**Title**
Social Implications of Vehicle Choice and Use



**Author**
Langer, Ashley Anne

**Publication Date**
2010

Peer reviewed|Thesis/dissertation

**Social Implications of Vehicle Choice and Use**

by

Ashley Anne Langer

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Economics

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Enrico Moretti, Chair
Professor David Card
Professor Kenneth Train
Professor Catherine Wolfram

Spring 2010

**Social Implications of Vehicle Choice and Use**

## Abstract

Social Implications of Vehicle Choice and Use

by

Ashley Anne Langer
Doctor of Philosophy in Economics

University of California, Berkeley

Professor Enrico Moretti, Chair

This dissertation explores three ways in which consumers' choices about which vehicle to purchase and how much to drive affect others. Chapter 2 seeks to understand how consumer preferences for new vehicles, which may vary in ways that are correlated with consumer demographics, can lead different demographic groups to pay different prices for new vehicles. By estimating consumer preferences for new vehicles by demographic group, I show that dealers do use the distribution of preferences within a demographic group to engage in third-degree price discrimination based on consumer demographics. Additionally, controlling for this third-degree price discrimination, I find that women and single buyers pay more for the same new vehicles than male and married buyers, suggesting that either there are differences in negotiating ability that are correlated with demographics or that dealers are engaging in taste-based discrimination that is not apparent when demographic groups' preferences are not controlled for. Chapter 3, which is coauthored with Nathan Miller, shows that manufacturers adjust new vehicle price incentives in response to changes in gasoline prices in a way that suggests that manufacturers believe that consumers care about vehicle operating costs. We show that these price adjustments would bias earlier estimates of consumer demand for fuel economy that assume that vehicle prices are constant, which implies that consumers have a higher demand for fuel economy than earlier estimates and that the optimal gasoline tax may be lower than earlier estimates. Finally, Chapter 4, which is coauthored with Clifford Winston, estimates the relationship between congestion, urban land use, and home prices in order to understand how congestion tolling would affect urban land use. We find that congestion tolling would lead to denser city centers and subcenters, which would reduce urban sprawl and substantially increase the social benefits of congestion tolling relative to the costs.

For my family: Holly, Steve, Hilary, and Asaf

# Contents

# List of Figures

# List of Tables

# Acknowledgments

This dissertation is the culmination of six years of lessons, conversations, advice, and guidance from a wide range of friends and advisors. As a whole, the Berkeley faculty has been incredible in their advice, both academically and professionally. In particular, my advisor, Enrico Moretti, has been a source of great guidance and feedback. Additionally, David Card has more than lived up to his reputation as an amazing advisor, even for a paper that was slightly out of his field. Kenneth Train was constantly encouraging, especially when computational issues threatened to derail the economics, and Catherine Wolfram helped to keep me grounded in the interesting questions rather than overwhelmed by methodological issues.

Yet one of the amazing things about Berkeley is how faculty who are not your primary advisor are interested in and willing to talk about students' research. Pat Kline, Fred Finan, and Justin McCrary were particularly helpful in getting my job market paper out the door, offering to read drafts or give up computer time when crunch time hit. I'm not sure I'll be able to repay the favor, but I hope to pass on the good karma to my students in the future.

Of course, the major thanks for this dissertation have to go to my co-authors for chapters 3 and 4. Cliff Winston has been a source of encouragement and guideance since I graduated from college and was always there to point out the forest in all those trees. Nate Miller is an amazing researcher and is unbelieveably motivated to move projects from random discussions in the hallways to finished projects that are out the door. I would not have made it nearly this far without either of them.

The ups and downs of grad school are evened out by one's peers, and my peers have provided me endless support both academically and personally. From my first-year study group to my research officemates, countless students have provided insight, feedback, and encouragement along the way. In particular, Willa Friedman, Jane Herr, Brad Howells, Pam Jakiela, Erin Johnson, Matt Levy, Vikram Maheshri, Owen Ozier, Marit Rehavi, Kevin Stange, Erin Syron, and many, many others have read drafts, been sounding boards, and generally listened to me gripe extensively over the last six years, and for that I am greatful.

Finally, there is no real way to explain what my family has meant to me, not only in this crazy endeavor that is grad school, but in the rest of my life as well. My parents have put up with endless phone calls, some much later East Coast Time than than they should be, and still seem to put up with me. Hilary reminds me that there are bigger things out there in the world than academia and my windowless office and she gets my neverending respect for actually getting things done on the ground that I merely intellectualize about. Joan and Fred have been a great local support for dinners and wine tasting and general escapes from the UC bubble. And of course the newest but most immediate member of my family, Asaf is the only one who sees me at my worst both academically and personally but is still supportive and encouraging and wonderful day after day. I love you all and wouldn't be anywhere without you.

# Chapter 1

# Introduction

Personal vehicles play a critical role in the United States economy. In 2008, at the start of a major recession, U.S. consumers still spent nearly $370 billion on over 13 million new vehicles (Parker (2009). These new sales contributed to the approximately 137 million registered vehicles on the road in 2008 that drove a total of 1.6 trillion vehicle miles (2.6 trillion passenger miles).[1] This massive consumption allows drivers to commute to work and school, access shopping and recreational activities, and much more.

Yet there are social reprecussions of this consumption. Pollution from driving is a major issue at both the local and global levels. Personal vehicles consumed 71.5 billion gallons of gasoline in 2008, resulting in 693 million tons of carbon dioxide ($CO_2$) emissions.[2] Overall, personal vehicles account for 19% of total U.S. carbon dioxide emission and 16% of total U.S. greenhouse gas emissions, making transportation the second largest domestic contributor to global warming after energy generation (US Environmental Protection Agency (2010)). In addition, emissions from driving contribute to smog and harmful airborne pollutants at the local level. These emissions are classic externalities in that the consumer's decision to drive does not incorporate the cost to society of the pollution caused by that driving, yet since the consumer reaps the benefits of driving she chooses to drive more than the socially optimal amount.

Emissions are one externality of personal vehicle use, but there are many more. One of the most visible externalities of vehicle use is congestion. Commuters in the 90 largest US cities lost, on average 41 hours to delay in 2007, and commuters in Los Angeles, the most congested city in the country, lost 70 hours.[3] This congestion externality results from the fact that each driver imposes

[1] US Department of Transportation (2008)

[2] 71.5 billion gallons of gasoline times 19.4 pounds of $CO_2$ per gallon.

[3] These numbers define "lost" hours as the difference between the amount of time spent driving and the amount of time that driving would have taken at "free flow" speeds. While

marginal delay on the other vehicles on the road, while only experiencing the average delay. Thus the driver's personal congestion cost at peak periods is less than the cost that she is imposing on the other drivers on the road. While this congestion externality leads to excess time spent driving, the fact that personal commute cost is less than the social cost of that commute means that commutes are longer than is optimal and cities are more sprawling than is optimal, both of which further exacerbate the pollution externality of driving.

It is important to understand the implications and magnitudes of these externalities of personal vehicle choice and use in order to understand what policies might bring personal incentives more in line with social incentives, thus protecting the environment and improving economic efficiency. This dissertation focuses on three particular ways in which the consumer decisions about which vehicle to purchase and how much to drive it have implications for the rest of society. Chapter 2 seeks to understand how heterogeneity in consumers' preferences for new vehicles across demographic groups can lead different demographic groups to pay different prices for the same new vehicles. The Chapter 3, which is coauthored with Nathan Miller, explores manufacturer's beliefs about consumer demand for fuel economy and highlights the implications of those beliefs for the estimation of the optimal gasoline tax. Finally, the Chapter 4, which is coauthored with Clifford Winston, asks how an optimal congestion toll would affect urban land use and calculates how the inclusion of land use affects the costs and benefits of congestion tolling. Each chapter takes a different component of consumer demand for new vehicles or their use and seeks to understand the implications of these choices for society and continue the discussion on the optimal policies for aligning personal and social incentives.

---

this is a reasonable measure of delay, it does not imply that free flow speeds are the efficient allocation of scarce roadways. Thus these numbers overestimate the extent of the congestion externality. The congestion externality would be measured by the difference between the cost of the delay that a driver imposes on others and the benefit to the driver of using that road at that time.

# Chapter 2

# Demographic Preferences and Price Discrimination in New Vehicle Sales

Understanding why different demographic groups pay different prices is central to questions of consumer welfare and equity in many markets. In a monopolistically competitive setting like the new car market, sellers have an incentive to charge higher prices to consumers with more inelastic demand. Unlike seller animus or differences in bargaining skills, such "third-degree" price discrimination implies a distinctive pattern of *product-specific* price differentials across groups. This chapter proposes and implements a simple test for the importance of third-degree price discrimination in the new vehicle market in the U.S. Specifically, I use micro data for a large sample of recent buyers to estimate separate random-coefficient discrete choice models for married and unmarried men and women, and calculate optimal markups for each group. Across 230 different vehicle models I find that observed price differences between groups closely track the predicted relative markups: a one-dollar increase in the predicted relative markups leads to a 30-45 cent rise in relative prices. This suggests that firms are partially successful in discriminating by gender and marital status, although arbitrage across groups or lack of co-ordination between dealers limits the extent of this discrimination. The estimates imply that the elimination of third-degree price discrimination would reduce the consumer surplus of single women by 5.6% of their total vehicle expenditure, and raise the surplus of married men by 5.4% of their total expenditure.

## 2.1 Introduction

Why do different demographic groups pay different prices for the same goods? Three primary explanations have been proposed in the economics literature, with different implications for consumer welfare and equity. Sellers may dislike interacting with certain groups and in equilibrium charge them higher prices (Becker, 1957). Alternatively, some groups may have different negotiating abilities and therefore pay more for all vehicles (Babcock and Laschever, 2003). Finally, sellers may use consumers' demographics to infer their preferences and practice third-degree price discrimination. In a market with a homogeneous product, all three explanations predict a difference between groups in average prices and are hard to distinguish empirically. In a differentiated product market, however, third-degree price discrimination implies a distinctive pattern of product-specific price differences across groups that are related to their relative demand elasticities. Thus, even if average prices paid by women (for example) reflect some degree of seller animus, or limited negotiating skills, the impact of third degree price discrimination can be identified from the relative variation of prices and elasticities across different products.

This paper tests for the presence of third-degree price discrimination in the market for new automobiles by estimating group-specific product demand functions and comparing observed relative prices paid by different groups to the predicted relative markups implied by the model. The new vehicle market is well suited to this investigation because prices are set by individual negotiation and vary across consumers, there are many different (but closely substitutable) models, and different buyers appear to have distinct valuations for each product attribute. Moreover, previous research has shown that a random-coefficient model of consumer choice among vehicle models can capture many of the most important features of the structure of demand in this industry.[1] When combined with a simple model of price-setting, these models yield good descriptions of the determinants of price and quantity in the market.

This approach extends the literature on discrimination in new vehicle sales. Earlier papers have focused on measuring the difference in the average price paid by demographic groups. This work has included both paired audit studies (Ayers and Siegelman (1995)) and cross-sectional investigations (Goldberg (1996)). Researchers have looked at whether dealer profit varies over consumers of different demographic groups (Harless and Hoffer (2002)) and whether negotiating online rather than in person changes prices paid differentially across demographic groups (Scott Morton, Zettelmeyer, and Silva-Risso (2003)). All of these papers have acknowledged that while they are measuring differences in the treatment

[1] Berry, Levinsohn, and Pakes (1995, henceforth BLP), Berry, Levinsohn, and Pakes (2004, henceforth MicroBLP)

of different demographic groups in the market, there are multiple explanations for why those differences may arise. Yet these papers look at the average price paid (or profit made) by demographic group controlling for attributes, rather than looking at the product-specific price differences across demographic groups. My approach of estimating demand functions for different demographic groups investigates price differences at a finer level than has been done in the past.[2] This approach could also be useful for understanding the determinants of price differences in markets like housing (Yinger (1998)) and loans (Charles, Hurst, and Stephens (2008)).

My estimation relies on the assumption that the demographics of the person purchasing the new vehicle are the same as the demographics of the person whose preferences lead to the vehicle choice. If consumers rely on friends or family members to purchase a new vehicle for them, this assumption could be problematic. I circumvent this concern by using a unique survey dataset of new vehicle purchasers from the second quarter of 2005. I use only those 10,703 new vehicle purchasers who confirmed that they were "both the principle buyer and driver" of one of 230 new vehicles, provided full demographic information, and reported the price they paid for their new vehicles. I augment this data with information from the Current Population Survey on the total number of Americans in each demographic group in order to include information on consumers who chose not to purchase a new vehicle in the second quarter of 2005.

Using this dataset, I estimate separate random coefficient discrete choice models for married women, married men, single women, and single men.[3] Consumers are assumed to have utility specifications similar to Berry, Levinsohn, and Pakes (1995, henceforth BLP) and Berry, Levinsohn, and Pakes (2004, henceforth MicroBLP), but estimation is via maximum likelihood in the style of Train and Winston (2007) paired with the Berry (1994) inversion. This approach allows each demographic group to value all vehicle attributes, including those unobserved to the econometrician, differently from other demographic groups. It also allows for consumer heterogeneity within each demographic group. Using these demand estimates, I calculate the optimal markup for each firm to charge each demographic group for each vehicle. If firms can perfectly price discriminate between demographic groups, then the pattern of differences in a vehicle's price between any two demographic groups should exactly equal the pattern of differences in the predicted markup between the groups.[4] I estimate the level of

[2] Goldberg (1995) looks at the variance in prices paid by demographic groups, but does not look at the prices for each vehicle individually.

