrast, with average delinquency rates of 3.7% of mortgaged single family homes, Minneapolis falls in the bottom 5th percentile of mortgage delinquency rates during my sample period.

Referring back to table IX, we see that the average increase in APRs is higher in Miami. Employing a one-sided unpaired two-sample t-test, I reject the null of equal average change in APRs across both markets at the 0.1 level.[16] That is, in average terms, APRs have increased by more so in the Miami. This leads me to suggest that competition between captive finance companies and traditional banks leads banks to under-price risk. Alternatively, banks price risk less aggressively when they compete against other traditional banks.

In summary, I find that removing captives leads to lower prices and higher APRs, relative to those in the actual market configuration. I also find that changes in prices and in APRs are inversely correlated, suggesting the Consumer Financial Protection Bureau should anticipate vehicle price increases should it aggressively regulate dealer lending markups. I also find that APRs increase more in relative terms in riskier

[16]There is a caveat in order: the variances used to construct the test statistic do not reflect variance transmitted from the original demand estimates, through to marginal costs, and through the counterfactual experiment.

markets when I remove captive finance companies from the game. This suggests that competition between captives and banks may lead banks to under-price risk.

The direction of change in APRs and in prices is certainly believable and is consistent with intuition. The magnitudes implied by my counterfactual - especially for consumer and producer surplus - are subject to some caveats. The results are partially driven by the main conclusion of my demand estimates - that consumers are more responsive to changes in loan principal than to changes in APRs. In turn, this result is at least slightly attributable to my decision to model consumers as static decision-makers. Consumers in my model pay the full cost of the loan at the time of purchase, rather than pay monthly payments over the term of the loan. This modelling decision was driven by (1) the absence of appropriate income data, (2) the absence of granular data on auto loan delinquency and default, and (3) the complexity associated with estimating a dynamic decision problem for consumers given the size of the choice set present in my data.

I also find that consumers are better off in the counterfactual scenarios, as measured by compensating variation. Firms are also more profitable in the counterfactual outcome. The magnitude of the compensating variation seems large. This is partly driven - again - by the fact that consumers are modelled as static decision-makers. All benefits of smaller loans in the counterfactual are realized immediately, versus being discounted over a longer time horizon. In actuality, we may also suspect that firms would adjust loan term and down payment. Allowing for this fact would change the entire equilibrium profile. Unfortunately, absent observing variation in loan terms over my sample, and without access to data on minimum down payment rules at lenders, I would not be able to estimate the effect of these features on demand let alone re-solve for an equilibrium which treated term and down payment as variables.

Finally, the manufacturer profit figures also appear large. There are considerable fixed costs of manufacturing that are unaccounted for in the profit calculations. Also, it is known that captives are typically more profitable than the sales and manufacturing units at their parent carmakers. Estimated profits of captives relative to manufacturers do not reflect this, likely due to the absence of manufacturer level cost data in my estimates.

1.8 Conclusion

This paper is the first to empirically investigate automotive captive finance companies. In particular, I investigate how manufacturers' use of captive finance companies affects competition and risk in the new vehicle and auto lending markets. The question is highly relevant in the wake of the US Auto Bailout. The US Treasury nearly let GMAC fail, and did let Chrysler Financial fail. Furthermore, both the Consumer Financial Protection Bureau and the Federal Deposit Insurance Corporation have taken an interest in aspects of auto lending and captive finance.

To investigate captive finance companies and their role in the new vehicle and auto lending markets, I specify and estimate an equilibrium model. This model explicitly accounts for two highly relevant features of captive finance: (1) carmakers use captive finance companies to subsidize interest rates on loans and (2) carmakers use their captives to extend loans to subprime borrowers. Model estimates imply that consumers are more responsive to changes in loan principal than to changes in interest rates.

Using estimates from the model, I conduct a counterfactual in which I remove captive finance companies from the game. The counterfactual allows me to evaluate how captives affect competition in this market. I find that captives affect competition by depressing interest rates and by propping up vehicle prices. By examining the relationship between changes in APRs and prices, I can make a statement relevant to the Consumer Financial Protection Bureau's pending regulation of auto lenders and car dealers. Given my finding that changes in prices and APRs are inversely correlated, I would conclude that the CFPB take into account the impact of new regulation on dealer lending markups on vehicle prices. By comparing interest rates set by banks across markets both before and after captives are removed, I am able to make a statement regarding how competition between the two types of lenders leads banks to under-price risk. I find that interest rates change more aggressively in riskier markets when I remove captive finance companies. This suggests that traditional banks under-price risk in the presence of competition with captive finance companies.

The results are not without qualification. The direction of the effect of captive finance on prices and APRs is accurate and intuitive. The magnitudes of consumer and producer surplus seem large and in some cases do not accord with intuition. Results follow from the demand estimates which indicate that consumers respond more to changes in loan principal than to changes in interest rates. This result is at least partly driven by the assumption that consumers pay the full amount of their loan at the time of purchase. A consumer model which captures richer dynamics and accounts for unobserved heterogeneity in income is likely to yield improvements in this regard.

1.9 Appendix

1.9.1 Solving for Markups (Captives Owned by Manufacturers)

With some abuse of notation, let t index a market - a combination of geographic market and time. Please note that the notation differs slightly from that in the main text. Carmaker product and lending profits are given by,

$$\Pi_f = \sum_{j\in J_f} Ms_{jt}\left(p_{jt} - mc_{jt}^f\right) + \sum_{j\in J_f} Ms_{jft}\left(L\left(T_{jft}, z_{jft}, p_{jt}\right) - p_{jt} + D_{jft}\right) - \sum_{j\in J_f} vc_f\left(Ms_{jft}\right)$$

Where total loan $L(\cdot)$,

$$L\left(T_{jft}, z_{jft}, p_{jt}\right) \equiv \frac{T_{jft} z_{jft} p_{jt}}{1 - \left(1 + z_{jft}\right)^{-T_{jft}}}$$

$$\begin{aligned}\frac{d\Pi_f}{dp_{jt}} &= Ms_{jt} + \sum_{r\in J_f} M\left(\frac{ds_{rft}}{dp_{jt}} + \frac{ds_{rBt}}{dp_{jt}}\right)\left(p_{rt} - mc_{rt}^f\right) + Ms_{jft}\left(\frac{dL\left(T_{jft}, z_{jft}, p_{jt}\right)}{dp_{jt}} - 1\right) + ... \\ & ... + \sum_{r\in J_f} M\frac{ds_{rft}}{dp_{jt}}\left(L\left(T_{rft}, z_{rft}, p_{rt}\right) - p_{rt}\right) - \sum_{r\in J_f} M\frac{ds_{rft}}{dp_{jt}} mc_f\left(Ms_{rft}\right)\end{aligned}$$

Define lending and vehicle markups,

$$\begin{aligned}\Delta_{rft}^L\left(Ms_{rft}\right) &\equiv L\left(T_{rft}, z_{rft}, p_{rt}\right) - p_{rt} + D_{jlt} - mc_f\left(Ms_{rft}\right) \\ \Delta_{rt}^V &\equiv p_{rt} - mc_{rt}^f\end{aligned}$$

Then (suppressing market size), and defining for simplicity

$$\begin{aligned}\frac{d\Pi_f}{dp_{jt}} &= s_{jt} + \sum_{r\in J_f}\left(\frac{ds_{rft}}{dp_{jt}} + \frac{ds_{rBt}}{dp_{jt}}\right)\Delta^V_{rt} + s_{jft}\left(\frac{dL\left(T_{jft}, z_{jft}, p_{jt}\right)}{dp_{jt}} - 1\right) + \sum_{r\in J_f}\frac{ds_{rft}}{dp_{jt}}\Delta^L_{rft}\left(Ms_{rft}\right)\\ &= s_{jBt} + \sum_{r\in J_f}\left(\frac{ds_{rft}}{dp_{jt}} + \frac{ds_{rBt}}{dp_{jt}}\right)\Delta^V_{rt} + s_{jft}\frac{dL\left(T_{jft}, z_{jft}, p_{jt}\right)}{dp_{jt}} + \sum_{r\in J_f}\frac{ds_{rft}}{dp_{jt}}\Delta^L_{rft}\left(Ms_{rft}\right)\end{aligned}$$

In vectorized form,

$$0 = \frac{d\Pi_f}{dp} = s_B + \left(\frac{ds_f}{dp} + \frac{ds_B}{dp}\right)\Delta^V + \frac{dL_f}{dp}s_f + \frac{ds_f}{dp}\Delta^L_f\left(Ms_f\right) \tag{1.4}$$

Where $\frac{dL_f}{dp}$ is a matrix with diagonal entries equal to $\frac{dL(T_{jft},z_{jft},p_{jt})}{dp_{jt}}$ and off diagonal entries equal to zero. Now, w.r.t. to interest rate,

$$\begin{aligned}\frac{d\Pi_f}{dz_{jft}} &= \sum_{r\in J_f} M\frac{ds_{rft}}{dz_{jft}}\left(p_{rt} - mc^f_{rt}\right) + \sum_{r\in J_f} M\frac{ds_{rmt}}{dz_{jft}}\left(L\left(T_{rft}, z_{rft}, p_{rt}\right) - p_{rt}\right) + ...\\ &\quad ... + Ms_{jft}\frac{dL\left(T_{jft}, z_{jft}, p_{jt}\right)}{dz_{jft}} - \sum_{r\in J_f} M\frac{ds_{rmt}}{dz_{jmt}}mc_f\left(Ms_{rmt}\right)\end{aligned}$$

I assume that in setting its own interest rates, the captive does not internalize the effects of changes in its own rates on carmaker sales with loans originated by the bank. In words, I assume that $\frac{ds_{jB}}{dz_{jl}} = 0$. Re-writing and suppressing market size,

$$\frac{d\Pi_f}{dz_{jft}} = \sum_{r\in J_f}\frac{ds_{rft}}{dz_{jft}}\Delta^V_{rt} + \sum_{r\in J_f}\frac{ds_{rft}}{dz_{jft}}\Delta^L_{rft}\left(Ms_{rft}\right) + s_{jft}\frac{dL\left(T_{jft}, z_{jft}, p_{jt}\right)}{dz_{jft}}$$