[3] Data limitations require that the demographic groups be large in order to facilitate estimation, which excludes some interesting demographic variables, such as race, from analysis.

[4] This assumes that each vehicle's marginal cost is the same regardless of the demographic group that purchases the vehicle.

effective price discrimination: the extent to which firms can convert differences in predicted markups between demographic groups into differences in observed transaction prices. The discrete choice framework also allows me to calculate the change in consumer surplus for each demographic group that would result from the elimination of third-degree price discrimination.

I find that preferences vary substantially across demographic groups. On average, women are more price sensitive than men and single consumers are more price sensitive than married consumers. Some of this difference is driven by income differences: on average men are from richer households than women and married people are from richer households than single people. All demographic groups substitute substantially between vehicles within the same vehicle type (car, truck, SUV, or van), but married women strongly prefer SUVs to cars, while single women have the opposite preference, on average. Men, both single and married, prefer vehicles with high curb weight, although the preference heterogeneity for curb weight indicates a preference for both large, heavy vehicles, and lighter, sportier cars. Women, on the other hand, are fairly indifferent to curb weight after controlling for other vehicle characteristics and do not have much heterogeneity in their taste for curb weight. I discuss the extent of preference differences between groups in more detail in section 2.5.

Using this variation in preferences, I find optimal product markups for each demographic group that are consistent with earlier results in BLP and MicroBLP. Married men have the highest optimal markups, averaging approximately 40% of transaction prices, while single women have the lowest optimal markups, averaging approximately 20% of transaction prices. These markups are consistent with the estimates prepared for the US Environmental Protection Agency on the ratio of total vehicle price to vehicle costs.[5]

When I compare the differences in predicted optimal markups across demographic group pairs to the differences in observed average prices, I find that firms do engage in third-degree price discrimination. A $1 increase in the difference in optimal markups between two groups leads to a statistically significant 30 to 45 cent increase in the difference in observed average prices. Additionally, once preference differences between demographic groups are considered, women and single buyers appear to pay more on average for new vehicles than their male and married counterparts. This leaves open the possibility that animus and differences in the taste for negotiation are operating simultaneously with third-degree price discrimination in this market.

To understand the impact of third-degree price discrimination on the consumer surplus of each demographic group, I use the estimated demand functions

[5]In a report for the US Environmental Protection Agency, RTI International (2009) calculates weighted average markups of approximately 32% for new vehicles based on the observed prices and costs in the industry.

to ask how the consumer surplus of each group would change in the absence of third-degree price discrimination. I find that eliminating third degree price discrimination would increase the consumer surplus of married men by 5.4% of their total new vehicle expenditures and decrease the consumer surplus of single women by 5.6% of their total new vehicle expenditures, thus explicitly hurting the groups that are already paying more for new vehicles conditional on optimal markups.[6]

The remainder of the paper is organized as follows: in the next section I describe the literature on price discrimination. In section 2.3 I describe my empirical specification. The data used is explained in section 2.4. I then present results of the demand estimation and comparison of optimal markups to observed prices in section 2.5 and the results of the consumer surplus calculations in section 2.6. Section 2.7 concludes.

## 2.2 Literature Review

The literature on price discrimination has moved from theoretically proving the potential for price discrimination in monopolistically competitive markets to empirically investigating the causes of price differences between demographic groups. Beginning in the late 1970s, authors such as Salop and Stiglitz (1977) showed that price discrimination could exist even with competitive firms if consumers had different preferences for search. Later, Borenstein (1985) and Holmes (1989) formally showed that price discrimination is possible in differentiated product environments.

However, it wasn't until the early 1990s that the literature began to empirically show that price discrimination might exist in actual markets. Some of the first empirical papers (Borenstein (1991), Shepard (1991)) used the market for gasoline to show that differences in prices between leaded and unleaded or full and self-service gasoline might in fact be driven by price discrimination. At the same time, other authors (Lott and Roberts (1991)) questioned whether observed differences in prices were the result of price discrimination or unobserved cost differences, foreshadowing a debate over the causes of price differences that continues today.

Issues of fairness often arise when price discrimination is correlated with consumer demographics. Becker's 1957 theories of taste-based discrimination provided a theoretical foundation for research into price (or wage) discrimination based on animus towards a certain demographic group. In his model with perfect competition for identical employees, the only reason for differences in prices

[6] Results of the change in consumer surplus calculations, in both dollars and as a percent of total demographic group expenditures, are presented in Table 2.7.

would be employer animus. However, in imperfectly competitive markets or those with differentiated products or employees, the issue of untangling the cause of price differences between demographic groups has been central to discussions of the policy implications of demographically-based price discrimination. Phelps (1974) provided a simple model of statistical discrimination in which employers use consumer demographics to predict unobservable employee productivity but argued that regardless of the cause of discrimination, policy might still aim to eliminate it in the marketplace.

Identifying price discrimination based on consumer demographics became a popular topic in the late 1990s when the literature embarked on both cross-sectional and experimental investigation of price differences across demographic groups. In the cross-sectional literature, Goldberg (1996) studied car discounts by race and gender and Graddy (1997) investigated fast food prices in neighborhoods with different demographics. Scott Morton, Zettelmeyer, and Silva-Risso (2003) asked whether price differences between demographic groups decrease when consumers purchase online. Yinger (1998) summarizes a body of literature measuring price discrimination in housing using both cross-sectional and experimental data. Paired audit studies, in which experimental consumers of different demographic groups but with otherwise similar observable attributes attempt to purchase a car or rent a home provided a different approach. Ayers and Siegelman (1995) use a paired audit to investigate discrimination in car pricing. Heckman and Siegelman (1993) and Heckman (1998) discuss the advantages and disadvantages of paired audits at length, including discussions of the fact that statistical discrimination may still be the cause of differences between pairs and that discrimination in the market as a whole is determined by the marginal seller or employer, which paired audits may not capture.[7]

While these studies often find differences across demographic groups in the price paid (or negotiated), the cause of those price differences is generally unclear. Potential explanations for differences in prices include animus against certain demographic groups and profit-maximizing third-degree price discrimination, but also differences in preferences for negotiation (Babcock and Laschever (2003), Gneezy, Leonard, and List (2009)) and differences in consumer behavior in the market (Akerlof and Kranton (2000)). Researchers have attempted to untangle these explanations with varying degrees of success. Graddy (1995) finds that Asians pay less than whites for fish in the Fulton Street Fish Market and hypothesizes that price sensitivity and negotiating power are stronger explanations than racial animus given that all of the fish sellers are white. Altonji and Pierret (2001) focuses on the role of employer learning about worker productivity to better understand statistical discrimination in hiring. List (2004) is the first to attempt

[7]Empirical investigations of discrimination at the margin include Charles, Guryan, and Pan (2009) and Charles and Guryan (2008).

to experimentally disentangle the causes of price differences across groups. He convincingly argues that statistical discrimination, rather than animus, is what drives minority groups to pay more for a particular sports card at trade shows.

In addition to understanding why different demographic groups may pay different prices in a market, a literature builds on the original work of Schmalensee (1981) who calculated the welfare implications of third-degree price discrimination in a monopoly. Corts (1998) provides a scenario under which prices could fall to all consumers under price discrimination, leading to a positive consumer surplus effect for all consumers and a negative profit effect for firms. A newly-developing empirical literature attempts to understand how price discrimination affects total welfare in multi-firm markets. Graddy and Hall (2009) estimates a structural model of the Fulton Street Fish Market to attempt to understand how welfare would change if sellers were forced to post prices. In a demand analysis similar to this paper, Villas-Boas (2009) estimates a single discrete choice model of demand to better understand the welfare implications of banning wholesale price discrimination.

## 2.3 Empirical Specification

The empirical model aims to compare the firms' optimal price for each demographic group under third-degree price discrimination to the observed price for that group in order to identify the extent of price discrimination in the market. This requires estimating the vehicle demand functions of each demographic group for each vehicle and pairing these estimates with a model of manufacturer and dealer pricing behavior. I will describe the demand and supply approaches before turning to my test for third-degree price discrimination.

### 2.3.1 Demand Functions

The demand function follows directly from MicroBLP but is estimated using maximum likelihood as in Train and Winston (2007). I use the Berry (1994) inversion to reduce the dimensionality of the coefficient space.

Consumers are each assumed to belong to a single demographic group, $d = 1, ..., D$. Within these demographic groups, consumers are heterogeneous along both observable and unobservable individual characteristics. Consumer $i$'s utility for vehicle $j = 0, 1, ..., J$ is assumed to be:

$$U_{idj} = p_{jd}\tilde{\alpha}_{id} + \sum_k x_{jk}\tilde{\beta}_{idk} + \xi_{dj} + \epsilon_{idj} \tag{2.1}$$

where $p_{jd}$ is the price charged to $i$'s demographic group $d$; $x_{j1}, ..., x_{jK}$ are the non-price attributes of vehicle $j$; $\xi_{dj}$ is the preference of demographic group $d$ for the unobservable attributes of vehicle $j$; and $\epsilon_{idj}$ is an extreme value type 1 residual preference parameter. The $\tilde{\alpha}_{id}$ and $\tilde{\beta}_{idk}$ are the individual's preference for vehicle attributes $p_{jd}$ and $x_k$ respectively, which are assumed to have the form:

$$\tilde{\alpha}_{id} = \bar{\alpha}_d + \sum_r z_{idr}\alpha^o_{dr} + \nu_{idp}\alpha^u_d \quad (2.2)$$
$$\tilde{\beta}_{idk} = \bar{\beta}_{dk} + \sum_r z_{idr}\beta^o_{dkr} + \nu_{idk}\beta^u_{dk}$$

Thus the individual's preference for vehicle attribute $x_k$ is decomposed into a component ($\bar{\beta}_{dk}$) that is constant within that individual's demographic group, a component ($\beta^o_{dkr}$) that varies with consumer characteristics $z_{idr}$ that are observed by the econometrician,[8] and a component ($\beta^u_{dk}$) that varies with consumer characteristics $\nu_{idk}$ that are unobserved by the econometrician, but are assumed to have a known distribution.[9] These unobserved consumer characteristics capture the fact that there is heterogeneity in preferences for different vehicle attributes in every demographic group, although we may not have variables that allow us to identify those consumers who get particular utility from horsepower, for instance, rather than side air bags.

Combining equations (2.1) and (2.2) leads to the consumer's choice model:

$$U_{idj} = \delta_{dj} + \sum_r p_{jd}z_{idr}\alpha^o_{dr} + \sum_{k,r} x_{jk}z_{idr}\beta^o_{dkr} + p_{jd}\nu_{idp}\alpha^u_d + \sum_k x_{jk}\nu_{idk}\beta^u_{dk} + \epsilon_{idj} \quad (2.3)$$
$$\text{where} \qquad \delta_{dj} = p_{jd}\bar{\alpha}_d + \sum_k x_{jk}\bar{\beta}_{dk} + \xi_{dj} \quad \text{for each} \quad j = 1, ..., J \quad (2.4)$$

The consumer chooses the vehicle $j = 1, ..., J$ or the outside option ($j = 0$, not purchasing a new vehicle) that maximizes this utility function.[10] As this notation makes clear, there is a component ($\delta_{dj}$) to each individual's utility for each vehicle that is common across all members of his or her demographic group $d$. Additionally, the term $\sum_r p_{jd}z_{idr}\alpha^o_{dr} + \sum_{kr} x_{jk}z_{idr}\beta^o_{dkr}$ allows consumers with different observable characteristics to have different tastes for certain vehicle attributes, and thus specifies the extent to which vehicle substitution varies with

[8] I assume that the dealer does not observe the consumer characteristics $z_{idr}$, and therefore they are not allowed to enter into the vehicle's price.

[9] I will generally assume that unobserved consumer characteristics have normal distributions.

[10] The outside option of not purchasing a new vehicle is assumed to have utility equal to $U_{id0} = \beta^u_{d0}\nu_{id0} + \epsilon_{id0}$, where $\nu_i$ is a draw from a standard normal distribution and $\epsilon_{id0}$ is a draw from an EV1 distribution.

observable consumer demographics. Finally, there is a component of consumer preference $(p_{jd}\nu_{idp}\alpha^u_{dk} + \sum_k x_{jk}\nu_{idk}\beta^u_{dk})$ that is unobserved by the econometrician but helps explain why certain consumers have stronger preferences for some vehicle attributes rather than others, and helps to explain why individuals may substitute more strongly between certain vehicles. The $\beta^u_{dk}$ and $\alpha^u_d$ coefficients can be thought of as representing the standard deviation in the unobserved preference within demographic group $d$ for the vehicle attribute. For notational ease, I define the vector of distributional coefficients $\theta_d \equiv [\alpha^o_{dr},\ \alpha^u_{dr},\ \beta^o_{dkr},\ \beta^u_{dkr}]'$.

I estimate the $\theta_d$ and $\delta_d$ coefficients via maximum-likelihood. The extreme-value error term guarantees that the probability of vehicle $j$ maximizing consumer $i$'s utility conditional on the observable attributes of the vehicle $(p_{jd}, x_{jk})$ and the consumer's observable $(z_{idr})$ and unobservable $(\nu_{id} = [\nu_{idp}, \nu'_{idk}]')$ characteristics is:

$$Pr_{idj}(p_{jd}, x_{jk}, z_{idr}, \nu_{id}; \theta_d, \delta_d) \quad = \frac{\exp(V_{idj}(p_{jd}, x_{jk}, z_{idr}, \nu_{id}; \theta_d, \delta_d))}{\sum_{l=0}^{J} \exp(V_{il}(p_{ld}, x_{lk}, z_{idr}, \nu_{id}; \theta_d, \delta_d))} \tag{2.5}$$

where $V_{idj}(p_{jd}, x_{jk}, z_{idr}, \nu_{id}; \theta_d, \delta_d)$ is the non-stochastic component of consumer $i$'s utility for vehicle $j$ from equation (2.3). To condense notation, I will write $V_{idj}(\nu_{id}; \theta_d, \delta_d)$ and $Pr_{idj}(\nu_{id}; \theta_d, \delta_d)$ with the understanding that the non-stochastic utility and probability are also a function of the observable data. Because $\nu_{id}$ is unobserved to the econometrician, the expected value of the probability unconditional on $\nu_{id}$ is:

$$Pr_{idj}(\theta, \delta_d) \quad = \int \frac{\exp(V_{idj}(\nu_{id}; \theta_d, \delta_d))}{\sum_{l=0}^{J} \exp(V_{idl}(\nu_{id}; \theta_d, \delta_d))} f(\nu) d\nu \tag{2.6}$$

where again $Pr_{idj}(\theta_d, \delta_d)$ is understood to also be a function of the observable data.

Because the $\theta_d$ coefficients determine how consumers substitute between vehicles as attributes change, information on consumers' first and second choice vehicles aids identification of $\theta_d$. Thus, the joint probability that consumer $i$ chooses vehicle $j = 1$ out of the full choice set, and $j = 2$ out of the choice set with $j = 1$ and the outside good removed is:[11]

[11] I remove the outside good from the second-choice choice set because the second choice information is based on the vehicle the consumer said she considered but did not purchase. It is not clear whether she would have purchased the second choice vehicle if the first choice were not available (she may not have purchased any vehicle), but it is her preferred alternative out of the set of vehicles once her first choice is removed.

$$
\begin{aligned}
&Pr_{i1}(\theta_d, \delta_d)Pr_{i2}(\theta_d, \delta_d|1) = \\
&\qquad \int \frac{\exp(V_{id1}(\nu_{id}; \theta_d, \delta_d))}{\sum_{l=0}^{J}\exp(V_{idl}(\nu_{id}; \theta_d, \delta_d))} \left( \frac{\exp(V_{id2}(\nu_{id}; \theta_d, \delta_d))}{\sum_{l=2}^{J}\exp(V_{idl}(\nu_{id}; \theta_d, \delta_d))} \right) f(\nu) d\nu \quad (2.7)
\end{aligned}
$$

Since the probability of observing a particular first and second choice for an individual is conditional upon the individual's $\nu_{id}$ vector, the integration over the distribution of $\nu$ must be over the joint probability of the first and second vehicle choices. I approximate this integral using simulation, and then sum the log of this probability over consumers $i$ in demographic group $d$ to calculate the log-likelihood function.[12]

The log-likelihood function is maximized over $\theta_d$. For each value of $\theta_d$, I choose $\delta_d$ to set the predicted market shares for each demographic group equal to the observed market shares for that group as in Berry (1994):

$$
\begin{aligned}
S_{dj} &= \int_{z_{idr}} \int_{\nu} Pr_{idj}(\theta_d, \delta(\theta_d)) f(\nu) f(z_{idr}) d\nu dz_{idr} \qquad (2.8) \\
&= Pr_{dj}(\theta_d, \delta(\theta_d))
\end{aligned}
$$

where $f(z_{idr})$ is the pdf of the consumer characteristics $z_{idr}$ in the demographic group $d$. Therefore it should be understood that the $\delta_d$ vector is estimated conditional on $\theta_d$ and is thus formally $\delta(\theta_d)$. The maximum-likelihood procedure solves for the value of $\theta_d$ that maximizes the likelihood function subject to a market-share constraint that is a function of both $\theta_d$ and $\delta(\theta_d)$.

This model differs from previous random-coefficient demand models in an important way: the preferences of each demographic group $d$ are assumed to be completely independent of the preferences of every other demographic group.[13] While this means that demographic groups may value the observable (to the econometrician) attributes of the vehicles differently, it is particularly important that the unobservable (to the econometrician) characteristics of a vehicle ($\xi_{dj}$) are allowed to be valued differently by members of different demographic groups. A prime example of such varying preference for vehicle unobservables would be the vehicles that are commonly referred to as "chick cars" or "guy cars"[14] such that the opposite gender might be interested in the vehicle for its physical attributes, but dissuaded from buying the car because of its social connotation. Additionally,

[12] Simulation uses 50 scrambled Halton draws to approximate the integral for each consumer.

[13] This also means that the preferences in the population (as estimated in the previous discrete choice literature) are a combination of the preferences in each demographic group.

[14] See, for instance: http://www.cartalk.com/content/features/Guy-Chick-Cars/index.html

options packages that appeal to one group rather than another (for instance spoilers or wheel rims) would potentially change the unobservable quality of the car for different groups differently.

Once I have estimated $\theta_d$ and calculated $\delta(\theta_d)$, I can use the $\delta_d$ vector to extract information about the $\bar{\alpha}_d$ and $\bar{\beta}_{dk}$ coefficients rather than just the $\theta_d$ coefficients. Recall that:

$$\delta_{dj} = p_{jd}\bar{\alpha}_d + \sum_k x_{jk}\bar{\beta}_{dk} + \xi_{dj}$$

The problem is that unobservable vehicle quality, $\xi_{dj}$ may include vehicle attributes that allow firms to charge more for the vehicle. Therefore, an OLS regression of the $\delta_{dj}$ vector on vehicle price and attributes will estimate that consumers are less price sensitive than they actually are. In order to correct for this bias, I run a weighted IV regression of $\delta_{dj}$ on the vehicle price and attributes.[15] I use the standard Bresnahan (1987)/BLP instruments:

$$\sum_{l \in f_j, l \neq j} x_{lk} \qquad \text{and} \qquad \sum_{l \notin f_j} x_{lk} \tag{2.9}$$

which are the sum of each vehicle attribute for competing vehicles produced by the same firm as vehicle $j$, $f_j$, and the sum of each vehicle attribute for competing vehicles produced by other firms. These instruments are intended to capture the extent of price competition faced by vehicle $j$ in the market. For instance, if a vehicle is competing with a set of vehicles that have particularly high horsepower, then competitive pressure will keep the vehicle's price fairly low conditional on its attributes. If the observed price is actually high conditional on attributes, it must be that the vehicle has a high level of of unobservable quality that is increasing its demand.

Because demographic groups face different prices and value vehicle attributes differently, the competitive pressure on price created by competing vehicles' attributes should vary over demographic groups. Therefore the instrumental variables regression is run separately for each demographic group.

The estimated demand coefficients allow me to calculate demand elasticities. Because the predicted demand of demographic group $d$ for vehicle $j$ is just the number of people in group $d$ times the predicted market share of group $d$ for vehicle $j$, the own-price elasticity of demand is just:

[15] The weights are equal to the number of consumers of demographic group $d$ who chose vehicle $j$, which is an approximation of the inverse of the variance of $\hat{\delta}_{dj}(\theta_d)$ from the maximum-likelihood estimation.

$$\frac{\partial Pr_{dj}(\hat{\theta}_d)}{\partial p_{dj}} \left( \frac{p_{dj}}{Pr_{dj}(\hat{\theta}_d)} \right) \tag{2.10}$$

which is straightforward to calculate given $\hat{\theta}_d$. This is the key formulation of demand for firms that are choosing prices to maximize profits.

### 2.3.2 Supply

I pair this demand specification with a model of vehicle supply that closely follows BLP, MicroBLP and Bresnahan (1981). Firms maximize profits over the set of vehicles they offer, which are assumed to have constant marginal cost. The equilibrium is Nash in prices. The complication from the standard supply model is that firms choose an optimal price for each vehicle to offer to *each* demographic group. I explicitly assume that "firms," which include both the manufacturer and its dealer network, are able to charge optimal prices to each demographic group. This means that there is perfect contracting between the manufacturer and its dealers; dealers are perfectly able to identify the demographic group of each consumer, and consumers cannot engage in across-demographic-group arbitrage. This final assumption is stronger than the typical no-arbitrage assumption where consumers are assumed to not participate in a secondary resale market.[16] In this case consumers are also assumed not to obscure their demographic characteristics by sending someone of a different demographic group to purchase the new vehicle for them.[17]

Thus firms $f = 1, ..., F$ set prices to maximize profits over the vehicles they sell:

$$\pi_f = \sum_{d=1}^{D} \sum_{j \in f} Q_{dj}(p_d)(p_{dj} - c_j - D_d) \tag{2.11}$$

where the demand function of demographic group $d$ for vehicle $j$, $Q_{dj}(p_d)$, is a function of the vector of the demographic group's prices for all vehicles, $p_d$. I allow for the possibility of animus or differences in average bargaining abilities by including a fixed cost of selling to demographic group $d$, $D_d$. The maximization

[16] Note that the used car market would not function as a secondary resale market in this case because the price difference between a new and a barely used vehicle is substantial. In particular, the gain from reselling a car to a different demographic group is small relative to the loss of selling a "used" car rather than a new one.

[17] Recall that the in the estimation I will use data on those survey respondents who said that they are both the principle buyer and driver of the new vehicle.

of this set of profit functions for all firms leads to the vector of optimal prices given the vector of marginal costs, $c$:

$$
\begin{aligned}
P_d^* &= c - \Omega_d^{-1} Q_{dj} + D_d && (2.12)\\
&\equiv c + M_d + D_d && (2.13)
\end{aligned}
$$

where $P_d^*$ is the optimal price vector for group $d$, $M_d$ is the vector of optimal markups, and $\Omega_d$ is the matrix of own and cross-price derivatives of demand:

$$
[\Omega_{djk}] = \begin{cases} \frac{\partial Q_{dk}(\theta_d, p_d)}{\partial p_{dj}} & \text{if } j \text{ and } k \in F \\ 0 & \text{otherwise} \end{cases}
$$

From the demand estimation, I have estimates of $\hat{\theta}_d$ and $\bar{\alpha}_d$, and I can therefore construct estimates of the demographic group's demand and price derivative matrix, $Q_{dj}(\hat{\theta}_d)$ and $\Omega_d(\hat{\theta}_d, \bar{\alpha}_d)$.[18] Thus I have enough information to construct estimates of the optimal markups for each vehicle $j$ sold to demographic group $d$, $\hat{M}_{dj}$. While I do not have information on the costs of vehicle $j$, I do assume that the marginal vehicle costs are the same for all demographic groups, and therefore that the difference in the optimal price between demographic groups is equal to the difference in the optimal markup between groups plus any difference in the animus between groups: $P_A^* - P_B^* = M_A - M_B + D_A - D_B$.

This supply model assumes that the firms are able to observe the demographic groups, $d$, perfectly, but that they don't observe any other consumer characteristics, such as those included in the $z$ vector in the demand specification. I will discuss this assumption, along with the potential for consumers to obscure their demographic group, in the context of my specific choice of $d$ and $z$ characteristics.

### 2.3.3 Understanding Price Discrimination

This relationship between the prices paid by different demographic groups provides the basis for my estimation of the extent of third-degree price discrimination in the market. While there may be many other considerations in price setting, the extent to which observed price differences track differences in the

[18] In BLP and MicroBLP, the authors use a moment similar to equation 2.12 to estimate their model, allowing the cost of each vehicle to be a linear combination of the vehicle's observed attributes. I do not exploit this moment, and therefore do not assume that the observed prices, $p_{dj}$, are optimal. This leaves open the possibility of animus or differences in average bargaining ability across groups in the observed data.

predicted optimal markup is a measure of effective third degree price discrimination. Therefore, in order to understand the extent to which observed price differences between demographic groups follow differences in predicted markups between groups, I run the regression:

$$\bar{p}_{Aj} - \bar{p}_{Bj} \quad = \gamma_0 + \gamma_1(\hat{M}_{Aj} - \hat{M}_{Bj}) + e_j \tag{2.14}$$

where $\bar{p}_{Aj}$ (or $\bar{p}_{Bj}$) is the average transaction price for vehicle $j$ for demographic group $A$ ( or $B$), and $e_j$ is the measurement error in the predicted price differences plus prediction error in $\hat{M}_{Aj} - \hat{M}_{Bj}$ and unmodeled variation in the difference in average vehicle prices between groups. In this regression, the coefficient of primary interest is $\gamma_1$, which is the effective amount of third-degree price discrimination, or the amount of the difference in the optimal markup across groups that the firms are able to extract from consumers. If the assumptions of the supply model (perfect contracting between manufacturers and dealers, perfect identification of consumers' demographic groups, and no arbitrage) hold exactly, then I would expect $\gamma_1 = 1$. However, there are reasons to believe that these assumptions will not hold perfectly. Competition between dealers of the same vehicles is unlikely to be perfectly contracted away, especially as internet retailing and phone negotiations replace face-to-face dealer interactions. Individual dealers may know the price that they are supposed to charge a given consumer for a given vehicle, but may be tempted to deviate from that price when the consumer threatens to walk out of the showroom. Additionally, the dealers may not be able to perfectly identify a consumer's demographic group either because a married consumer may arrive at the showroom without her spouse or because she may send a spouse or family member with different demographics to purchase the vehicle in the hopes of getting a better deal. The breakdown of these assumptions will lead firms to have an effective rate of third-degree discrimination that is less than 1. In the extreme, firms would be unable to engage in third-degree price discrimination and I would estimate a coefficient of $\gamma_1$ that is indistinguishable from 0.