In vectorized form,

$$0 = \frac{d\Pi_f}{dz_f} = \frac{ds_f}{dz_f}\Delta^V + \frac{ds_f}{dz_f}\Delta^L_f\left(Ms_f\right) + \frac{dL_f}{dz_f}s_f$$

Assuming that $\frac{ds_f}{dz_f}$ is invertible,

$$\Delta^V = -\Delta^L_f\left(Ms_f\right) - \left(\frac{ds_f}{dz_f}\right)^{-1}\frac{dL_f}{dz_f}s_f \tag{1.5}$$

Or, manufacturer's vehicle markups are equal to their loan markups plus something reminiscent of an (inverse) interest rate elasticity of loan demand. Substituting into the result from the other FOC,

$$\begin{aligned}
0 &= s_B - \left(\frac{ds_f}{dp} + \frac{ds_B}{dp}\right)\left(\Delta_f^L\left(Ms_f\right) + \left(\frac{ds_f}{dz_f}\right)^{-1}\frac{dL_f}{dz_f}s_f\right) + \frac{dL_f}{dp}s_f + \frac{ds_f}{dp}\Delta_f^L\left(Ms_f\right) \\
&= s_B - \left(\frac{ds_f}{dp} + \frac{ds_B}{dp}\right)\left(\frac{ds_f}{dz_f}\right)^{-1}\frac{dL_f}{dz_f}s_f - \frac{ds_B}{dp}\Delta_f^L\left(Ms_f\right) + \frac{dL_f}{dp}s_f \\
&= s_B - \frac{ds_f}{dp}\left(\frac{ds_f}{dz_f}\right)^{-1}\frac{dL_f}{dz_f}s_f - \frac{ds_B}{dp}\left(\frac{ds_f}{dz_f}\right)^{-1}\frac{dL_f}{dz_f}s_f - \frac{ds_B}{dp}\Delta_f^L\left(Ms_f\right) + \frac{dL_f}{dp}s_f
\end{aligned}$$

Or,

$$\Delta_f^L\left(Ms_f\right) = \left(\frac{ds_B}{dp}\right)^{-1}s_B - \left(\frac{ds_B}{dp}\right)^{-1}\left(\frac{ds_f}{dp}\right)\left(\frac{ds_f}{dz_f}\right)^{-1}\frac{dL_f}{dz_f}s_f - \left(\frac{ds_f}{dz_f}\right)^{-1}\frac{dL_f}{dz_f}s_f + \left(\frac{ds_B}{dp}\right)^{-1}\frac{dL_f}{dp}s_f \tag{1.6}$$

Plugging (1.6) into (1.5) yields the final expression for vehicle markup.

Next, I obtain an expression for bank lending markups. The bank's profit function follows.

$$\Pi_B = \sum_j Ms_{jBt}\left(L\left(T_{jBt}, z_{jBt}, p_{jt}\right) - p_{jt} + D_{jBt}\right) - vc_B\left(Ms_{jBt}\right)$$

The bank only optimizes w.r.t. interest rates,

$$\frac{d\Pi_B}{dz_{jBt}} = s_{jBt}\frac{d\Delta_{jBt}^L}{dz_{jBt}} + \sum_{r\in J_m}\frac{ds_{rBt}}{dz_{jBt}}\Delta_{rBt}^L\left(Ms_{rBt}\right) = 0$$

Where,

$$\Delta_{rBt}^L\left(Ms_{rBt}\right) \equiv L\left(T_{rBt}, z_{rBt}, p_{rt}\right) - p_{rt} + D_{jBt} - mc_B\left(Ms_{rBt}\right)$$

Yielding the following expression for the vector of bank lending markups,

$$\Delta_B^L\left(Ms_B\right) = -\left(\frac{ds_B}{dz_B}\right)^{-1}\frac{dL_B}{dz_B}s_B$$

1.9.2 Solving for Markups (Captives Independent From Carmakers)

Carmaker f's manufacturing and sales profits, Π_{MS}^{f}

$$\begin{aligned}\Pi_{MS}^{f} &= \sum_{j\in J_f} M\left(s_{jB}+s_{jf}\right)\left(p_j-mc_j^f\right)\\ &\sum_{j\in J_f} M\left(s_{jB}+s_{jf}\right)\Delta_j^V\end{aligned}$$

Yielding price FOCs,

$$\frac{d\Pi_{MS}^{f}}{dp_j}=\sum_{r\in J_f}\left(\frac{ds_{rB}}{dp_j}+\frac{ds_{rf}}{dp_j}\right)\Delta_r^V+\left(s_{jB}+s_{jf}\right)=0$$

Or, stacking elements over vehicles and manufacturers, we obtain the following expression for vehicle markups,

$$\Delta_r^V=-\left(\frac{ds_B}{dp}+\frac{ds_f}{dp}\right)^{-1}\left(s_B+s_f\right)$$

The now independent captive finance company's profits, who still lends exclusively to firm f,

$$\Pi_C^f=\sum_{j\in J_f} s_{jf}\left(L\left(T_{jf},z_{jf},p_j\right)-p_j+D_{jf}\right)-vc_f\left(Ms_{jf}\right)$$

Then, captive's first order conditions with respect to interest rates,

$$\frac{d\Pi_C^f}{dz_{jf}}=s_{jf}\left(\frac{dL_{jf}}{dz_{jf}}\right)+\sum_{r\in J_f}\frac{ds_{rf}}{dz_{jf}}\Delta_{rf}^L\left(Ms_{rf}\right)$$

Or,

$$\Delta_f^L\left(Ms_f\right)=-\left(\frac{ds_f}{dz_f}\right)^{-1}\frac{dL_f}{dz_f}s_f$$

For traditional banks, their lending markup is unchanged, so,

$$\Delta_B^L\left(Ms_B\right)=-\left(\frac{ds_B}{dz_B}\right)^{-1}\frac{dL_B}{dz_B}s_B$$

1.9.3 Elasticity Formulas

Point elasticity formulas for the own-price elasticity of vehicle demand and own-interest rate elasticity of loan demand,

$$\begin{aligned}
\mathcal{E}_{D,p} &= \frac{ds_j}{dL_j}\frac{dL_j}{dp_j}\frac{p_j}{s_j} = \alpha\left(1-s_j\right)p_j\frac{zT}{1-\left(1+z\right)^{-T}} \\
\mathcal{E}_{D,z} &= \frac{ds_j}{dL_j}\frac{dL_j}{dz_j}\frac{z_j}{s_j} = \alpha\left(1-s_j\right)z_j\frac{T\left(p-d\right)}{1-\left(1+z\right)^{-T}}\left[1-\frac{zT\left(1+z\right)^{-(T+1)}}{1-\left(1+z\right)^{-T}}\right]
\end{aligned}$$

1.9.4 Robustness Checks

This section presents robustness specifications for the demand estimates. Specifically, to examine whether estimating a single coefficient on total loan is too restrictive, I let price, monthly interest rate (z_{jlmt}) and term (T_{jlmt}) enter separately.

$$\ln\left(\frac{s_{jlmt}}{s_{0mt}}\right) = X_{jt}\beta + \alpha p_{jmt} + \gamma z_{jlmt} + \eta T_{jlmt} + \xi_{jlmt}$$

I instrument price with the prices of the same good in nine different geographic markets and instrument interest rates with mortgage delinquency rates. The following table demonstrates that the IV parameter values for the physical characteristics of the vehicle are comparable with the preferred specification presented in the estimation section.

Table XI : OLS and IV demand robustness specification

	Coefficient	SE	Coefficient	SE
	OLS - Total Loan		IV - Delinquency	
Variable				
Price	-4.0E-05	4.4E-07	-3.7E-05	5.9E-07
Monthly APR	-74.82	1.46	-26.5	4.01
Term	-2.7E-03	4.7E-04	-8.5E-03	6.0E-04
Horsepower/Weight	3.3E-02	0.26	0.78	0.27
Miles/Gallon	1.4E-02	6.0E-04	1.8E-02	6.2E-04
Size (in^2)	1.1E-04	2.0E-06	1.2E-04	2.1E-06
Fixed Effects				
Market Fixed Effects	Y		Y	
Quarter Fixed Effects	Y		Y	
R^2	0.26		0.24	
N	213192		212226	

The table below demonstrates that the implied price elasticities are roughly comparable - over half of the implied elasticities are greater than one in absolute value.

Table XII : Price Elasticity of Vehicle Demand

	Mean	Std. Dev.	Min	Max	% Elastic
OLS	-1.13	0.38	-4.28	-0.41	43%
IV	-1.06	0.36	-4.01	-0.38	51%

Finally, the interest rate elasticities are more directly comparable to those from the preferred specification, where we obtain that all implied elasticities are less than one in absolute value.

Table XIII : Interest Rate Elasticity of Loan Demand

	Mean	Std. Dev.	Min	Max
OLS	-0.37	0.16	-1.39	-6.0E-04
IV	-0.13	0.06	-0.49	-2.2E-04

1.9.5 Counterfactual Output

Table XIV: Percent Change in APRs and Prices in Counterfactual Relative to Actual for Miami

Brand	Change in APR (%)	Change in Price (%)
Mitsubishi	90.6	-21.1
Suzuki	68.4	-12.2
Saturn	65.4	-8.1
Pontiac	58.5	-13.3
Subaru	56.6	-24.6
Mazda	54.3	-7.5
Mercury	42.5	-3.9
Volkswagen	40.9	-12.6
Kia	35.7	-10.7
Honda	17.8	-16.9
Buick	17.7	-6.0
Lincoln	14.1	-10.4
Hyundai	11.6	-3.0
Chrysler	11.5	-13.6
Toyota	10.7	-13.0
Nissan	10.3	-12.8
Jeep	7.2	-16.3
GMC	7.0	-9.5
Ford	5.9	-11.7
Chevrolet	1.2	-7.6
Scion	-0.9	-8.2
Dodge	-3.0	-17.8
Cadillac	-14.1	-4.3
RAM	-23.1	-7.5

Table XV: Percent Change in APRs and Prices in Counterfactual Relative to Actual for Minneapolis

Brand	Change in APR (%)	Change in Price (%)
Pontiac	48.0	-10.5
Suzuki	43.9	-13.7
Kia	36.1	-20.3
Scion	35.7	-24.7
Mercury	35.6	-6.4
Mitsubishi	35.5	-15.5
Mazda	33.5	-12.0
Volkswagen	30.1	-6.9
Saturn	27.9	-6.6
Honda	27.7	-8.2
Nissan	25.3	-11.9
Hyundai	17.1	-4.7
Subaru	13.9	-5.5
Toyota	13.1	-8.2
Chevrolet	12.1	-6.8
Buick	6.8	-5.4
Jeep	3.4	-17.4
Dodge	2.1	-14.8
Chrysler	-0.1	-13.4
Ford	-0.5	-10.4
GMC	-2.2	-8.8
Lincoln	-16.1	-9.6
RAM	-34.1	-7.3
Cadillac	-46.3	-1.5

1.10 References

1. Ackerberg, Daniel A. & Marc Rysman (2005). "Unobserved product differentiation in discrete-choice models : estimating price elasticities and welfare effects," *RAND Journal of Economics,* 36(4) : pp. 771-788.