The identification of $\gamma_1$ hinges upon the assumption that $\hat{M}_{Aj} - \hat{M}_{Bj}$ is uncorrelated with $e_j$. The primary concern would be that there is unmodeled variation in animus or bargaining across vehicles that would bias $\hat{\gamma}_1$ upwards, making it appear that firms are engaging in third-degree price discrimination when they are not. Yet for this to be true, differences in animus between groups would have to vary with the difference in the demographic groups' optimal markups. While there may be models in which this occurs, those models of animus are very different from the Becker (1957) model upon which most models of animus or taste-based discrimination are based. Alternatively, $\hat{\gamma}_1$ would be biased upwards if consumers exerted different amounts of effort in bargaining depending upon their preference for the vehicle. Again, this is a very different model of bargaining than one in which consumers have different tastes for bargaining or negotiation skills, which would predict that consumers who are particularly skilled bargainers would pay lower prices for all vehicles. A secondary concern may be that there is measurement error in the $\bar{p}_{Aj} - \bar{p}_{Bj}$ that is correlated with the $\hat{M}_{Aj} - \hat{M}_{Bj}$. Again, this would require a very unique pattern of measurement error in the average prices, which seems unlikely in this context.

The $\gamma_0$ coefficient serves multiple purposes in this regression. First, it measures the average price difference between the demographic groups for a vehicle with identical optimal markups. In that role, it captures any animus discrimination and differences in average bargaining ability between the two demographic groups. However, $\gamma_0$ is also the intercept in a linear regression that would capture any systematic differences in the ability of dealers to effectively price discriminate against a particular demographic group.[19] Finally, the intercept would include any differences in the average cost of selling to the two demographic groups that has not been modeled. These roles complicate attempts to interpret $\gamma_0$ as the difference in demographic group animus and average bargaining ability.

## 2.4 Data

The primary data for this analysis is a survey of new vehicle buyers conducted by a major market research firm. This data is augmented with data from the Current Population Survey, the Automotive News Market Data book, and Autodata Solutions.

The survey of new vehicle buyers includes 25,875 respondents who purchased new vehicles in the second quarter of 2005.[20] The survey includes information on the model of vehicle purchased and the other models considered,[21] but does not include information on the trim level or the options packages of the vehicle. The survey asks respondents a series of questions about their purchase, including the price they paid for the vehicle, and whether they paid cash for the vehicle, leased it, or secured a loan. Additionally, respondents indicated their age, gender,

[19]For instance, if the total amount of the possible price difference between groups were capped by consumers' arbitrage opportunities, then the effective price discrimination would not be perfectly linear and $\gamma_0$ may not equal 0 even in the absence of animus.

[20]Because the survey is limited to consumers who purchased new vehicles in the second quarter of 2005 I abstract from concerns about the changing demographic sales patterns over the calendar year that are raised in Aizcorbe, Bridgman, and Nalewaik (2009).

[21]I will follow the standard practice of assuming that the other models considered are listed in the order in which they were considered in order to identify the consumer's second choice vehicle. I only use the second choice information rather than the third and fourth because the number of respondents who entered a third or fourth choice is low.

marital status, education, household income, and race on the survey.[22] In a particularly relevant question, the survey asks whether the respondent is both the "principle buyer and driver" of the vehicle. 21,085 respondents indicated that he or she was both the principle buyer and driver, and I will limit my analysis to these respondents in order to assure that the demographic information matches the driver of the vehicle and the person who physically purchased the vehicle.[23]

My analysis will focus on four demographic groups: married women, married men, single women, and single men. These groups are large enough to estimate demand functions for each. Gender and marital status are attractive groups to use for this analysis because they are fairly evenly distributed geographically, so it is likely that all dealerships interact with consumers of all demographic groups. Additionally, gender is a readily observable variable to dealers and is often thought of as a dimension along which vehicle preferences may vary. Marital status may be less observable to dealers, so any differences in the amount of price discrimination based on marital status relative to gender might be related to consumers' ability to obscure their demographic group. Additionally, to the extent that married consumers are more likely than single consumers to be older and have larger households that potentially include children, I would expect married consumers' preferences to differ from single consumers of the same gender.[24]

I remove from consideration any consumers who purchased a vehicle with an average sales price of over 75 thousand dollars in order to limit the analysis to commonly purchased vehicles. In order to calculate prices for each demographic group for every vehicle, I only include vehicles which at least one survey respondent of each demographic group purchased. When combined with the restriction that all of the relevant questions were answered, these restrictions bring my dataset down to 10,703 consumers. 58% of my sample is male and 64% is married.[25] New car buyers tend to be wealthier than the average American, with

[22] Consumers were asked to indicate the range in which their education and household income fell, rather than the exact amount.

[23] Of course, many people may take a friend or family member with them to purchase a vehicle, in which case the dealer may not be completely sure who the primary driver of the vehicle will be.

[24] These groups have the advantage of being fairly observable to dealers, but gender and marital status are clearly only a subset of the demographics that a dealer may observe or infer. In this analysis, differences in average income, age, education, and race across these four demographic groups will enter into the mean preference coefficients, $\bar{\alpha}$ and $\bar{\beta}$. I use household income as an observed determinant of consumer heterogeneity within demographic groups, but assume that prices are set for the demographic group as a whole rather than for different income classes within the demographic group. Age, education, and race differences within demographic groups will contribute to unobserved consumer heterogeneity while differences across demographic groups will enter into the mean preference coefficients.

[25] Of the 10,735 observations in my sample, 2,513 are married women, 4,314 are married men, 1,950 are single women, and 1,958 are single men.

38% of respondents coming from households making over \$100,000 and only 25% coming from households making less than \$50,000. 53% of respondents in my sample have a college degree.

By limiting the data to those consumers who are both principle buyers and drivers of their new vehicles, I do introduce some selection into my analysis. This selection is most likely strongest for married consumers, who have another adult in the household who might negotiate for the new vehicle in the consumer's place, while single consumers may not have another adult who could easily replace him or her in purchasing the vehicle. In fact, of the respondents with complete data, 22% of married men report not being the principle buyer and driver and 11% of married women report not being the principle buyer and driver. For single men, only 5% of respondents are not the principle buyer and driver, and for single women 4% are not the principle buyer and driver. If a survey respondent reports that she is not the principle buyer and driver of the new vehicle, it is impossible to tell whether she is the principle buyer but someone else is driving the vehicle or whether someone else purchased the vehicle for her to drive. Research on women's propensity to avoid negotiation (Babcock and Laschever, 2003) would indicate that married men may be purchasing vehicles for their wives and then sometimes filling out the accompanying survey. This might mean that the women who purchase vehicles for themselves are particularly strong negotiators, leading selection to bias my results towards finding men paying more for all new vehicles. Regardless of the predicted sign of the selection bias, the fact that single people generally are both the principle buyer and driver of the new vehicle would indicate that selection will be low for single consumers.

Before discussing my treatment of consumers who purchase the outside good, Table 2.1 displays the correlations between the log of the price paid and the consumer demographic groups. All of the regressions include fixed effects for each vehicle model to control for differences in marginal costs. The coefficients show that there is almost no significant difference between demographic groups in the mean price paid for new vehicles. Single women pay slightly more than married men for new vehicles on average, but the coefficient is only significant at the 10% level. No other coefficients are statistically significant. These results are similar to Goldberg (1996) and Harless and Hoffer (2002), yet, because of third-degree price discrimination, the absence of a statistically significant difference in the average price paid does not imply that there is no animus in this market. If, for instance, single women were willing to pay less than married men for new vehicles on average, then with third-degree price discrimination we would expect their average price paid to be less than married men's. The fact that the two groups pay approximately the same amount would then indicate that animus or bargaining differences are driving up the price paid by single women.

The other type of selection that may enter my analysis is the selection in-

Table 2.1: Correlations Between Log Price Paid and Demographic Group

| Dependent Variable: Ln(Price Paid) | (1) | (2) | (3) |
|---|---|---|---|
| Female | 0.0034 | | |
| | (0.0043) | | |
| Single | | 0.0070 | |
| | | (0.0049) | |
| Single Female | | | 0.0103* |
| | | | (0.0059) |
| Single Male | | | 0.0038 |
| | | | (0.0073) |
| Married Female | | | -0.0001 |
| | | | (0.0055) |
| Vehicle Fixed Effects | 230 | 230 | 230 |

Regression of ln(price paid) on consumer characteristics and vehicle fixed effects. Robust standard errors in parentheses. Coefficients marked with a * are significant at the 10% level. All regressions include 10,703 observations.

curred by limiting the dataset to only those consumers who provide full price and demographic information on the survey. If a consumer feels that she got a particularly bad deal on a car she might choose to "forget" how much she paid when it comes time to complete the survey. As long as this selective omission is similar for different demographic groups and similar over vehicles for which a group has different preference intensities, I would not expect missing data to impact my conclusions. However, future work might benefit from using a dataset that is matched to actual transaction data both to remove the possibility of missing values and to confirm that consumers' recollections of the price they paid for their new vehicles matches the actual transaction price.[26]

In order to account for customers who decided not to purchase a new vehicle in the second quarter of 2005, I append observations to my sample with consumers from each demographic group who purchased the outside good. This approach is valid given four conditions:

1. The population distributions of the observed consumer characteristics, $d$ (demographic group) and $z$ (other consumer characteristics), are known.

2. The consumer characteristics, $d$ and $z$, have discrete distributions.

[26]The Scott Morton, Zettelmeyer, and Silva-Risso (2003) dataset would avoid the issues of self-reported prices and has substantially more observations, but does not have information directly from the consumer on demographics, consumers' second choices, or whether a consumer is both the principle buyer and driver.

3. The fraction of non-buyers is known for each $\{d, z\}$ cell.

4. Non-buyers all receive utility $U_{id0} = f(z_{id}) + \epsilon_{id0}$.

Working through the requirements, I begin by using information from GfK Automotive Research which says that approximately 20% of Americans considered buying a new vehicle in the previous year. I therefore assume that 10% of Americans considered buying a new vehicle in the second quarter of 2005.[27] I assume that that consumer characteristics of this 10% of Americans mirror the attributes of the American population as a whole, as measured by the Current Population Survey for May 2005.[28] These assumptions give me the population distributions of $d$ and $z$. Since $d$ is by definition a discrete demographic group and $z$ is the consumer's response to discretely-valued survey questions, the consumer characteristics have discrete distributions.

In order to know the fraction of non-buyers in each $\{d, z\}$ cell, I start with information from the Automotive News Market Data Book on the total number of vehicles of each model sold in the second quarter of 2005.[29] I assume that the distribution of consumer characteristics for purchasers of each vehicle is the same as the distribution of consumer characteristics for purchasers of that vehicle in my survey data. By summing the characteristics of consumers over the total number of vehicles sold in the quarter, I then know the number of new vehicle purchasers in each $\{d, z\}$ cell. Thus I have both the total number of consumers in each cell who considered purchasing a new vehicle and the total number who did purchase a new vehicle. The difference is the weight I place on the non-purchase observation for that consumer characteristic cell, thus satisfying requirement 3 above. Finally, I assume that all non-buyers receive utility $U_{id0} = \beta^u_{d0}\nu_{id0} + \epsilon_{id0}$.[30] Thus my data satisfies the four requirements above and the data augmentation procedure gives me a sample of vehicle purchasers and non-purchasers.

I pair this data with data from AutoData Solutions on the attributes of model year 2005 vehicles. This data includes extensive information on the vehicle, including the manufacturer's suggested retail price (MSRP), horsepower, curb

[27] Assuming 5% or 15% of Americans considered buying a new vehicle in the second quarter of 2005 does not change the results substantially.

[28] This assumption mirrors assumptions in BLP and MicroBLP that all Americans consider buying a new vehicle each year, but limits the population to a more relevant group. The other extreme assumption would be that consumers who consider buying a new vehicle have the characteristics of the consumers who do buy new vehicles. I will test this alternative assumption in a future version of this paper.

[29] Note that in the second quarter of the year almost every vehicle has the model year equal to the calendar year, which alleviates the issue of the mix of model years of vehicles sold.