2. Adams, William , Liran Einav & Jonathan Levin (2009). "Liquidity Constraints and Imperfect Information in Subprime Lending," *American Economic Review,* 99(1) : pp. 49-84.

3. Ashcraft, Adam , Allan Malz & Zoltan Pozsar (2012). "The Federal Reserve's Term Asset Backed Securities Loan Facility." *FRBNY Economic Policy Review,* November 2012 : pp.29-66.

4. Beresteanu, Arie & Shanjun Li (2011) "Gasoline Prices, Government Support and the Demand for Hybrid Vehicles in the US," *International Economic Review,* 52(1) : pp. 161-182.

5. Berry, Steven T., James Levinsohn & Ariel Pakes (1995). "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4) : pp. 841-890.

6. Berry, Steven T. (1994). "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics,* 25(2) : pp. 242-262

7. Boczar, Gregory E. (1978) "Competition Between Banks and Finance Companies : A Cross Section Study of Personal Loan Debtors," *The Journal of Finance,* 33(1) : pp. 245-258.

8. Bresnahan (1987) "Competition and Collusion in the American Automobile Industry : The 1955 Price War," *The Journal of Industrial Economics,* 35(4) : pp. 457-482.

9. Cetorelli, Niccola & Stavros Peristiani (2012). "The Role of Banks in Asset Securitization." *FRBNY Economics Policy Review,* July 2012 : pp. 47-63.

10. Crain, Keith. "Cars run on money and credit, not gas." *Automotive News,* September 22, 2008. Retrieved November 8, 2012.

    <http://www.autonews.com/apps/pbcs.dll/article?AID=/20080922/OEM/309229848>

11. Dieckman, Christian (2008). "Automotive Finance : The Case for an Industry-Specific Approach to Risk Management," in Georg Bo, Svetlozar Rachev & Reinhold Würth (eds) *Risk Assessment* : pp. 1-9.

12. FICO, *What's in my FICO Score.* Retrieved November 21, 2012.

<http://www.myfico.com/CreditEducation/WhatsInYourScore.aspx>

13. Goldberg, Pinelopi Koujianou (1995) "Product Differentiation and Oligopoly in International Markets : The Case of the US Automobile Industry," *Econometrica*, 63(4) : pp. 891-951.

14. Gowrisinkaran, Gautam & Marc Rysman (2012) "Dynamics of Consumer Demand for New Durable Goods," working paper.

15. Hausman, J. (1996). "Valuation of New Goods under Perfect and Imperfect Competition," in T. Bresnahan & R. Gordon, eds., *The Economics of New Goods,* Studies in Income and Wealth, Vol. 58, Chicago: NBER.

16. Hausman, J., Gregory Leonard, & J. Douglas Zona (1994). "Competitive Analysis with Differentiated Products," *Annales d'Economie et de Statistique,* 34, pp: 159-180.

17. Henry, Jim. "No captive? No problem - so far." *AutoNews Online,* October 12, 2011. Retrieved April 28, 2013.

<http://www.autonews.com/article/20111012/BLOG13/310139999/1142#axzz2RgopqoJx>

18. Henry, Jim. "Flying without a captive - and with one." *AutoNews Online,* April 3rd, 2013. Retrieved, April 28, 2013.

<http://www.autonews.com/article/20130403/FINANCE_AND_INSURANCE/130409952/...

flying-without-a-captive-and-with-one#axzz2RgopqoJx>

19. Hirsh, Evan, Jan Miecznikowski, & Robert Reppa (2004). *Rebundling of the Auto Finance Industry,* Booz Allen & Company report. Retrieved November 8, 2012.

20. IXIS, Corporate & Investment Bank. "Spreads and Credit : Credit Research - Automobile/Banks." December 5, 2005. Retrieved November 8, 2012.

<http://www.docstoc.com/docs/80767710/Automobile-Banks : investment banking document regarding captive finance>

21. Krebs, Michelle. "2007 : A Historical Year for US Vehicle Sales." *Edmunds Auto Observer,* January 4, 2008. Retrieved November 21, 2012.

<http://www.edmunds.com/autoobserver-archive/2008/01/...

2007-a-historic-year-for-us-vehicle-sales.html>

22. Lunby, Tami. "Consumers cut up credit cards, but buy cars." *CNN Money,* August 29, 2012. Retrieved November 21, 2012.

<http://money.cnn.com/2012/08/29/news/economy/credit-card-debt/index.html>

23. Nevo, Aviv (2000 a). "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand," *Journal of Economics & Management Strategy,* 9(4) : pp. 513-548.

24. Nevo, Aviv (2000 b) "Mergers with Differentiated Products : the Case of the Ready-to-Eat Cereal Industry," *RAND Journal of Economics,* 31(3) : pp. 395-421.

25. Nussel, Philip & Donna Harris. "GMAC auto finance business loses $1.31 billion." *Automotive News,* February 3, 2009. Retrieved November 21, 2012.

<http://www.autonews.com/apps/pbcs.dll/article?AID=/20090203/RETAIL02/302039741>

26. Rattner, Steven (2010). *Overhaul : An Insider's Account of the Obama Administration's Emergency Rescue of the Auto Industry.* New York, NY : Houghton Mifflin Harcourt Publishing.

27. Sousanis, John. "Sandy Should Push November SAAR Past 15 million." *Wards Auto*, November 21, 2012. Retrieved November 21, 2012.

<http://wardsauto.com/sales-amp-marketing/sandy-should-push-november-saar-past-15-million>

28. Tedlow, Richard S. (1988). "The Struggle for Dominance in the Automobile Market : The Early Years of Ford and General Motors," in W. J. Hausman (ed.). *Business and Economic History: Papers Presented at the Thirty-Fifth Annual Meeting of the Business History Conference, March 31-April 2, 1989 Boston, Massachusetts,* Williamsburg, VA : College of William and Mary, Department of Economics, Business History Conference : pp.49-62.

29. US Commodities Future Trading Commission (2011). *Re: Margin and Capital Requirements for Covered Swap Entities/File Numbers RIN: 1557-AD43;RIN: 7100 AD74; RIN : 3064-AD79; RIN 2590-AA45; RIN : 3052-AC69.*

30. US Government Accountability Office (2009). *Auto Industry : Summary of Government Efforts and Automakers' Restructuring to Date.*

31. Vlasic Bill (2011). *Once Upon A Car : The Fall and Resurrection of America's Big Three Automakers : GM, Ford & Chrysler.* New York, NY : HarperCollins Publishers.

32. Warren, Elizabeth, Paul S. Atkins, Richard H. Neiman, Damon Silvers, & J. Mark McWatters (2010). *The Unique Treatment of GMAC Under the TARP.* Congressional Oversight Panel, March Oversight Report.

33. Wilson, James A., Sandra L. DiChiaria, Edward B. Butler, Kevin J. Glueckert "Supervisory Insights : The Changing Landscape of Indirect Automobile Lending." FDIC 2005, retrieved November 21, 2012. <http://www.fdic.gov/regulations/examinations/supervisory/insights/...

sisum05/article04_auto_lending.html>

Chapter 2

Gas Price Expectation & Vehicle Replacement:

Evidence & Analysis Using Survey Data

2.1 Introduction

Since at least the 1970s oil crises, economists have sought to estimate how consumer demand for gasoline responds to fluctuations in oil prices.[1] Even the earliest analyses in this literature acknowledged that in the long run, consumers can respond to gas price fluctuation by changing the cars that they own. More recent analysis has sought to explicitly model the dynamic adjustment process whereby consumers change their vehicles in response to exogenous factors such as gasoline prices.[2] Given that adjustment is dynamic, consumers respond not only to current gasoline prices, but their forecasts for future gasoline prices. Yet, it has proven extremely difficult to obtain data on consumer gasoline price forecasts, let alone expectations data combined with consumer vehicle choice data.[3] To overcome this missing data problem, researchers have used realized gasoline price data and paired it with the assumption that consumer gasoline price forecasts follow the rational expectations model. However, the analysis of Anderson, Kellogg, Sallee & Curtin (2011)

[1] Bassom & Oum (2007).

[2] Shiraldi (2011), Gillingham (2013), Liu (2013).

[3] A notable exception is Allcott (2011) who has such data for a single sample period, after the 2008-2009 retail gas price spike.

shows that an surveyed expectations can differ considerably from rational expectations[4] In light of this, it is evident that omitting surveyed expectations data from dynamic car choice analysis can lead to biased estimates of consumer responsiveness to gasoline price fluctuation.