[30] In a specification test, I allowed non-buyers' utility to be $\beta^o_d(\frac{1}{\text{Income }_i}) + \epsilon_{id0}$, and although some estimated coefficients changed, the final estimates of the extent of third-degree price discrimination were quite similar.

weight, wheel base, fuel economy, turning radius, and whether the vehicle has stability control, traction control, or side airbags. This data is at the vehicle trim level, which allows it to differ for the same vehicle model based on differences such as engine type (e.g. V6 vs V8) or body style (e.g. hatchback vs sedan). Since my consumer choice data only specifies a consumer's purchase decision at the model level, I use the vehicle attributes of the trim with the lowest MSRP as the model attributes and consider any deviations from this unobserved quality. This reinforces the idea that consumers of different demographic groups might have different valuations of unobserved quality, since not only the vehicle's styling may be valued differently but also the average trim level chosen may vary by demographic group. To the extent that many options such as leather seats, rear spoilers, or sunroofs may be fairly inexpensive to produce but command a large markup, these options packages may be a way for firms to encourage consumers to self-select into options packages that are priced to further price discriminate.[31]

## 2.5 Results

The results are presented in three steps: the demand coefficients are presented first and include the $\delta$s (the mean preference of each demographic group for each vehicle), mean preference coefficients (how different vehicle attributes contribute to the $\delta$ vector for each demographic group), and the coefficients governing observed and unobserved preference heterogeneity within demographic groups. I then present the elasticities and optimal markups that are calculated from these demand coefficients, and finally I compare the predicted markup differences between pairs of demographic groups to the observed average price differences for those groups.

### 2.5.1 Demand Estimation Results

The $\delta$ vector of mean preference parameters contains the values of the mean preference for each vehicle for each demographic group that set the predicted market share for each vehicle equal to its observed market share, conditional upon the consumer heterogeneity coefficients. In this respect, the $\delta$ vector acts as an adjusted market share, where the adjustment comes from the fact that some consumers with extreme preferences will buy certain vehicles frequently enough to match a substantial portion of the vehicle's market share even when the

[31] Although this is technically second-degree price discrimination, where firms configure product offerings such that consumers will sort by willingness to pay, I will estimate it as a part of what I call third-degree price discrimination. Ireland (1992) provides an interesting discussion of the identification of this type of second-degree price discrimination using the assumption that costs are linear in vehicle attributes.

mean consumer strongly dislikes the vehicle. A good example of such a scenario occurs with extremely expensive luxury sedans. While the average consumer of any demographic group would find such vehicles far too expensive, there are some consumers in each group with a low price sensitivity or a high demand for vehicle performance who find these cars attractive. The extent to which these adjustments in market shares occur can be measured by the correlation between the $\delta$ vector and the log of observed market shares. I find that for all four demographic groups these correlations are between 0.636 and 0.820. Single women have the lowest correlation at 0.636, indicating that heterogeneity appears to be relatively more important in explaining their observed market shares, while married men have the highest correlation at 0.820. Single men have a higher correlation between $\delta$ and observed market share (0.818) than married women (0.713).[32]

These mean preference coefficients order vehicles by the preference of the average consumer of each demographic group, making them a useful reality-check before moving on to the coefficient estimates. The highest and lowest five vehicles for each demographic group are listed in Table 2.2. For both single men and married men, the five highest mean preference vehicles are all pickup trucks. Women, however, vary much more substantially by marital status. The top five vehicles for married women are all SUVs, while three of the five *lowest* preference vehicles for single women are SUVs (and the other two are vans). Single women prefer sedans, with the Toyota Camry Sedan topping the list. At the bottom of all demographic groups' lists are luxury cars and SUVs that most likely appeal to a wealthy minority. For instance, married women have both the Cadillac Escalade Sport Utility Truck (a SUV with a pickup truck bed) and the Hummer H2 Sport Utility Truck in the bottom five. As expected, three of the four groups (all except for single women) have the ultra-expensive Audi A8 in the bottom five. Generally, these results correspond with our expectations about the types of vehicles that different demographic groups prefer.

[32] All correlations are for the 230 vehicle choices excluding the outside good, which is standardized to have a $\delta$ of zero.

Table 2.2: Mean Preferences by Demographic Group

| | Demographic Group | | | |
|---|---|---|---|---|
| | Married | | Single | |
| | Women | Men | Women | Men |
| Top Five Vehicles: | Ford Escape | Chevrolet Silverado | Toyota Camry Sedan | Toyota Tacoma |
| (in order, | Jeep Liberty | Ford F-150 | Chevrolet Cobalt | Chevrolet Colorado |
| highest to lowest) | Chevy Equinox | Toyota Tacoma | Nissan Altima | Chevrolet Silverado |
| | Toyota Highlander | Chevrolet Colorado | Honda Accord Sedan | Ford F-150 |
| | Hyundai Tucson | Dodge Ram Pickup | Toyota Corolla | GMC Sierra |
| Bottom Five Vehicles: | Chevrolet SSR | Audi A8 | Toyota Land Cruiser | Audi A8 |
| (in order, | Audi A8 | Jaguar XJ | Infiniti QX56 | Jaguar XJ |
| lowest to highest) | Cadillac Escalade EXT | Lexus SC 430 | Lexus LX 470 | Audi Allroad |
| | Hummer H2 SUT | Mercury Monterey | Mercury Monterey | Lexus LX470 |
| | Chevy Corvette | Jaguar S-Type | GMC Safari | Lexus SC430 |

Vehicles are ordered by the estimated mean preference of each demographic group, controling for heterogeneity in group preference. The mean preference is estimated using the Berry (1994) inversion in the maximum likelihood estimation of the coefficients that describe within-demographic group heterogeneity in preferences.

Regressing these deltas on vehicle attributes using weighted instrumental variables generates the mean preference coefficients for each demographic group. In these regressions, I include price (instrumented with the BLP instruments as discussed earlier)[33] and the vehicle types with cars as the excluded group. Because the $\delta$ vector is scaled such that the $\delta$ for the outside good is zero and not included in the regression, the constant term captures the preference for cars relative to the outside good. I include vehicle attributes including fuel use, curb weight, horsepower, and turning radius (which can proxy for the inverse of vehicle performance) in the mean preference specification.

The results of the mean preference regression are reported in Table 2.3. Married women are more price sensitive than married men, single women are more price sensitive than single men, and single people of either gender are more price sensitive than their married counterparts. This most likely reflects the fact that single people generally have lower household incomes than married people and women have lower household incomes than men.

While all groups other than single women prefer SUVs to cars, married women have a particularly large preference for SUVs. Similarly, men have a strong preference for pickups over cars while women are more indifferent. Although no group significantly prefers vans to cars, married women do have a fairly high van coefficient relative to other groups. All of the groups except single women have negative, significant constant terms, reflecting the low numbers of new vehicle buyers conditional on income and choice set attributes in any given quarter.

In terms of vehicle attributes, all groups dislike high fuel use vehicles.[34] Men like high curb weight vehicles and women are indifferent. Single women have a surprisingly strong preference for horsepower, and married men have a significant preference for low turning radius vehicles, although other groups have similar estimates of their preferences that are not statistically significant.

The specification of consumer demand heterogeneity is similarly comprised of three sets of coefficients. I specify price as having a normally distributed unobservable heterogeneity component as well as a component that varies with the demeaned inverse of consumers' household income. This captures the fact that price sensitivity may depend upon the vehicle's price relative to a consumer's income. I then include normally distributed unobservable heterogeneity terms for

[33] I do not use the instruments constructed from every mean preference variable. The van instrument has very little variation such that it is primarily picking up whether the vehicle is produced by a major manufacturer. I exclude the curb weight instrument because the combination of curb weight, horsepower, and fuel use are nearly colinear, and the deviations from colinearity are likely picking up some of the unobserved quality of the vehicle.

[34] Single females have an estimated mean preference for fuel use that's of the same magnitude as the other groups, but the high standard error means that the estimate is not statistically significant.

Table 2.3: Mean Preference Coefficients by Gender and Marital Status

| | Gender and Marital Status | | | |
|---|---|---|---|---|
| Variable | Married Females | Married Males | Single Females | Single Males |
| Price | -1.55*** | -1.34*** | -1.90*** | -1.75*** |
| (tens of thousands of dollars) | (0.43) | (0.29) | (0.58) | (0.33) |
| SUV | 4.06*** | 1.90*** | -3.58*** | 1.12** |
| | (0.54) | (0.41) | (0.58) | (0.46) |
| Pickup | 0.05 | 3.30*** | -0.38 | 4.09*** |
| | (0.72) | (0.48) | (1.33) | (0.77) |
| Van | 1.33 | -0.32 | -0.15 | 0.40 |
| | (1.22) | (1.00) | (3.95) | (2.16) |
| Constant | -4.82*** | -6.87*** | 0.94 | -6.29** |
| | (1.41) | (1.31) | (3.96) | (3.20) |
| Fuel Use | -65.09** | -48.85*** | -46.43 | -64.18*** |
| (gallons per 100 miles) | (30.37) | (14.31) | (51.17) | (16.35) |
| Curb Weight | 0.34 | 1.33** | -0.05 | 1.11* |
| (thousands of pounds) | (1.02) | (0.59) | (1.12) | (0.63) |
| Horsepower | -0.09 | 0.47 | 2.26*** | 0.42 |
| (hundreds) | (0.53) | (0.38) | (0.87) | (0.56) |
| Turning Radius | -0.15 | -0.24** | -0.28 | -0.21 |
| (feet) | (0.11) | (0.10) | (0.22) | (0.20) |
| Number of Observations | 230 | 230 | 230 | 230 |

Instrumental variables regression of the mean preference of each group for each vehicle on vehicle characteristics. Instruments are functions of the vehicle attributes of competing vehicles, as in BLP. Weighted instrumental variables standard errors in parentheses, where the weights are equal to the number of observations for that demographic group that purchased that vehicle in the maximum likelihood stage. Significance level indicated by: *=10%, **=5%, ***=1%.

the vehicle's type (SUV, pickup truck, van, car) and the outside good. This set of coefficients allows consumers to substitute more intensely within vehicle types, even conditional on vehicle attributes, and makes the model a generalization of a nested-logit framework (e.g. Goldberg (1998)). Finally, I allow for normally distributed unobservable heterogeneity in each demographic group's demand for vehicle attributes including fuel use (measured in gallons per hundred miles), curb weight, horsepower, and whether the vehicle has side air bags. The coefficients on all of these normally distributed unobservable heterogeneity terms can be interpreted as the standard deviation in the demographic group's preference for the vehicle attribute.

Table 2.4 presents the coefficient estimates for these consumer heterogeneity terms for all four demographic groups. I find that all groups except married females have small but statistically significant heterogeneity in their price preference, even after controlling for income differences. All groups except for single women exhibit substantial heterogeneity in price preferences based on income, and the signs are as expected: wealthier consumers are less price sensitive than average and poor consumers are more price sensitive than average. Consumers of all demographic groups display high variation in their preference for different types of vehicles, which indicates that consumers of all groups substitute substantially within vehicles of the same type. Of particular interest is the fact that married women show relatively little heterogeneity in their demand for SUVs and single men show relatively little heterogeneity in their demand for pickup trucks. Since I estimated that both groups have fairly high mean preferences for these vehicle types, the lack of heterogeneity in these preferences indicates that the demographic group is surprisingly united in its taste for these types of vehicles. The final component of the vehicle nests is the heterogeneity in preference for the outside good. Only married men and single women appear to have substantial heterogeneity in their demand for the outside good relative to a new vehicle, which may be a result of these groups' incomes being at the two extremes of the distribution.

Table 2.4: Preference Heterogeneity Coefficients by Gender and Marital Status

| | | Gender and Marital Status | | | |
|---|---|---|---|---|---|
| Variable | Variable Type | Married Females | Married Males | Single Females | Single Males |
| Price | Std Dev | 0.06 | 0.13*** | 0.13** | 0.37*** |
| (tens of thousands of dollars) | | (0.12) | (0.03) | (0.06) | (0.07) |
| | Divided by | -6.20*** | -5.79*** | 0.15 | -5.28*** |
| | Income | (0.88) | (0.35) | (0.44) | (0.61) |
| SUV | Std Dev | 0.60*** | 2.29*** | 3.76*** | 2.24*** |
| | | (0.21) | (0.23) | (0.67) | (0.59) |
| Pickup | Std Dev | 3.59* | 2.80*** | 1.84*** | 0.36 |
| | | (2.13) | (0.27) | (0.47) | (0.22) |
| Van | Std Dev | 3.81*** | 3.85*** | 2.81*** | 2.39** |
| | | (0.63) | (0.51) | (0.85) | (0.99) |
| Car | Std Dev | 4.07*** | 4.22*** | 1.14*** | 4.26*** |
| | | (0.48) | (0.43) | (0.16) | (0.73) |
| Outside Good | Std Dev | 0.38 | 0.55*** | 0.64** | 0.43 |
| | | (0.27) | (0.17) | (0.32) | (0.26) |
| Fuel Use | Std Dev | 0.24** | 0.33*** | 0.23*** | 0.36*** |
| (gallons per hundred miles) | | (0.10) | (0.04) | (0.03) | (0.09) |
| Curb Weight | Std Dev | 0.21 | 0.21*** | 0.01 | 0.63*** |
| (thousands of pounds) | | (0.17) | (0.05) | (0.10) | (0.13) |
| Horsepower | Std Dev | 1.18*** | 0.02 | 0.01 | 0.37** |
| (hundreds) | | (0.22) | (0.03) | (0.07) | (0.19) |
| Side Air Bag Dummy | Std Dev | 0.10 | 0.09 | 0.35 | 0.26 |
| | | (0.25) | (0.15) | (0.31) | (0.29) |
| Number of Observations | | 2513 | 4314 | 1950 | 1958 |

Standard errors in parentheses. Significance level indicated by: *=10%, **=5%, ***=1%. Coefficients estimated with maximum likelihood.

Table 2.5: Estimated Elasticities and Markups

| | | Gender and Marital Status | | | |
|---|---|---|---|---|---|
| Variable | Statistic | Married Females | Married Males | Single Females | Single Males |
| Elasticity | Mean | -3.25 | -2.81 | -5.80 | -3.87 |
| | Min | -6.57 | -4.71 | -13.66 | -5.42 |
| | Max | -1.70 | -1.66 | -2.41 | -1.99 |
| $ Markup | Mean | 10,114 | 11,689 | 5,537 | 8,470 |
| | Min | 7,324 | 7,308 | 5,219 | 5,727 |
| | Max | 15,292 | 18,413 | 6,238 | 16,676 |
| % Markup | Mean | 35.76 | 40.62 | 20.37 | 29.22 |
| | Min | 15.44 | 21.54 | 7.37 | 19.29 |
| | Max | 61.01 | 70.75 | 41.91 | 51.46 |

Descriptive statistics are over the 230 vehicles in the sample. All numbers are calculated using the demand coefficients presented in tables 2.3 and 2.4. Percent Markup is the markup divided by the average price paid for the vehicle by that demographic group.

Finally, I estimate varying amounts of heterogeneity in consumer demand for different vehicle attributes. Consumers of all demographic groups exhibit heterogeneity in their demand for fuel use, which might not be surprising in a country where Toyota Priuses and Chevrolet Suburbans increasingly share the road. Men, both married and single, display heterogeneity in their preference for curb weight, a likely result of heterogeneity in their preferences for large, heavy trucks and SUVs and smaller, sportier performance cars. Perhaps similarly, married women exhibit large differences in their demand for horsepower, perhaps displaying variation in preference for power and ease of driving. There is not much heterogeneity within any demographic group in the demand for side air bags.

### 2.5.2 Elasticities and Optimal Markups

When converting these coefficient estimates into estimated markups, a useful statistic with some economic intuition is the aggregate own-price demand elasticity for each demographic group for each vehicle. The first three rows of Table 2.5 give some descriptive statistics on the distribution of own-price elasticities across vehicles. Generally, these elasticities average from just under 3 (in absolute value) for married men to almost 6 for single women. These appear to be of the same magnitude as the elasticities reported in MicroBLP.

The second and third panels of Table 2.5 provide descriptive statistics for the vehicle markups for each demographic group in terms of both dollars and as a percent of the transaction price. As we would expect given the elasticities, single women have the lowest average elasticities and married men have the highest. The markups average between 20 and 40 percent of transaction prices. There are a few reasons to think that markups of this magnitude would be reasonable. First, a 2009 report for the US Environmental Protection Agency used manufacturer cost information to calculate that automobile costs should be inflated by 1.46 to estimate retail prices, which implies an average markup of 31.5% (RTI International (2009)). Additionally, this is an industry with high fixed costs for each model produced, and these are markups over marginal costs. Therefore, we would expect a firm to only bring a vehicle to the market if it believed that the vehicle would have high marginal profits.

### 2.5.3 Comparing Actual and Predicted Price Differences

The goal in calculating these optimal consumer markups is to understand whether firms are engaging in third-degree price discrimination between demographic groups. I test this by comparing the observed difference in average prices for a pair of demographic groups to the predicted price difference. Recall from Section 2.3 that, under third-degree price discrimination,

$$\begin{aligned} p_{jA} - p_{jB} &= C_j + M_{jA} + D_A - (C_j + M_{jB} + D_B) \\ &= (D_A - D_B) + M_{jA} - M_{jB} \end{aligned}$$

where $p_{jA}$ ( or $p_{jB}$) is the price charged for vehicle $j$ to demographic group $A$ ( or $B$). $M_{jA}$ is the optimal markup for vehicle $j$ to demographic group $A$, and $D_A$ is animus or bargaining differences that change the prices of all vehicles. When estimated with average prices and predicted optimal markups, the regression function is

$$\bar{p}_{jA} - \bar{p}_{jB} = \gamma_0 + \gamma_1(\hat{M}_{Aj} - \hat{M}_{Bj}) + e_j$$

where the $e_j$ captures the measurement error in the average price difference, the estimation error in the predicted markup difference and any unmodeled variation in prices. The measurement error in the average price difference has a variance that is proportional to the number of purchasers of vehicle $j$ from demographic groups $A$ and $B$, and I therefore weight the each observation in this regression function by $\frac{1}{N_{Aj}+N_{Bj}}$ where $N_{Aj}$ ( or $N_{Bj}$) is the number of consumers of group $A$ ( or $B$) who purchased vehicle $j$.[35] The $\gamma_1$ term measures the effective price discrimination allowing for the fact that firms may not be able to extract the

[35]Since $\hat{M}_{jA}$ and $\hat{M}_{jB}$ are consistently estimated, measurement error in the $\hat{M}_{jA} - \hat{M}_{jB}$ term will be small relative to the measurement error in the average price difference.

Figure 2.1: Price and markup Comparison for Married Men vs Married Women

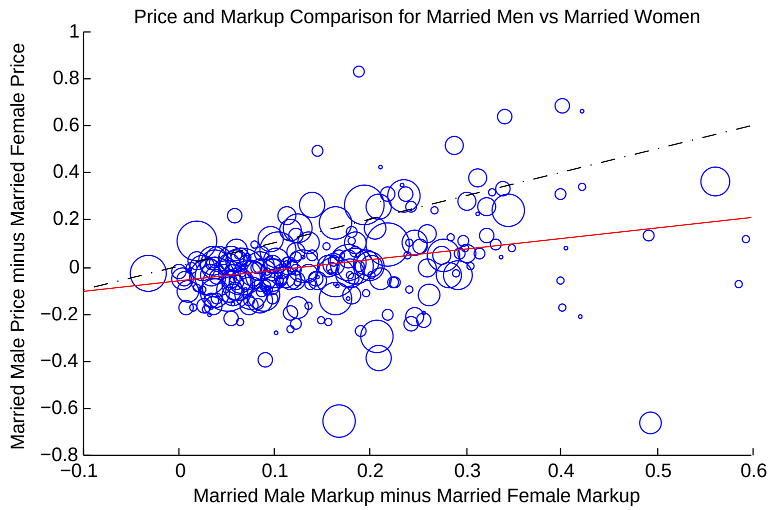


Figure displays 230 vehicle observations. Circle sizes indicate the weight on the observation, $\frac{1}{N_{Aj}+N_{Bj}}$, where group $A$ is married men and group $B$ is married women. Dashed line is the 45° line. The solid line is the best fit line described in table 2.6.

full third-degree price discriminating markup difference from consumers. If I can reject the null hypothesis that $\gamma_1 = 0$, I will say that dealers engage in third-degree price discrimination, and if I reject the secondary null hypothesis that $\gamma_1 \geq 1$, I will say that dealers are unable to perfectly third-degree price discriminate. The $\gamma_0$ term will capture both any animus discrimination or bargaining differences that remain once I control for the differences in optimal markups and any systematic differences across demographic groups in the ability to price discriminate between the demographic groups.[36]

Table 2.6 gives the estimates of $\gamma_1$ and $\gamma_0$ for four demographic group pairs. The first set of results is for married people, comparing the prices and markups of married men minus married women. Figure 2.1 shows the plot of the weighted data and the regression line with the dotted 45 degree line included for reference. The slope of the regression line is 0.450, meaning that for every dollar increase in the difference between married men and married women's optimal markups, firms are able to increase the difference between the groups' average price paid by 45 cents. This coefficient is statistically different from both zero and one, so firms are engaging in third-degree price discrimination but are not able to extract the full third-degree price discriminating markup differences.

[36]For instance, if dealers are supposed to charge higher prices to one group than another but are capped in the total price difference between the groups that they can extract without enducing arbitrage, the $\gamma_0$ coefficient may capture aspects of the nonlinearity of effective price discrimination.

Table 2.6: Comparison of Observed Price Differences to Predicted Markup Differences

| | Markup Difference ($\gamma_1$) | Intercept ($\gamma_0$) |
|---|---|---|
| **Difference by Gender:** | | |
| Married Men minus Married Women | 0.450*** | -0.059*** |
| | (0.106) | (0.018) |
| Single Men minus Single Women | 0.315** | -0.066** |
| | (0.130) | (0.031) |
| **Difference by Marital Status:** | | |
| Married Women minus Single Women | 0.367*** | -0.151*** |
| | (0.127) | (0.053) |
| Married Men minus Single Men | 0.417*** | -0.145*** |
| | (0.106) | (0.038) |

These regressions estimate the coefficients in equation 2.14. The dependent variable is the difference in the averge price for each vehicle between the two demographic groups, $\bar{p}_{Aj} - \bar{p}_{Bj}$. The markup difference is the difference in the estimated optimal markup for each vehicle to each group, $\hat{M}_{Aj} - \hat{M}_{Bj}$. Each regression is over 230 vehicle choices. Weighted standard errors in parentheses, where the weight is equal to $\frac{1}{N_{Aj}+N_{Bj}}$ and $N_{Aj}$ ($N_{Bj}$) is the number of observations of consumers in group $A$ ($B$) who purchased vehicle $j$. Significance level indicated by: **=5%, ***=1%. All variables are in tens of thousands of 2005 dollars.

Figure 2.2: Price and Markup Comparison for Single Men vs Single Women

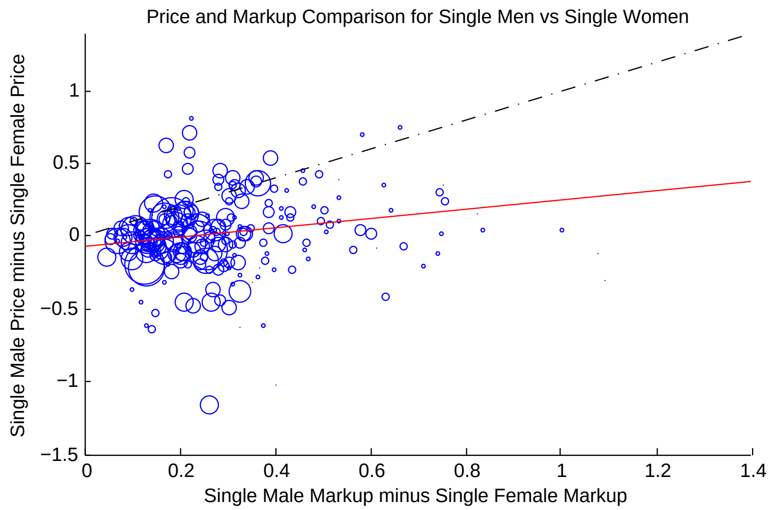


Figure displays 230 vehicle observations. Circle sizes indicate the weight on the observation, $\frac{1}{N_{Aj}+N_{Bj}}$, where group $A$ is single men and group $B$ is single women. Dashed line is the 45° line. The solid line is the best fit line described in table 2.6.

The remainder of Table 2.6 gives the regression coefficients for single men minus single women, married women minus single women, and married men minus single men. Figures 2.2, 2.3, and 2.4 represent these results graphically. While the first two sets of results obviously represent discrimination based on gender holding marital status constant, the second two represent discrimination based on marital status holding gender constant. The recurring result is that firms do engage in third-degree price discrimination, although not to the full extent predicted by the model. None of the slope coefficients are statistically different from any of the others, but all of them are statistically different from 0 and 1. Thus, third-degree price discrimination based on consumer demographics does contribute to the differences between demographic groups in the prices paid for new vehicles.

The fact that the $\gamma_1$ estimates are statistically indistinguishable across demographic groups suggests that consumers may not be engaging in extensive arbitrage in this market. Married consumers, by definition, have another adult in the household who is typically a member of a different demographic group. Thus married consumers should have lower costs than single consumers of sending a member of a different demographic group to purchase a vehicle for them. This lower cost is reflected in the differential rates at which married consumers report being the principle buyer and driver of their new vehicles (recall that married consumers are 2 to 4 times more likely than single consumers to report that they are not both the principle buyer and driver). Yet the effective price discrimina-

Figure 2.3: Price and Markup Comparison for Married vs Single Women

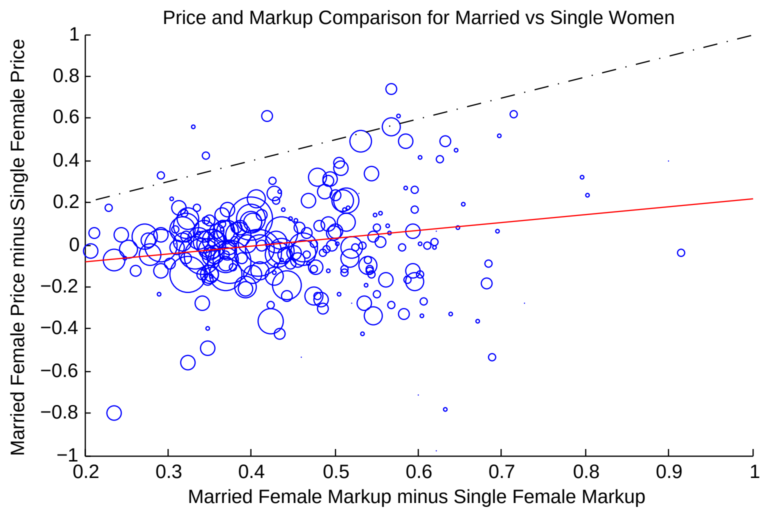


Figure displays 230 vehicle observations. Circle sizes indicate the weight on the observation, $\frac{1}{N_{Aj}+N_{Bj}}$, where group $A$ is married women and group $B$ is single women. Dashed line is the 45° line. The solid line is the best fit line described in table 2.6.