In this paper, I introduce entirely new survey data which combines consumers' forecasts for future gasoline prices with their intended replacement vehicles. These data allow me to overcome the missing data hurdle encountered previously in this line of research. I use these data to disentangle the relative contributions to vehicle replacement decisions of persistent vehicle preferences and beliefs over future gasoline prices. This analysis is partly motivated by the ongoing discussion of the 'energy paradox.'[5] The basic premise of this paradox is that consumers purportedly fail to invest in cost-saving, energy efficient technologies. The possibility that gasoline price expectations may play a role in determining under-investment has been discussed in this literature.[6] Yet, to date no empirical analysis has used surveyed expectations to address the question.[7]

Fundamentally, the energy paradox literature is a discussion about technology adoption. In the case of energy-efficient vehicles, both firms and the government alike have a strong interest in accurately calculating the determinants of consumer demand for different technologies. Yet, without detailed micro-data, accurate assessments of future demand are difficult to obtain. One needs to look no further than to the recent introduction of electric vehicles to find practical examples. Accurate assessments of consumers' future demand for energy efficient vehicles are a vital input into the investment decision-making of the firms developing these technologies and local, state and federal governments offering incentives for producers and consumers of these technologies. In a prominent example, General Motors over-anticipated demand for the Chevy Volt that cost nearly $1.2 billion to develop and market. Demand was so low that in uncharacteristic honesty, in 2012 after two years on the market one GM executive was quoted as saying, "It's true, we're not making money yet."[8] Similarly, the Federal government has been left on the hook in such instances as

[4] Anderson, Kellogg, Sallee, Curtin (2011) show that surveyed gas price expectations are on average consistent with a "no change" forecast, which they posit as the rational expectations process for this setting. However, for the period from 2008-2009, their data reject the "no change" forecast model, or surveyed expectations do not necessarily follow a rational expectations model.

[5] Jaffe & Stavins (1994) and Allcott & Wozny (2012).

[6] Jaffe & Stavins (1994), Allcott & Wozny (2012), and Anderson, Kellogg, Sallee, & Curtin (2011).

[7] Again, the notable exception being Allcott (2011).

[8] Reuter News Service, 2012.

well. In the case of Michigan's LG Chem, after receiving over $325 million in local, state, and federal grants and tax relief, the company has yet to produce a single battery for use in American autos - in particular Chevy Volts which the factory was intended to supply. In these examples, firms and governments alike may have over-anticipated consumers' responsiveness to gasoline price fluctuation and therefore adoption rates of alternative fuel vehicles.

This paper looks at a very specific set of questions at the crux of this issue. First do gas price expectations impact consumers decisions over which cars to buy and hold? Second, how do persistent preferences for vehicles - or state dependence - figure into the decision-making of consumers to hold or replace vehicles? The primary contribution of this paper is to introduce and analyze entirely novel data to answer these questions. The data come from a national survey of vehicle owners in the US, and span a highly relevant period of gasoline price fluctuation. The second major contribution of this paper is to use the data to estimate a model of vehicle replacement which takes as given consumers' current vehicles and their forecasts over future gasoline prices. The estimated model addresses both of the papers' questions and is a laboratory for examining the relative contributions of state dependence and expectations in determining vehicle replacement.

The data are perfect for the question at hand, both in terms of providing identification but also in terms of providing direct insight into expectations and preferences. To identify how current and expected gasoline prices affect vehicle replacement decisions, one would necessarily need to observe variation in both current and expected gasoline prices. My data certainly exhibit this property. They span the emergence of a new gasoline price level - post - Hurricane Katrina in 2005, the rapid run-up in prices in 2008, their subsequent decline and return to the new price level through 2010. Further, the data allow me to disentangle true underlying preferences from transactional details (e.g. financial incentives provided by manufacturers, trade-in value of current vehicle, availability of credit). To get a clean measure of preference at a given point in time, the survey asks consumers to choose their replacement as if their current vehicle were lost or stolen.

Beyond providing a basis for identification, the data are sufficiently informative to provide direct insight into my two primary research questions. As a preliminary, the data demonstrate that expectations deviate from the rational expectations model in periods where gasoline prices rise dramatically. Next, I can directly observe an interesting correlation between preferences and expectations. Survey respondents who drive or who intend to replace to small cars have on average the most aggressive expectations for future gasoline prices. In contrast, respondents who currently own and who intend to hold on to large vehicles tend to expect gasoline prices to change relatively little, on average. The data also demonstrate that while expectations do matter, they must contend with strong state dependence : at least 70% of respondents would prefer to hold on to their current vehicle if offered the choice to switch.

To more formally address these factors, I specify and estimate a discrete choice model wherein consumers choose which vehicle they would like to choose as a replacement conditional on their current vehicle and given their beliefs over how gasoline prices will evolve over the coming year. The model captures relevant features of consumers' dynamic decision problems without solving and estimating a structural dynamic programming problem. Often in this particular setting of vehicle replacement, authors choose to specify and estimate structural models precisely because they do not observe micro data linking consumers' gas price forecasts, current vehicle holdings, and replacement vehicles. In particular, a structural model allows one to (1) link aggregate data on current vehicles and replacement vehicles and (2) assume a process for expectations so that average retail gasoline prices can be employed. My data link consumers current vehicle with their intended replacement, directly addressing concerns regarding the accuracy of linking aggregate data on trade-ins and new car purchases. The data are also linked to consumers' forecasts for future gasoline prices. As I note, the expectations in my survey data demonstrate that gasoline price expectations are not consistent with the rational expectations hypothesis would predict. In this light, one benefit of using data on forecasts directly - rather than imputed values implied by a model - is that I avoid making a priori assumptions on the expectations process which appear to be invalidated in my data.

The estimated model confirms the intuition from the raw data. Both expected prices and current vehicle segment are statistically significant determinants of replacement vehicles. The effects of expected prices and current vehicle segment differ across choice options - expected prices have a different effect on the choice of SUVs than they do for midsize cars. While both factors determine replacement decisions, persistent preferences have a larger overall effect. Without controlling for current vehicle, an increase in average gas price expectations by about $0.75 (a one standard deviation increase) implies that a consumer would be 40% more likely to choose a compact car relative to an SUV. However, taking into account current vehicle, increasing average gas price expectations by the same amount only leads to an increase in the share of SUV owners switching to small cars from 9% to 12%. The results are robust to a number of alternative specifications, including the addition of, (1) vehicles to the choice set, (2) vehicle characteristics to the model, (3) demographic characteristics to the model.

The parameter estimates - and experiments based on the parameter estimates - suggest that low adoption of fuel efficient technologies is less driven by consumers having inaccurate gas price expectations and more by persistent vehicle preferences. This result is somewhat at odds with what is hypothesized in the Energy Paradox literature, yet consistent with standard results from consumer theory.

The results also have implications for policy-makers. The results speak to the relative merit of alternative new vehicle incentives offered by local, state, and federal governments. The conclusion of the paper is that fuel prices - both current and expected - do matter in determining consumers vehicle replacement decisions, yet state dependence is the dominant determinant. That is, by and large, small car drivers are small car drivers. The policy implication is that government incentives attached to fuel efficient vehicles may simply be a transfer to consumers already inclined to small drive and fuel-efficient vehicles. If the policy intent is to incentivize consumers away from large vehicles, then incentives should be linked to both the trade-in vehicle and the new vehicle. As an example, incentives could be increasing in the difference between miles-per-gallon on the trade-in vehicle and miles-per-gallon on the new vehicle.

The paper proceeds as follows. First, the data are introduced in general terms. Then, details regarding the expectations component of the data and the vehicle choice component of the data are provided in separate sections. Subsequently, I specify the baseline econometric model as well as alternative specifications. After discussing identification of model parameters, I present results for both the baseline and alternative specifications. The final section concludes.

2.2 Data

2.2.1 Fuel Prices & Intentions Survey

The data are from the Fuel Prices & Intentions Survey (FPIS), conducted by an automotive market research firm. The FPIS data are comprised of thirty samples of at least 750 respondents, spanning the period from September 2005 to April 2010. Note that this time period includes **both** the post-Katrina burst in gas prices (summer/fall 2005) and the major run-up in gas prices over 2008. This is the only known source of data to span this time period *and* include both measures of respondents' current and intended replacement vehicle, as well as their gas price expectations.

The survey firm conducts multiple surveys concurrently and targets potential respondents through multiple channels (e.g. mass-mailings, advertising in automotive magazines). Once a respondent opts into a given survey pool, actual surveying is conducted through a self-directed survey via an online portal.

In addition to gas price and vehicle variables, the data include some demographic variables for the second half of the survey sample. In particular, I have data on respondent age, education, income (range-valued) male (binary), married (binary), and presence of children in the household. The tables below provide summary statistics for these demographic data.

Table 1 : Demographic Summary Statistics

	Mean	50th Percentile
Income	$ 110,975.00	$ 95,000.00
Age	55	56
Education	College Grad	NA

Table 2 : Demographic Summary Statistics

	Percentage
Children in Household	0.21
Married	0.78
Male	0.79

Some conclusions emerge immediately from these tables. Respondents tend to be older, wealthier and more educated than the average US citizen. Furthermore, the sample is skewed towards males and marriage rates are high. The presence of children in household being low is not surprising given average age of respondent. Beyond the fact that the sample is not statistically representative of the US population, there is a slight bias towards automotive enthusiasts in the sample. This is a result of how the survey firm obtains survey volunteers - via automotive magazines.

The sample selection bias may impact results in the following ways. First, given the average level of education and income of respondents (thus increase likelihood that they conduct personal investment), one may expect that survey respondents will give more accurate assessments of future gasoline prices. Because of the bias towards automotive enthusiasts, one could also reasonably expect that respondents understand differences in vehicle characteristics (including fuel efficiency) better than the average US consumer. Third, given the average income of survey respondents, all else equal I would not expect the average respondent to be experiencing significant current hardship as a result of gasoline prices. That said, were gas prices to be placing strain on survey respondents, given average income and other demographic averages, the respondents from this survey would most likely still have had access to credit should it have been necessary to obtain a new vehicle. In that sense, their replacement responses are probably more credible than responses from a more representative survey population.

2.2.2 Gasoline Price Expectations

The survey collects from respondents data on current gas price as well as their expected gas price level a year from now. Prices are measured in \$0.25 or \$1 intervals : \$1 intervals are reserved for the right tails of the expected gas price distribution. In all specifications, the constructed variable $\Delta p^e_{g,i,t} = p^e_{g,i,t} - p^c_{g,i,t}$ is used, rather than either current gas prices, expected gas prices, or both measures separately. The reason for this is two-fold. First, on a conceptual level, when agents make their forecasts, it is reasonable to assume that they condition on current gasoline prices. Mechanically, as current prices and expected prices are correlated in the data, including both in our regressions will lead to imprecise parameter estimates. Second, agents may experience high current gas price levels, and thus form higher forecasts, because they live in a region predisposed to higher current gas prices. Differencing mitigates against the effects of confounding variation due to geography.

Table 3 : Expected Gas Price Summary Statistics Change by Sample

Sub-Sample	Mean	Std. Dev.	t-test
2005	0.047	0.566	3.49
2006	0.277	0.553	42.81
2007	0.263	0.463	26.91
2008	0.695	0.924	52.42
2009	0.607	0.806	53.11
2010	0.402	0.671	28.47
Overall	0.427	0.735	90.05

The table above provides summary statistics of the constructed variable $\Delta p^e_{g,i,t}$. Summaries are provided for each year in the sample as well as over all years. On both a per-sample and overall basis, expectations are consistently and statistically significantly biased above current retail prices. The extent of this bias is most pronounced for samples in the years 2008 and 2009. The fact that this is the case for 2008 is not surprising - that year saw the the largest and most rapid run-up in retail gasoline prices in recent history. That we see such bias in 2009 is somewhat surprising, as this year saw on of the most rapid declines in retail gasoline prices.

In the figure below, I plot the time series of the national average retail gasoline price from June 2005-July 2010, using standard data reported by the Energy Information Agency. Call this series $\bar{p}^c_{g,t}$. Against this series, I also plot what the average of my survey population stated would be one year from now (black squares). That is, I take the average of the $\Delta p^e_{g,i,t}$ for a given survey sample $\left(\Delta \bar{p}^e_{g,t}\right)$ and I add this to the current actual price, and obtain $\bar{p}^c_{g,t} + \Delta \bar{p}^e_{g,t}$

Figure 1 : Expected Retail Gasoline Price Level vs. Actual, 2005-2010

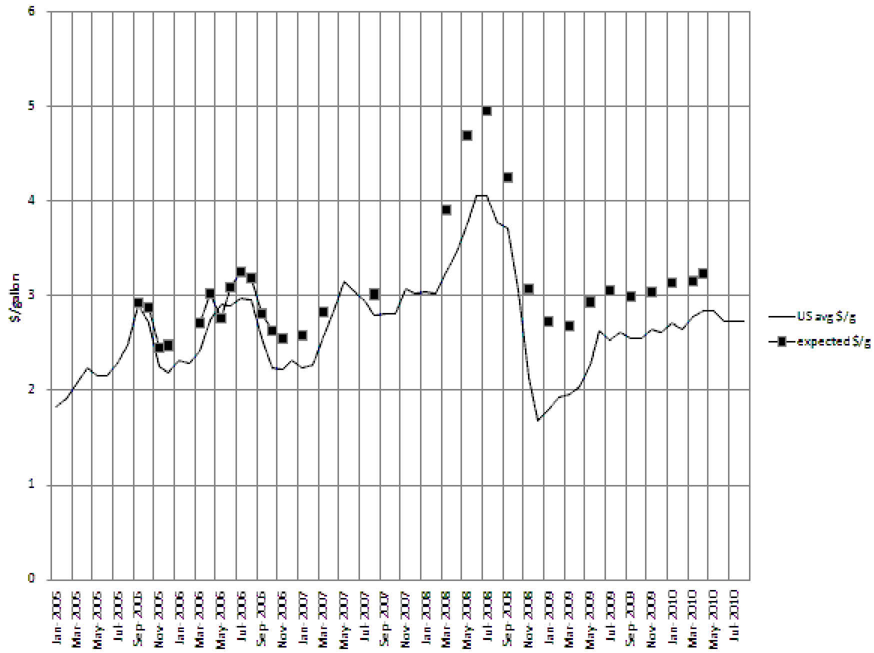

This figure serves to highlight three features of expectations during this time period. First, forecasts were consistently biased above current actual values. Second, this bias ramped up considerably in the run-up of gasoline prices in 2008. Third and finally, the dampening of bias was slow even as prices declined in 2009. That is, when underlying prices are volatile, consumers forecast that prices will be biased above current levels and forecasts respond asymmetrically to rapid increases and decreases in underlying prices.

These findings corroborate and expand upon what Anderson, Kellogg, Sallee & Curtin (2011) find with regard to gas prices over this condensed time horizon. In particular, I find that consumers' expectations over future gas prices differ from the no-change process which Anderson, Kellogg, Sallee & Curtin contend to be the rational expectations process when gas prices are relatively calm. While the data are incomplete from a time series perspective, it seems reasonable to conclude that the data would reject standard models of expectation formation. It stands to reason that to accurately estimate consumer's vehicle choice problem, one should use their surveyed expectations rather than modeled expectations.

2.2.3 Vehicle Choice and Gas Prices in the Raw Data

In the survey, both current and replacement vehicles are defined by their segmentation. Each vehicle segment is associated with three example cars which are meant to be representative of the vehicle segment in question (see table below). Notably for this study, the intent to replace question is stated in such a manner so as to minimize the role of such confounding factors as current vehicle resale price and optimal replacement time as a function of state variables. Agents are asked to imagine that their car has been either lost or stolen, and that they **must** choose a replacement **today.**

Figure 1 : FPIS Vehicle Replacement Survey Question

4a: If your primary vehicle had to be replaced right now (for example, if it was stolen), what type of vehicle would you choose to replace it?

Please tick **only one** of the following:

- Luxury or Large Car (e.g. Mercedes S-Class, Lexus ES330, Mercury Grand Marquis)
- Mid-Size Car (e.g. Toyota Camry, Honda Accord, Chevrolet Impala)
- Small Car/Compact Car (e.g. Ford Focus, Honda Civic, Hyundai Elantra)
- Sports/Sporty Car (e.g. Porsche Boxster, Ford Mustang, Mazda MX-5 Miata)
- Pickup Truck (e.g. Chevrolet Silverado, Dodge Ram, Nissan Frontier)
- Sport Utility Vehicle (SUV) (e.g. Ford Explorer, GMC Envoy, Toyota RAV4)
- Minivan (e.g. Dodge Caravan, Toyota Sienna, Honda Odyssey)

In the primary specification, respondents who currently own or intend to replace to a large/luxury car, minivan, or sports car are dropped. To precisely estimate model parameters, I need to observe sufficiently many responses for each combination of current segment and replacement segment. It happens to be the case that in my data, combinations involving large/luxury cars, minivans, and sports cars as either the current or replacement tend to yield insufficient observations. As a proportion of actual vehicle ownership and vehicle transactions, these segments tend to be relatively small. Their exclusion from the following analysis is therefore not expected to affect this paper's broader implications.[9]

The stated vehicle preference regression results are foreshadowed in the raw data. Specifically,

- Higher expected gas prices among those *replacing* to more fuel efficient segments.
- Higher expected gas prices among those *currently* in more fuel efficient segments.
- Lower expected gas prices among those who currently own and would replace to trucks or SUVs.
- High degree of state dependence in vehicle segment preferences.

[9]Robustness of results is investigated by including the large/luxury category in an alternative specifciation. Results from the baseline specification are unaffected.

For evidence of the first two points, note the table below. To construct the table, current vehicle, replacement vehicle, and Δp_g^e were pooled over individuals and over time. The table entries are the average expected change in gas prices, for a given current vehicle/replacement vehicle combination. We see from the "Small" column, that those replacing to the most fuel efficient segment tend to have higher expected changes in gas price, regardless of current segment. In the "Average" column, we see evidence of the second point - current owners of smaller vehicles tend to expect larger price changes in the future. Looking at the cells associated with truck and SUV owners who hold on to their vehicles, we see that these types of consumers report the lowest expected change in gasoline prices.

Table 4 : Expected Gas Price Change by Current vs. Replacement Vehicle Segment

Current Vehicle	Replacement Vehicle				
	Medium	Small	Truck	SUV	Average
Medium	\$ 0.40	\$ 0.60	\$ 0.38	\$ 0.38	\$ 0.44
Small	\$ 0.32	\$ 0.56	\$ 0.44	\$ 0.43	\$ 0.52
Truck	\$ 0.41	\$ 0.59	\$ 0.35	\$ 0.35	\$ 0.39
SUV	\$ 0.46	\$ 0.60	\$ 0.32	\$ 0.32	\$ 0.39

The fourth point - that there is considerable state dependence in the data - can be seen in the table below. Data were pooled over time and individuals. Each entry represents the share of those currently in segment A who choose B as their replacement. For instance, 20% of those who currently own a medium car indicate that they would go down to a small car if they have to replace today. Diagonal entries in the table demonstrate the amount of state-dependence for a given current segment. At least 70% stick with their current vehicle segment, regardless of which current segment we are considering. Comparing those in currently larger segments (truck or suv) to those currently in smaller segments (medium or small), we see that large-to-small transitions are slightly more likely than small-to-large.

Table 5 : Current vs. Replacement Vehicle Segment Transitions

Current Vehicle	Replacement Vehicle			
	Medium	Small	Truck	SUV
Medium	72%	20%	2%	6%
Small	12%	82%	2%	4%
Truck	7%	10%	75%	8%
SUV	12%	9%	4%	75%

2.3 Estimation & Results

2.3.1 Specification

My objective is to assess empirically the effects of gasoline price expectations and persistent vehicle preferences on the choice of replacement vehicle. I follow a standard econometric approach to achieve this objective. It is natural to think of agents as making discrete choices over the vehicles they intend to own. Gas price expectations and the vehicle an agent currently owns are certainly factors which impact their choice of replacement vehicle. Additionally, it is reasonable to hypothesize that these factors have a differential effect across choices - e.g. gas price expectations may have a different effect on the probability of choosing an SUV versus a small car.

To accomplish my empirical objective, and to flexibly incorporate the differential impact on choice vehicle of factors such as gasoline price expectations and current vehicle ownership, I employ the multinomial logit model. The preferred econometric specification addresses whether observed variation in stated replacement vehicles can be explained by variation in expected gas prices, conditional on agents' current vehicles. The replacement outcome is discrete and individual covariates do not vary by segment. The multinomial logit model with choice-invariant regressors is flexible and perfect for the choice problem at hand. It allows me to estimate parameters by choice rather than across choices. This implies that there will be separate coefficients for gasoline price expectations and current vehicle (Midsize, Small, SUV, or Truck) across replacement choices.