Figure 2.4: Price and Markup Comparison for Married vs Single Men

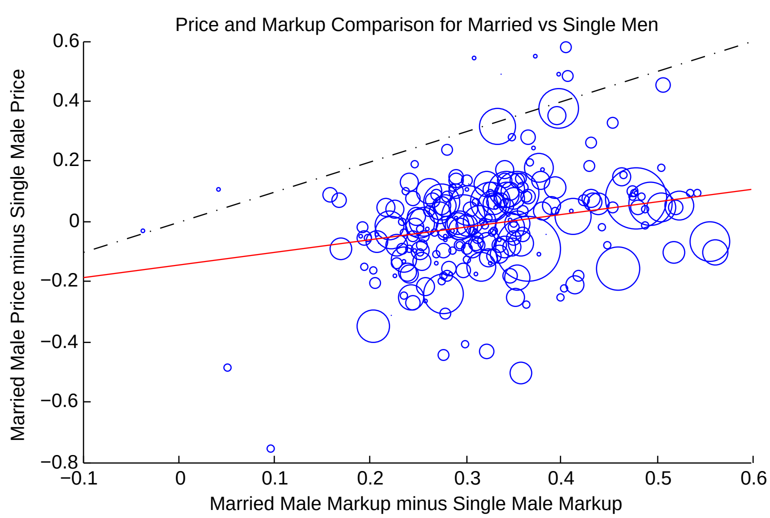


Figure displays 230 vehicle observations. Circle sizes indicate the weight on the observation, $\frac{1}{N_{Aj}+N_{Bj}}$, where group $A$ is married men and group $B$ is single men. Dashed line is the 45° line. The solid line is the best fit line described in table 2.6.

tion based on gender for married people is statistically indistinguishable from the discrimination based on gender for single people. Additionally, if we think that gender is more readily observable to firms than marital status, we might expect firms to price discriminate more effectively based on gender than marital status. Yet the estimated effective price discrimination based on marital status is statistically indistinguishable from the discrimination based on gender. Two things therefore seem to be true: dealers are able to identify consumer demographics and consumers are not particularly adept at circumventing third-degree price discrimination by sending a member of another demographic group to purchase a vehicle for them.

The interpretation of the intercept coefficients, $\gamma_0$, is complicated by its role as both the price difference between two groups when their optimal markups are identical and the intercept in a linear regression measuring effective third-degree price discrimination. However, the fact that the intercepts for the differences across genders are similar (for both married men versus married women and single men versus single women) and the intercepts for the differences across marital statuses are similar (married men versus single men and married women versus single women) is reassuring that the intercepts may be picking up a tendency for single people to pay more than married people and women to pay more than men for vehicles which have the same optimal markup for both groups. These intercepts are not particularly robust to specification changes, however, and there are often few vehicles with similar predicted markups for both groups, as seen in figures 2.1-2.4, so interpreting the intercept as the animus or bargaining differences when markups are identical is typically an out-of-sample prediction. Thus the intercept should be interpreted as animus discrimination or bargaining differences with caution.

## 2.6 Consumer Surplus Implications

By identifying third-degree price discrimination, I have sought to better understand the causes of price differences across demographic groups. But in order for this understanding to be useful to policy-makers, one must also know how third-degree price discrimination affects the consumer surplus of each demographic group. If firms did not price discriminate based on demographic group preferences, then they would charge the same optimal markup to each group. This would increase the price that some groups pay and decrease the price that other groups pay for the same vehicle. Thus the expected sign of the welfare change for each demographic group is uncertain.[37]

[37] Calculating the change in profits to firms from eliminating third-degree price discrimination would be a useful calculation, but would require a measure of the marginal cost of each vehicle,

In order to estimate the effect of eliminating third-degree price discrimination on demographic groups' consumer surplus, I first need to calculate the counterfactual price that firms would charge consumers in the absence of third-degree price discrimination. To do this I make two assumptions. First, I assume that everything about the demographic group's price that I cannot explain, $D_d$, remains constant. Thus, if there is animus in the current prices that I observe, there will be animus in the counterfactual prices with third-degree price discrimination removed. Second, since I observed that firms are only able to extract approximately 40% of the difference in optimal markup between groups, I assume that firms will only change prices by 40% of the difference between the current and counterfactual markups.

With these assumptions and the demand coefficient estimated in section 2.5, the calculation of the counterfactual price, $P_{1dj}$, is relatively straightforward. I use the monopolistic competition markup rule:

$$P_1 = C - \Omega^{-1} Q$$

where $Q$ is the total quantity demanded and

$$\Omega_{ab} = \begin{cases} \frac{\partial Q_b}{\partial P_a} & \text{if } a \text{ and } b \text{ are produced by the same firm} \\ 0 & \text{otherwise} \end{cases}$$

In this case, $Q_j = \sum_d N_d \Pr_{dj}$, so $\frac{\partial Q_b}{\partial P_a} = \sum_d N_d \frac{\partial Pr_{db}}{\partial P_a}$. All of these terms are calculable using the estimated coefficients.

To calculate the counterfactual price keeping all of the $D_d$ constant, I exploit the assumed additive separability of $D_d$ to get:

$$\begin{aligned} P_{1dj} - P_{0dj} &= C_j + \hat{M}_{1j} + D_d - (C_j + \hat{M}_{0dj} + D_d) \\ &= \hat{M}_{1j} - \hat{M}_{0dj} \\ \Rightarrow P_{1dj} &= P_{0dj} + \hat{M}_{1j} - \hat{M}_{0dj} \end{aligned}$$

where $P_{0dj}$ is the observed average price for group $d$ for vehicle $j$; $\hat{M}_{1j}$ is the counterfactual markups for vehicle $j$, and $\hat{M}_{0dj}$ is the third-degree price discrimination markup for group $d$ and vehicle $j$ from section 2.5.[38]

Finally, incorporating the fact that firms can only extract 40% of differences

which is not estimated in this model.

[38] Note that this is actually an approximation. If $D_d$ is the added cost of selling a vehicle to a consumer in group $d$, then the optimal counterfactual price includes a term $\bar{D} = \frac{1}{\partial Q/\partial P} \sum_d \frac{\partial Q_d}{\partial P_1} D_d$ that incorporates the animus cost of selling a particular vehicle. If animus and bargaining differences are small, then the approximation of the counterfactual price is likely close to the truth.

Table 2.7: Estimated Change in Consumer Surplus from Elimination of 3° Price Discrimination

| Demographic Group | $\Delta E[CS]$ | Percent of Group's Expenditures |
|---|---|---|
| Married Men | \$1.7 billion | 5.4% |
| Married Women | \$636 million | 3.7% |
| Single Men | -\$51 million | -0.3% |
| Single Women | -\$1.1 billion | -5.6% |
| Total | \$1.2 billion | 1.5% |

Estimates of the change in consumer surplus from the elimination of third-degree price discrimination follow the standard change in expected consumer surplus rule for random-coefficients logit (Train (2003)). The change in prices from eliminating third-degree price discrimination is calculated assuming that the optimal markup is the only thing that changes (differences in animus and bargaining across groups are assumed to stay constant). Percent of each demographic group's second quarter 2005 expenditures.

in markups, I define the counterfactual price as:

$$P_{1d} = P_{0d} + .4(\hat{M}_1 - \hat{M}_{0d})$$

I then use these price vectors in the calculation of the change in expected consumer surplus for each demographic group:

$$\Delta E[CS] = \sum_{i=1}^{N_d} \int \frac{1}{\tilde{\alpha}_i} \left( \ln \sum_{j=1}^{J} \exp(V_{1ji}(\theta, \nu)) - \ln \sum_{j=1}^{J} \exp(V_{0ji}(\theta, \nu)) \right) f(\nu) d\nu$$

where $\theta$ is the full vector of coefficients estimated in section 2.5, which contains the individual-specific price coefficient, $\tilde{\alpha}_i$ ; $V_{1ji}(\theta, \nu)$ is the non-stochastic component of utility of person $i$ for vehicle $j$ given the counterfactual vehicle prices $P_{1d}$, and $V_{0ji}(\theta, \nu)$ is the non-stochastic component of utility given the original, third-degree price discrimination prices, $P_{0d}$.[39]

The results are presented in Table 2.7. I find that overall, eliminating third-degree price discrimination would have increased consumer surplus by \$1.2 billion

[39]This calculation assumes that the changes in producer profits from the no third-degree price discrimination prices do not change the number or characteristics (other than price) of the vehicles offered. This may not be the case and would be an interesting avenue for future research.

in the second quarter of 2005, or 1.5% of quarterly vehicle expenditures. Married men would have benefited the most, with their consumer surplus increasing by \$1.7 billion, or 5.4% of their quarterly vehicle expenditures. However, single women's consumer surplus would have decreased by \$1.1 billion, or 5.6% of their quarterly vehicle expenditures. Single men's consumer surplus would have been nearly the same, decreasing by only \$51 million (0.3% of expenditures), and married women would have gained \$636 million in consumer surplus (3.7% of expenditures). While the overall change in consumer surplus is not particularly large relative to total expenditures, the interesting result is that single women, the group that pays more for new vehicles controlling for third-degree price discrimination, would be the most hurt by the elimination third-degree price discrimination, and would, to some extent, be transferring consumer surplus to married men, the group that pays less for new vehicles, controlling for third-degree price discrimination. This suggests that third-degree price discrimination may actually help to offset some of the distributional consequences of other factors such as animus or differences in bargaining that lead demographic groups to pay different prices for the same vehicles.

## 2.7 Conclusion

This paper explores the extent to which differences in demographic groups' preferences may lead to third-degree price discrimination. I find that firms do engage in third-degree price discrimination, but that the effective rate of discrimination is between 30 and 45% of "perfect" third-degree price discrimination. Removing the ability to engage in third-degree price discrimination would benefit married men and hurt single women, but increase consumer surplus overall. This suggests that third-degree price discrimination could potentially be countering consumer surplus losses for some groups caused by taste-based discrimination or differences in consumer negotiating preferences.

Finally, this paper develops an approach to identifying the effect of preferences in market outcomes that could be taken to many other settings. Preferences could affect where consumers choose to live, work, and study, and could interact with home sellers, employers, and admissions committees in ways that obscure whether taste-based discrimination is occurring in a market. Understanding the role of preferences in consumers' decision-making would allow policymakers to better target markets where discrimination is leading to adverse distributional outcomes.

# Chapter 3

# Automakers' Short-Run Responses to Changing Gasoline Prices, and Their Imploications for Energy Policy

### Coauthored with Nathan Miller

In this chapter we provide empirical evidence that automobile manufacturers price as if consumers respond to gasoline prices. We estimate a reduced-form regression equation and exploit variation in nearly 300,000 vehicle-week-region observations on manufacturer incentives over 2003-2006. We find that vehicle prices generally fall in the gasoline price; the median price reduction associated with a $1 increase in gasoline prices is $779 for cars and $981 for SUVs. The prices of inefficient vehicles fall more substantially, and the prices of particularly efficient vehicles may rise. Models that ignore these effects underestimate preferences for fuel efficiency and the efficacy of market-based policy instruments.

## 3.1 Introduction

The combustion of gasoline in automobiles poses some of the most pressing policy concerns of the early twenty-first century. This combustion produces carbon dioxide, a greenhouse gas that contributes to global warming. It also limits the flexibility of foreign policy – more than sixty percent of U.S. oil is imported, often from politically unstable regimes. These effects are classic externalities. It is not clear whether, in the absence of intervention, the market is likely to produce efficient outcomes.[1]

We examine the empirical relationship between gasoline prices and the cash incentives offered by automobile manufacturers. The results we obtain suggest that a burdeoning structural literature may understate systematically both consumer preferences for fuel efficiency and the efficacy of market-based policy instruments such as carbon taxes and cap-and-trade regulation. We base our analysis on a formal theoretical model of Nash-Bertand competition with linear downstream demand. The model yields the insight that a change in the gasoline price affects an automobile's equilibrium price through two main channels: its effect on the automobile's fuel cost and its effect on the fuel costs of the automobile's competitors.[2] Consistent with intuition, the model suggests that the net price effect of an adverse gasoline price shock on vehicle prices is negative for most automobiles, but positive for automobiles that are particularly fuel efficient relative to their competitors.

The theoretical model provides a simple reduced-form expression for equilibrium automobile prices that we take to the data. We use a comprehensive set of manufacturer incentives to construct region-time-specific "manufacturer prices" for each of almost 700 automobiles produced by GM, Ford, Chrysler, and Toyota over the period 2003-2006. We combine information on these automobiles' attributes with data on retail gasoline prices to measure fuel costs. We then regress manufacturer prices on fuel costs and competitor fuel costs – identification is strengthened by the dramatic run-up in gasoline prices during the sample period. Overall, we exploit variation among nearly 300,000 automobiles-week-region observations; estimation is feasible even with automobile, time, and region fixed effects.