Formally, a survey respondent is endowed with a vehicle k from one of the four possible vehicle segments(Midsize, Small, SUV, or Truck)[10]. Given this endowment and his expected change in gas prices, $\Delta p^e_{g,i}$, the respondent chooses from one of the four possible replacement segments. Or,

[10]In alternative specifications, I expand the set of current and replacement vehicles to include large/luxury cars. This does not materially affect the results from the preferred specification.

$$
\begin{aligned}
u_{ikj} &= \beta_j + \gamma_j \Delta p_i^e + \eta_{kj} + \epsilon_{ij} \\
u_{ij} &= \beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \setminus \{1\}} \eta_{kj} 1\left(c_i = k\right) + \epsilon_{ij}
\end{aligned}
$$

The first line represents the utility to agent i, endowed with vehicle segment k, from choosing replacement segment j. The second line notes that the utility from replacement j depends on the vehicle with which agent i is endowed. The γ_j parameters capture differential impact of gas price expectations on the probability of choosing a given alternative. The η_{kj} parameters capture the effect of owning a current vehicle type on the probability of choosing a given replacement. This is intended to capture the effects of state dependent or persistent preferences - a feature which is allowed to differ across choice segments. Then, assuming that the ϵ_{ij} are distributed type 1 extreme value, the probability that i chooses j is,

$$
P_{ij} = \frac{\exp\left(\beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \setminus \{1\}} \eta_{kj} 1\left(c_i = k\right)\right)}{1 + \sum_{j \in J \setminus \{1\}} \exp\left(\beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \setminus \{1\}} \eta_{kj} 1\left(c_i = k\right)\right)}
$$

Agents in the model are required to currently own a vehicle from one of the four choice segments and also to choose a replacement from one of the four choice segments. I.e., there is no outside option for replacement and since respondents all start with an endowment vehicle, there is no new entry into the vehicle market. These assumptions on the model are entirely consistent with the survey design (respondents are required to choose from one of the possible vehicle segments) and the data (all respondents currently own a vehicle).

2.3.2 Alternative Specifications

To assess the robustness of the results from my primary specification, and in an attempt to incorporate more identifying variation, I specify and estimate three alternative specifications. In the first, I expand the choice set for agents to include large/luxury cars. This does not affect the baseline model in any major way. The two remaining alternative specifications do involve a material deviation from the baseline specification.

In another set of alternative specifications, I incorporate additional product characteristics. The survey question regarding replacement vehicles lists three example vehicles for each of the possible replacement segments. I incorporate product characteristics considered standard in automotive demand analysis (MSRP, size - length x width, horsepower over weight, and miles per gallon).[11] Specifically, as I only see the segment for both the current and replacement vehicle, I use the average of characteristics across example vehicles provided in the survey question, using data obtained from Ward's Automotive.

Since vehicle characteristics vary by choice, for this set of specifications, I am formally estimating a conditional logit model. The expression for the utility to agent i endowed with vehicle from segment k from choosing a vehicle from segment j is given by,

$$\begin{aligned} u_{ikj} &= \beta_j + \gamma_j \Delta p_i^e + \eta_{kj} + X_j \alpha + \epsilon_{ij} \\ u_{ij} &= \beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \setminus \{1\}} \eta_{kj} 1\left(c_i = k\right) + X_j \alpha + \epsilon_{ij} \end{aligned}$$

Note, X_j is the vector of segment average vehicle characteristics and α is the vector of coefficients reflecting the marginal utility contribution of each of the product characteristics. Again, under the assumption that the random taste shocks, ϵ_{ij}, are distributed Type 1 Extreme Value, I obtain the following expression for the probability of choosing a given vehicle segment.

$$P_{ij} = \frac{\exp\left(\beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \setminus \{1\}} \eta_{kj} 1\left(c_i = k\right) + X_j \alpha\right)}{1 + \sum_{j \in J \setminus \{1\}} \exp\left(\beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \setminus \{1\}} \eta_{kj} 1\left(c_i = k\right) + X_j \alpha\right)}$$

In addition, I estimate a specification of the model which includes the characteristics of the endowment vehicle. This entails the addition of four right-hand side variables and as many additional parameters to be estimated, *per replacement option.* I choose to incorporate current vehicle characteristics in this manner, rather than estimate separate effects for each of the possible endowment vehicles, across all replacement options. Formally, utility to i from choosing segment j given endowment k, is given by,

[11] See for instance Berry, Levinsohn, & Pakes (1995) and Beresteanu & Li (2011).

$$
\begin{aligned}
u_{ikj} &= \beta_j + \gamma_j \Delta p_i^e + \eta_{kj} + X_j \alpha + \sum_{l \in L} \phi_l X_{kl} + \epsilon_{ij} \\
u_{ij} &= \beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \backslash \{1\}} \eta_{kj} 1 \left(c_i = k\right) + X_j \alpha + \sum_{l \in L} \phi_{jl} X_{kl} + \epsilon_{ij}
\end{aligned}
$$

Where the ϕ_{jl} capture the marginal utility of owned vehicle characteristics, and are allowed to differ across replacement vehicle options. Again, given the assumption that the random taste shocks are distributed Type 1 Extreme Value, the probability of choosing a given replacement segment is given by,

$$
P_{ij} = \frac{\exp\left(\beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \backslash \{1\}} \eta_{kj} 1 \left(c_i = k\right) + X_j \alpha + \sum_{l \in L} \phi_l X_{kl}\right)}{1 + \sum_{j \in J \backslash \{1\}} \exp\left(\beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \backslash \{1\}} \eta_{kj} 1 \left(c_i = k\right) + X_j \alpha + \sum_{l \in L} \phi_l X_{kl}\right)}
$$

The final alternative specification incorporates survey data on demographics of the respondents. The hypothesis for this specification is that including demographic data will control for unobserved heterogeneity which affects replacement decisions, even after conditioning on current vehicle and expected gasoline prices. I include demographics in level.[12] The included demographic data are age of respondent, education level, gender, income, marital status and the presence of children in the household. This entails the following modification to the baseline model,

$$
\begin{aligned}
u_{ikj} &= \beta_j + \gamma_j \Delta p_i^e + \eta_{kj} z_i + \sum_{d} \tau_j z_{id} + \epsilon_{ij} \\
u_{ij} &= \beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \backslash \{1\}} \eta_{kj} 1 \left(c_i = k\right) + \sum_{d} \tau_j z_{id} + \epsilon_{ij}
\end{aligned}
$$

Which again yields the choice probabilities,

$$
P_{ij} = \frac{\exp\left(\beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \backslash \{1\}} \eta_{kj} 1 \left(c_i = k\right) + \sum_{d} \tau_j z_{id} + \epsilon_{ij}\right)}{1 + \sum_{j \in J \backslash \{1\}} \exp\left(\beta_j + \gamma_j \Delta p_i^e + \sum_{k \in J \backslash \{1\}} \eta_{kj} 1 \left(c_i = k\right) + \sum_{d} \tau_j z_{id} + \epsilon_{ij}\right)}
$$

[12] In unreported results, I estimated models with interaction terms between vehicle characteristics, expected gas prices and current segment dummies. These models generated overall model fit that was no better than the currently reported results. Individual parameter estimates tended to be insignificant. This is attributable to two facts : (1) there are less data in the sample with demographics, (2) interacting demographics with other variables increases the number of parameter estimates dramatically.

2.3.3 Identification

The main identification goal of this paper is to identify the effect of gas price expectations and the effect of persistent preferences on vehicle replacement probabilities. Furthermore, one hypothesis of the paper is that these factors differ in how they relate to given replacement choices. Given the nature of the survey - that a large number of respondents are polled each sample period and the sampling occurs over time - both cross-sectional variation and time series variation are employed to identify the models parameters.

Both cross-sectional and time series variation identify the effects of gas price expectations on replacement choices, conditional on current vehicle. The cross section yields me variation for a given time period and the time series yields me changes in the mean forecast over time. The data exhibit variation along both of these dimensions for each combination of current and replacement vehicle. Identifying persistent vehicle preferences follows a similar logic. Pooled cross-sectional and time series variation in conditional replacement probabilities identifies the effects of persistent preferences on replacement. In essence, the identifications comes from changes in the proportion of consumers who persist in or transition out of a given vehicle segment.

The raw data raise the issue of whether preferences and expectations are correlated. Of particular concern is whether preferences are evolving over time - i.e. whether the value a consumer places on holding onto a particular vehicle is changing over time. Were this to be the case, separately identifying expectations from preferences would be exceedingly challenging. In order to identify these two effects - gas price expectations versus persistent preferences - an additional assumption is necessary. It is conventional to assume that preferences are stable over time and I adopt this assumption here. Given this assumption on preferences, it is possible to claim that I am using the time series aspect of the data to identify expectations from preferences.

An additional assumption is necessary to pin down the mean of coefficients in the multinomial logit model. Specifically, it is necessary to assume the existence of a base option. Given an assumption on a base alternative, all model coefficients can be interpreted relative to the base option. For what follows, I assume a base with a natural interpretation for the current question. I assume that small cars are the base option. I do however re-run the model to assess the robustness of this assumption. The choice of base option does not affect overall fit of the model to the data, but can affect the parameter estimate values and their significance.

2.3.4 Results

Preferred Specification

This section provides results from both the preferred specification and from experiments based on model estimates. Model parameters are estimated via maximum likelihood. The multinomial logit model variants require specification of a base option to identify parameter estimates. In all instances, I use the small vehicle segment as the base option. For the preferred specification, I re-run the model using each possible choice segment as a base option. This allows me to examine the impact of the choice of base vehicle on model estimates. Parameter estimates from the baseline four-segment choice problem are included in the appendix in Table 10. I will highlight features of the estimates here and focus on the parameters associated with gas price expectations.

The overall fit of the baseline model is good - the model explains 43% of the variation present in the data. Most parameter estimates are significant at conventional levels and the direction of effects is as hypothesized. From the estimated model I conclude that there is strong evidence for persistent preferences. The estimated coefficient on the vehicle segment where current vehicle is the same as the replacement is always larger than the estimated coefficients for segments where current vehicle does not match the replacement option. Likelihood ratio tests reject the null that current vehicle effects are constant across choice options.[13]

The signs on estimated gas price expectation coefficients are as expected as well. In the table below I present estimates for the effects of gasoline price expectations on replacement choice probabilities. I include estimates for each possible choice of base option in the multinomial logit specification. A row corresponds to a replacement choice, a column to the base option. One can see that increasing Δp_g^e decreases the probability of choosing larger segments and increases the probability of choosing smaller segments. This effect is especially pronounced for the small segment. For instance, looking at the second column of the table below (where small is the base option), increases in expected gas prices decreases the probability of choosing

[13] The likelihood ratio chi-squared test statistic with two degrees of freedom for the test of constant current segment effects is 14131.46, leading to a rejection of the null of constant coefficients.

any other segment. The effect is most pronounced for the SUV category. Likelihood ratio tests reject the null that expected gas price effects are constant across choice options.[14]

Table 6 : Estimated Effects of Expected Gas Prices on Vehicle Replacement

Replacement	Base Option			
	Midsize	Small	Truck	SUV
Midsize		-0.35	0.09	0.12
		0.031	0.046	0.038
Small	0.35		0.44	0.46
	0.031		0.044	0.037
Truck	-0.09	-0.44		0.03
	0.046	0.044		0.048
SUV	-0.12	-0.46	-0.03	
	0.038	0.037	0.048	

While the estimated effects of persistent preferences and gas price expectations are statistically significant, they are difficult to interpret on their own. To increase interpretability, I run two sets of experiments using the model's parameter estimates. In the first, I consider how an increase in Δp_g^e changes the odds ratio of choosing one replacement segment relative to a base alternative. In the second, I examine how an increase in Δp_g^e is expected to change the sample shares of each replacement vehicle, conditional on currently owned vehicle. In both experiments I ask how stated choices would change if the average forecast increases by one standard deviation. In dollar terms, this equates to an increase from an average forecast of \$0.43/gallon to \$1.15/gallon.

For the first experiment, define the odds ratio as the odds of choosing a specific replacement vehicle j, relative to a specific *base alternative*, b. Now, consider the effect on this odds ratio of increasing Δp_g^e by one standard deviation. In order to generate these perturbed odds ratios, the model needs to be estimated four times - once for each base alternative. If the estimated parameter associated with Δp_g^e for choice j has a p-value of .01 or .05, this is reflected by the corresponding odds ratio in the table below.

[14] The likelihood ratio chi-squared test statistic with two degrees of freedom for the test of constant expected price effects is 10.16, leading to a rejection of the null of constant coefficients.

Table 7 : Experiment 1 - Effects of Increase of Expected Gas Price on Replacement Choice Odds Ratios

Replacement Vehicle	Base Vehicle			
	Medium	Small	Truck	SUV
Medium	--	-22%	7%	9%
Small	28%	--	37%	40%
Truck	-9%	-27%	--	2%
SUV	-8%	-28%	-2%	--

To read the table, consider a base option (column) and consider the effect of a change in Δp_g^e on the probability of choosing a given vehicle (row). We can see that the signs are as expected : increasing gas price forecasts increases the odds of choosing more fuel efficient vehicles relative to less efficient vehicles. The change in odds is especially pronounced for small cars - either as the choice vehicle or the base alternative vehicle. If one looks at the "SUV" column, we see that a one standard deviation increase in Δp_g^e implies that a consumer is 40% more likely to choose a small car relative to the SUV base option. In contrast, looking at the "small" column, we see that the same price increase would make a consumer 28% less likely to choose an SUV relative to the base option of a small car.

It should be noted that the change in odds ratio does not capture the change in probability of choosing a specific replacement vehicle for a given *current vehicle*. The change in odds ratio is in a sense averaged over current vehicles in the sample. With this in mind, I conduct a second set of experiments. I can condition on current vehicle k and consider the change in the probability of choosing j given a one standard deviation increase in Δp_g^e. The result is analogous to a marginal effect, except here we are considering a one standard deviation increase in Δp_g^e, rather than an increase of a unit. We can interpret these changes in conditional choice probabilities as changes in sample shares of replacement vehicles, conditional on current vehicle holdings. Unconditional shares of replacement vehicles can be calculated by taking a weighted average of column entries over the sample frequencies of current vehicles. The unconditional shares are only informative about market behavior if current vehicle holdings in the data are representative of the US vehicle fleet.

We see small but economically significant effects of gas price forecasts on conditional sample shares. Increases in Δp_g^e favor small cars, and to a lesser extent midsize cars. For instance, shares of small cars purchased by current suv holders increase from 9% to 12% of surveyed suv owners. The share of those choosing small cars given midsize cars jumps from 20% to 24.3%. The first table below indicates the raw increase/decrease, the second reproduces the initial stated transitions found in the raw data.

Table 8 : Experiment 2 - Change in Sample Shares of Replacement Conditional on Current Vehicle

Current Vehicle	Replacement Vehicle			
	Medium	Small	Truck	SUV
Medium	-3.4%	4.3%	-0.2%	-0.7%
Small	-2.2%	3.8%	-0.5%	-1.0%
Truck	-0.2%	3.1%	-3.0%	-0.4%
SUV	0.5%	2.9%	-0.1%	-3.3%

Taken together, these two tables are perhaps the most illustrative of the paper's conclusion. Increases in expected gasoline price expectations do increase the probability of choosing small, fuel efficient vehicles. However, the transitions are not dramatic - barely over 3% of current SUV owners shift to smaller vehicles. An increase of one standard deviation in gasoline prices is non-trivial and approximates the large run-up in prices in 2008.

Table 8 : Current vs Replacement Vehicle Segment Transitions (Raw Data)

Current Vehicle	Replacement Vehicle			
	Medium	Small	Truck	SUV
Medium	72%	20%	2%	6%
Small	12%	82%	2%	4%
Truck	7%	10%	75%	8%
SUV	12%	9%	4%	75%

Alternative Specifications

The mechanical details of estimation follow those from the base alternative. The differences are driven by specification and the data to be included. The first alternative specification is a simple extension to the baseline model. I include large/luxury as a segment both for current and replacement segments. This segment was initially excluded because of the small number of observations associated with either choosing or starting with a large/luxury vehicle. As is evident from Table 11, the baseline results are confirmed. Increasing expected gas prices decreases the probability of choosing larger vehicle segments. In addition, there is a high preference to replace with a vehicle which matches one's current segment, relative to all other segments. This specification demonstrates the robustness of my results to the inclusion of segments.

In the next set of alternative specifications, I consider how the inclusion of vehicle characteristics - both of the replacement vehicles and the current vehicle, affect choice probabilities. These alternative specifications allow me to assess the robustness of my conclusions regarding the effects of gas price expectations and persistent vehicle preferences on replacement demand. In the first specification, I consider the impact on model estimates from including characteristics for the intended replacement vehicle. Specifically, I include averages for the manufacturer suggested retail price (MSRP), horse power/weight, vehicle size and miles per gallon by replacement vehicle segments.

Table 12 in the appendix presents the estimated parameters. The estimates on expected gas price confirm results from the baseline specification. The parameter results are significant and indicate that relative to the base option of small cars, an increase in expected gas prices decreases the probability of choosing vehicles from another segment. Similarly, estimated parameters on persistent preferences are consistent with the results from the baseline specification. The estimated coefficient on the vehicle segment where current vehicle is the same as the replacement is always larger than the estimated coefficients for segments where current vehicle does not match the replacement option.

Parameter estimates on the choice-specific vehicle characteristics are less intuitive. The coefficient on the logarithm of MSRP (price) is positive and significant. Similarly inconsistent with what was expected, the estimated parameters for size and horsepower / weight are negative. However, it is necessary to view the coefficients as effects measured relative to the base option of small cars. The small car category consists of vehicles which are the lowest price, smallest in physical size, exhibit the largest horsepower to weight ratio, and the highest MPG. Through this lens parameter estimates make sense. In the case of MSRP, the model uses a positive parameter estimate to reconcile two facts. First, MSRPs were tending upward during the period in question. Given the high degree of persistence in segment demand, this tendency of increasing MSRPs will manifest as a positive price coefficient. Second, the model is trying to reconcile the fact that consumers choose categories other than the cheapest - which happens to be the base category, small vehicles. The model makes this reconciliation by yielding a positive price coefficient.

Applying this perspective to the size and HPW coefficients reconciles the counterintuitive signs. Relative to the physically smallest, highest HPW segment (small vehicles), choices of physically large, low HPW segments can only be rationalized by negative coefficients. Then, the only counterintuitive sign is that associated with miles per gallon. Given that the base category is the most fuel efficient, one might expect a negative coefficient on MPG. In actuality, the positive coefficient can be explained by the fact that MPG is increasing slightly for all vehicles over the time period in question. This fact explains the seemingly spurious coefficient on MPG.

The results and intuition provided for vehicle characteristic coefficients are further corroborated in the model where current vehicle characteristics are included alongside replacement vehicle characteristics (Table 13). Specifically, the intuition for price, size and HPW is confirmed. Looking at coefficients by choice category, we see that the largest, most expensive, and lowest HPW segment is associated with the largest parameter values (in absolute value terms). Further, the coefficients on current MPG confirm that their contribution to choice is either statistically or economically identical to zero. The negligible contribution of MPG to vehicle choice may further be explained by the fact that expected gas prices are soaking up all gas and fuel efficiency relevant effects. Finally, in the specification with both current and replacement vehicle characteristics, we see that expected gas price results are unaffected but that the current segment effects are muddled. This last fact is likely attributable to the difficulty in separately identifying segment dummies and segment characteristics.[15]

In the final alternative specification, I measure whether the inclusion of demographic variables alters the main results. As noted in Table 14, the inclusion of demographics does not alter the results from the baseline specification. The estimated effects of demographics do tell a believable story on their own. Older respondents are more likely to choose midsize cars over small cars. Higher income increases the tendency to choose all other categories over small cars. Finally, the choice of trucks tends to be dominated by males and is decreasing in level of education.

2.4 Conclusion

Accurately measuring consumer responsiveness to gasoline prices is a vital input into policy analysis. Researchers have recognized that long-run response to gas price changes includes changes in vehicle ownership.

[15] Nevo (2000 a), Nevo (2000 b).

Yet, studies examining the dynamics of vehicle adjustment have been stymied by the lack of data on both consumer gas price expectations and current vehicle ownership.

This paper's primary contribution is to introduce and analyze such data. I find that in contrast to the hypothesis of the Energy Paradox literature, low adoption rates for small and fuel efficient vehicles is less likely driven by fuel price expectations and more likely by persistent vehicle preferences. This conclusion also has policy implications.- fuel efficiency incentives for new vehicles should be attached to the trade-in as well as to the new car.

The current analysis was entirely focused on consumer decision-making. However, it would be interesting to examine the impacts of incorporating surveyed expectations on estimates of equilibrium models of the vehicle market. When considering the merits of investing in development of new, fuel-efficient vehicles, car manufacturers make assumptions on adoption rates of their new product offerings. Examining equilibrium outcomes under alternative assumptions of expectation formation / using alternative data for expectations would be an interesting exercise in its own right and would provide a valuable laboratory for policy analysis.

2.5 Appendix

Table 9:Current vs. Replacement Vehicle Expanded Segment Transitions (Raw Data)

Current Vehicle	Replacement Vehicle						
	Large/Luxury	Midsize	Small	Sports/Sporty	Truck	SUV	Van
Large/Luxury	68.42%	14.72%	6.10%	4.51%	1.21%	4.34%	0.70%
Midsize	2.90%	65.56%	18.50%	4.29%	1.91%	5.27%	1.56%
Small	0.70%	11.46%	76.48%	4.27%	2.02%	3.78%	1.30%
Sports/Sporty	2.63%	69.40%	9.66%	72.76%	1.85%	5.96%	0.20%
Truck	1.32%	67.90%	9.60%	2.59%	71.26%	7.75%	0.77%
SUV	3.24%	10.67%	8.19%	4.00%	3.54%	68.72%	1.65%
Van	2.22%	13.88%	10.19%	0.66%	3.62%	12.04%	57.39%

Covariate	Choice							
	Midsize	Small	Truck	SUV	Midsize	Small	Truck	SUV
Expected Price Change		0.35	-0.09	-0.12	0.09	0.44		-0.03
		0.031	0.046	0.038	0.046	0.044		0.048
Currently Small		3.17	1.8	1.41	5.9	4.28		3.23
		0.057	0.143	0.103	0.112	0.11		0.121
Currently Truck		1.61	5.9	2.67	4.1	5.66		2.85
		0.09	0.112	0.104	0.13	0.119		0.142
Currently SUV		1	2.43	4.39	3.47	2.85		5.19
		0.064	0.115	0.068	0.098	0.094		0.089
Constant		-1.44	-3.5	-2.48	-2.4	-2.22		-2.21
		0.036	0.092	0.057	0.069	0.063		0.065
Expected Price Change	-0.35		-0.44	-0.46	0.12	0.46	0.03	
	0.031		0.044	0.037	0.038	0.037	0.048	
Currently Midsize	3.17		1.38	1.77	4.39	3.39	1.95	
	0.057		0.14	0.098	0.068	0.075	0.122	
Currently Truck	1.56		5.66	2.82	2.98	5.15	2.34	
	0.096		0.119	0.112	0.095	0.089	0.142	
Currently SUV	2.18		2.81	5.15	1.72	2.34	5.19	
	0.072		0.128	0.089	0.096	0.092	0.089	
Constant	-1.74		-3.44	-2.81	-1.91	-2.35	-2.98	
	0.048		0.105	0.078	0.041	0.048	0.066	
Observations	20607							
Model Chi-square	24146.31							
Pseudo-R^2	0.43							

Table 10 : Results for Baseline Multinomial Logit Specification. Each of the four tiles represents a model with a different base category specified. The base category is indicated by the column with no reported coefficients. Note, when a base is specified, current segment dummy for that category is also excluded.

Table 11: Results for Alternative Specification with Expanded Segmentation

Covariate	Choice			
	Large/Luxury	Midsize	Truck	SUV
Expected Price Change	-0.49	-0.32	-0.42	-0.45
	0.043	0.028	0.042	0.035
Currently Large/Luxury	7.11	2.79	2.02	2.67
	0.187	0.089	0.187	0.125
Currently Midsize	2.85	3.17	1.38	1.76
	0.191	0.057	0.14	0.098
Currently Truck	2.73	1.56	5.66	2.81
	0.233	0.096	0.119	0.112
Currently SUV	3.78	2.17	2.81	5.15
	0.191	0.072	0.128	0.089
Constant	-4.49	-1.75	-3.45	-2.81
	0.175	0.048	0.105	0.078
Pseudo R^2		0.45		
Model chi-square		34931.3		
N		25047		

Table 12: Results for Alternative Specification with Choice-Specific Vehicle Characteristics (Replacement Characteristics Only)

Covariate	Choice Segment			
	Overall	Midsize	Truck	SUV
log_MSRP	3.62			
	0.502			
HPW	-180.67			
	25.523			
log_Size	-6.83			
	1.483			
MPG	0.16			
	0.029			
Expected Price Change		-0.31	-0.44	-0.46
		0.032	0.044	0.038
Currently Midsize		3.18	1.38	1.76
		0.057	0.14	0.098
Currently Truck		1.57	5.66	2.81
		0.096	0.119	0.112
Currently SUV		2.18	2.81	5.16
		0.072	0.128	0.089
Constant		-1.07	-2.36	-2.08
		0.238	0.57	0.323
Model chi-square		16622.19		
N		82428		

Table 13: Results for Alternative Specification with Choice-Specific Vehicle Characteristics (Replacement & Current Vehicle Characteristics)

Covariate	Choice Segment			
	Overall	Midsize	Truck	SUV
log_MSRP	4.3			
	0.571			
HPW	-172.22			
	30.95			
log_Size	-3.42			
	2.11			
MPG	0.18			
	0.039			
Expected Price Change		-0.31	-0.47	-0.48
		0.032	0.038	0.046
Current HPW		-165.14	-375	-636.87
		43.976	62.659	55.879
Current log_Size		-10.85	-10.71	-14.31
		3.362	3.589	3.489
Current log_MSRP		4.37	9.52	11.45
		0.756	1.015	0.987
Current MPG		0	0.16	-0.1
		0.057	0.088	0.111
Currently Midsize		3.71	1.66	0.53
		0.55	0.688	0.714
Currently Truck		2.13	2.07	0.62
		1.331	1.697	1.73
Currently SUV		1.69	3.77	-0.94
		0.75	1.071	1.231
Constant		66.96	21.63	57.29
		28.584	31.99	29.954
Model chi-square	16368.1			
N	82428			

Table 14: Results for Alternative Specification with Demographic Data

Covariate	Choice		
	Midsize	Truck	SUV
Currently Midsize	3.58	1.83	1.79
	0.114	0.247	0.183
Currently Truck	1.7	6.04	2.97
	0.207	0.226	0.217
Currently SUV	2.14	3.04	5.35
	0.15	0.238	0.165
Expected Price Change	-0.27	-0.2	-0.41
	0.059	0.075	0.07
Male	0.02	0.79	-0.25
	0.114	0.191	0.133
Education	-0.02	-0.2	0
	0.032	0.048	0.038
Age	0.01	0	0
	0.004	0.007	0.005
Income	0.04	0.05	0.06
	0.015	0.021	0.017
Married	-0.05	-0.08	-0.22
	0.114	0.168	0.136
Children at Home	0.04	0.05	-0.25
	0.124	0.186	0.146
Constant	-2.7	-3.38	-2.86
	0.316	0.489	0.381
Pseudo R-squared		0.5	
Model chi-square		8341.81	
N		6099	

2.6 References

1. Allcott, Hunt (2011). "Consumers' Perceptions and Misperceptions of Energy Costs," *American Economic Review : Papers & Proceedings*, 101(3) : pp. 98-104

2. Allcott, Hunt (2012). "The Welfare Effects of Misperceived Product Costs : Data and Calibrations from the Automobile Market," (forthcoming) *American Economic Journal : Economic Policy.*

3. Allcott, Hunt & Nathan Wozny (2012). "Gasoline Prices, Fuel Economy, and the Energy Paradox," unpublished manuscript.

4. Basso, Leonardo & Tae Hoon Oum (2007). "Automobile Fuel Demand : A Critical Assessment of Empirical Methodologies," *Transport Reviews,* 27(4) : pp. 449-484.

5. Beresteanu, Arie & Shanjun Li (2011) "Gasoline Prices, Government Support and the Demand for Hybrid Vehicles in the US," *International Economic Review,* 52(1) : pp. 161-182.

6. Busse, Meghan, Christopher R. Knittel, & Florian Zettlemeyer (2013). "Are Consumers Myopic? Evidence from New and Used Car Purchases," *The American Economic Review.*

7. Berry, Steven T., James Levinsohn & Ariel Pakes (1995). "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4) : pp. 841-890.

8. Gillingham, Kenneth (2013). "Selection on Anticipated Driving and the Consumer Response to Changing Gasoline Prices," unpublished manuscript.

9. Hughes, Jonathan E., Christopher R. Knittel, & Daniel Sperling (2008). "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand.," *The Energy Journal,* 29(1) : pp. 93-114.

10. Jaffe, Adam B. & Robert N. Stavins (1994). "The energy paradox and the diffusion of conservation technology," *Resource and Energy Economics*, 16 : pp. 91-122.

11. Liu, Yizao (2013). "Gasoline Price, Fuel Economy Efficiency and Automobile Replacement Dynamics," unpublished manuscript.

12. Nevo, Aviv (2000 a). "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand," *Journal of Economics & Management Strategy,* 9(4) : pp513-548.

13. Nevo, Aviv (2000 b) "Mergers with Differentiated Products : the Case of the Ready-to-Eat Cereal Industry," *RAND Journal of Economics,* 31(3) : pp. 395-421.

14. Shiraldi, Pasquale (2011). "Automobile Replacement : a Dynamic Structural Approach," *Rand Journal of Economics,* 42(2), pp.266-291

15. Woodall, Bernie, Paul Lienert, & Nen Klayman. "Insight : GM's Volt : The ugly math of low sales, high costs." *Reuters News Service,* September 10, 2012. Retrieved, March 30th, 2013.

    <http://www.reuters.com/article/2012/09/10/...

    us-generalmotors-autos-volt-idUSBRE88904J20120910>