By way of preview, the results are consistent with a strong and statistically significant manufacturer response to the retail price of gasoline. Manufacturer prices decrease in fuel costs but increase in the fuel costs of competitors. The median net manufacturer price change in response to a hypothetical one dollar increase in gasoline prices is a reduction of \$792 for cars and a reduction of \$981 for SUVs; the median price change for trucks and vans are modest and less statistically significant. Although the fuel cost effect almost always dominates the competitor fuel cost effect, the manufacturer prices of some particularly fuel efficient automobiles do increase (e.g., the 2006 Prius or the 2006 Escape Hy-

[1] Parry, Harrington and Walls (2007) review in detail the externalities of automobile use.

[2] By "fuel cost" we mean the fuel expense associated with driving the automobile. Notably, changes in the gasoline price affect the fuel costs of automobiles differentially – the fuel costs of inefficient automobiles are more responsive to the gasoline prices than the fuel costs of efficient automobiles. One can imagine that the gasoline price may affect equilibrium automobile prices through other channels, perhaps due to an income effect and/or changes in production costs. Our empirical framework allows us to control directly for these alternative channels; we find that their net effect is small.

brid). The manufacturer responses that we estimate are large in magnitude. Back-of-the-envelope calculations suggest that manufacturers often offset a sizable portion of the fuel costs that consumers expect to incur over the lifetimes of their automobiles.

The results provide strong empirical support for the hypothesis that consumer demand for automobile fuel efficiency is elastic with respect to gasoline prices – and therefore that market-based policy instruments such as cap-and-trade regulation and carbon taxes may prove powerful.[3] Importantly, our results suggest that manufacturers promote fuel inefficient automobiles when gasoline prices rise. This subsidy should dampen the shift towards fuel efficient automobiles in the short run, so that many models understate consumer elasticity when manufacturer price adjustments are unobserved in the data.[4] This omitted variable bias is ubiquitous in the recent literature, the bulk of which estimates consumer demand to be relatively inelastic (e.g., Goldberg (1998), Small and Van Dender (2007), Jacobsen (2008), Klier and Linn (2008), Bento et al (2009), Beresteanu and Li (forthcoming),[5] Li, Timmins, and von Haefen (forthcoming)). Even Gramlich (2009), which controls for endogenous automobile characteristics and produces more elastic estimates, may best be interpreted as providing a lower bound to consumer elasticity.

Finally, we note that our methodology and conclusions complement those presented in contemporaneous work by Busse, Knittel and Zettelmeyer (2009).[6] Whereas we examine the response of manufacturers incentives to gasoline prices, Busse, Knittel, and Zettelmeyer examine the response of transaction prices paid by consumers. The fact that these responses are quantitatively similar yields the insight, unavailable from either paper alone, that automobile manufacturers may accept much of the risk related to gasoline price fluctuations from dealerships.[7] Since the demand for fuel efficient automobiles appears to be less sensitive to gasoline prices, investments in the development of these automobiles can be interpreted as a hedge against gasoline price shocks.

[3]Perhaps most intriguingly, the short-run manufacturer price changes that we estimate should magnify the long-run incentives of manufacturers to develop and market fuel efficient automobiles. We speculate that long-run incentives may be of first order importance to the efficacy of market-based policy instruments. To our knowledge, the literature has yet to convincingly examine this difficult supply-side consideration.

[4]We formalize this argument for the specific case of logit demand in an appendix.

[5]Beresteanu and Li do instrument for operating cost in their analysis, and they find that the instrumented coefficient is substantially larger in absolute value than the OLS result. However, their instruments, average operating costs in other census regions and divisions, will likely still be correlated with national vehicle incentives resulting from national changes in gas prices.

[6]Busse, Knittel, and Zettelmeyer examine a large sample of consumer transactions over the period 1999-2008 and conclude that higher gasoline prices are associated with shifts in demand towards fuel efficient vehicles.

[7]See appendix 5.2.4 for a discussion of the two sets of results.

The paper proceeds as follows. We lay out the empirical model in Section 3.2, including the underlying theoretical framework and the empirical implementation. We describe the data and regression variables in Section 3.3, present the main regression results in Section 3.4, and discuss three extensions in Section 3.5. We conclude in Section 3.6.

## 3.2 The Empirical Model

### 3.2.1 Theoretical framework

We derive our estimation equation from a model of Bertrand-Nash competition between automobile manufacturers that a face a linear demand schedule. We take as given that there are $F$ automobile manufacturers and $J_t$ vehicles. Each manufacturer produces some subset $\Im_f$ of the vehicles and prices to maximize short-run profits:

$$\pi_{ft} = \sum_{j \in \Im_f} [(p_{jt} - c_{jt}) * q_{jt} - f_{jt}] \tag{3.1}$$

where for each vehicle $j$ and period $t$, the terms $p_{jt}$, $c_{jt}$, and $q_{jt}$ are the manufacturer price, the marginal cost, and the quantity sold respectively. We denote the fixed cost of production as $f_{jt}$. We assume that consumer demand depends on manufacturer prices, expected lifetime fuel costs, and certain exogenous demand shifters that include vehicle attributes, maintenance costs, and other factors:

$$q(p_{jt}) = \sum_{k=1}^{J_t} \alpha_{jk}(p_{kt} + x_{kt}) + \mu_{jt}, \tag{3.2}$$

where the $\alpha_{jk}$ is a demand parameter, $x_{kt}$ captures fuel costs, and $\mu_{jt}$ captures the demand shifters. We consider the case in which demand is well defined ($\partial q_{jt}/\partial p_{jt} = \alpha_{jt} < 0$) and vehicles are substitutes ($\partial q_{jt}/\partial p_{kt} = \alpha_{jk} \geq 0$ for $k \neq j$).[8]

The equilibrium manufacturer prices in each period are then characterized by $J_t$ first-order conditions. We solve these first-order equations to obtain equilibrium manufacturer prices as functions of the exogenous factors.[9] The resulting

[8]We assume that marginal costs are constant in quantity but responsive to certain exogenous cost shifters. Also, we abstract from the manufacturers' selections of vehicle attributes and fleet composition, as well as any entry and/or exit, which we deem to be more important in longer-run analysis.

[9]The solution technique is simple. Turning to vector notation, one can rearrange the first-order conditions such that $Ap = b$, where $A$ is a $J_t \times J_t$ matrix of demand parameters, $p$ is a $J_t \times 1$ vector of manufacturer prices, and $b$ is a $J_t \times 1$ vector of "solutions" that incorporate the fuel costs, marginal costs, and demand shifters. Provided that the matrix $A$ is nonsingular, Cramer's

manufacturer "price rule" is a linear function of the fuel costs, marginal costs, and demand shifters:

$$\begin{aligned} p^*_{jt} = \phi^1_{jt} x_{jt} + \sum_{k \notin \Im_f} \phi^2_{jkt} x_{kt} + \sum_{l \in \Im_f,\ l \neq j} \phi^3_{jlt} x_{lt} \\ + \phi^4_{jt} c_{jt} + \phi^5_{jt} \mu_{jt} + \sum_{k \notin \Im_f} \left( \phi^6_{jkt} c_{kt} + \phi^7_{jkt} \mu_{kt} \right) + \sum_{l \in \Im_f,\ l \neq j} \left( \phi^8_{jlt} c_{lt} + \phi^9_{jlt} \mu_{lt} \right). \end{aligned} \quad (3.3)$$

The reduced-form coefficients $\phi^1, \phi^2, \ldots, \phi^9$ are nonlinear functions of all the demand parameters. The price rule makes it clear that the equilibrium price of a vehicle depends on its characteristics (i.e, its fuel cost, marginal cost, and demand shifter), the characteristics of vehicles produced by competitors, and the characteristics of other vehicles produced by the same manufacturer. For the time being, we collapse the second line of the price rule into a vehicle-time-specific constant, which we denote $\gamma_{jt}$.

The sheer number of terms in Equation 3.3 makes direct estimation infeasible. With only $J_t$ observations per period, one cannot hope to identify the $J_t^2$ fuel cost coefficients, let alone the vehicle-time-specific constant. We move toward the empirical implementation by re-expressing the price rule in terms of weighted averages:

$$p^*_{jt} = \phi^1_{jt} x_{jt} + \phi^2_{jt} \sum_{k \notin \Im_f} \omega^2_{jkt} x_{kt} + \phi^3_{jt} \sum_{l \in \Im_f,\ l \neq j} \omega^3_{jlt} x_{lt} + \gamma_{jt}, \quad (3.4)$$

where the weights $\omega^2_{jkt}$ and $\omega^3_{jlt}$ both sum to one in each period; closer competitors receive greater weight.[10] Thus, the equilibrium price depends on its fuel cost, the weighted average fuel cost of vehicles produced by competitors, and the weighted average fuel cost of vehicles produced by the same manufacturer. The theory suggests that $\phi^1_{jt} < 0$ and $\phi^2_{jt} > 0$.[11]

The reduced-form imbeds the intuition that manufacturer prices can increase or decrease in response to adverse gasoline price shocks. Assume for the moment that the gasoline price does not affect the cost or demand shifters, and therefore

Rule applies and there exists a unique Nash equilibrium in which the equilibrium manufacturer prices are linear functions of all the fuel costs, marginal costs, and demand shifters.

[10] The weights have analytical solutions given by $\omega^i_{jkt} = \phi^i_{jkt}/\phi^i_{jt}$, and the coefficients $\phi^2_{jt}$ and $\phi^3_{jt}$ are the sums of the $\phi^2_{jkt}$ and $\phi^3_{jkt}$ coefficients, respectively. Mathematically, $\phi^i_{jt} = \sum \phi^i_{jkt}$.

[11] We derive this relationship in the working paper. Using a mild regularity condition, we show that 1) the equilibrium price of a vehicle decreases in its fuel costs and increases in the fuel costs of its competitors, 2) the equilibrium price of a vehicle is more responsive to changes in its fuel cost than identical changes in the fuel costs of its competitors, and 3) the relationship between the equilibrium price of a vehicle and the fuel costs of other vehicles produced by the same manufacturer is ambiguous (though if demand is symmetric these fuel costs have no effect).

does not affect the vehicle-time-specific constant (we relax this assumption in an extension). Denoting the gasoline price at time $t$ as $\text{gp}_t$, the effect of the gasoline price shock on the manufacturer price is:

$$\frac{\partial p^*_{jt}}{\partial \text{gp}_t} = \phi^1_j \frac{\partial x_{jt}}{\partial \text{gp}_t} + \phi^2_{jt} \sum_{k \notin \Im_j} \omega^2_{jkt} \frac{\partial x_{kt}}{\partial \text{gp}_t} + \phi^3_{jt} \sum_{l \in \Im_j,\ l \neq j} \omega^3_{jlt} \frac{\partial x_{lt}}{\partial \text{gp}_t}, \tag{3.5}$$

where fuel costs increase unequivocally in the gasoline price (i.e., $\partial x_{jt}/\partial \text{gp}_t > 0 \ \forall \ j$). The first term shows that manufacturers partially offset an increase in the fuel cost with a reduction in the vehicle's price. This reduction is greater for vehicles whose fuel costs are sensitive to the gasoline price (i.e., for fuel-inefficient vehicles). The second and third terms show that an increase in the fuel costs of other vehicles can raise demand through consumer substitution. Although the first effect tends to dominate, prices can increase for vehicles that are sufficiently more fuel efficient than their competitors.

### 3.2.2 Empirical implementation

The empirical implementation requires that we specify the fuel costs ($x_{jt}$), the weights ($\omega^2_{jkt}$ and $\omega^3_{jkt}$), and the vehicle-time-specific constants ($\gamma_{jt}$). We discuss each in turn.

We proxy expected lifetime fuel costs as a function of vehicle fuel efficiency and gasoline prices, following Goldberg (1998), Bento et al (2005) and Jacobsen (2007). The specific functional form is:

$$x_{jt} = \tau * \frac{\text{gp}_t}{\text{mpg}_j},$$

where $\text{mpg}_j$ denotes miles-per-gallon and $\tau$ is a discount factor that nests any form of multiplicative discounting; one specific possibility is $\tau = 1/(1-\delta)$, where $\delta$ is the "per-mile discount rate."[12] The fuel cost proxy is precise if consumers perceive the gasoline price to follow a random walk because, in that case, the current gasoline price is a sufficient statistic for expectations over future gasoline prices. As we discuss below, we fail to reject the null hypothesis that gasoline prices actually follow a random walk, but also provide some evidence that consumers consider both historical gasoline prices and futures prices when forming expectations.

To construct the weights, we assume that the intensity of competition between any two vehicles decreases in the Euclidean distance between their at-

[12]It may help intuition to note that the ratio of the gasoline price to miles-per-gallon is simply the gasoline expense associated with a single mile of travel.

tributes. To that end, we take a set of $M$ vehicle attributes, denoted $z_{jm}$ for $m = 1, \ldots, M$, and standardize each to have a variance of one. We then sum the squared differences between each attribute to calculate the effective "distance" in attribute space. We form initial weights as follows:

$$\omega^*_{jk} = \frac{1}{\sum_{m=1}^{M} (z_{jm} - z_{km})^2}.$$

To finish, we set the initial weights to zero for vehicles of different types and then normalize the weights to sum to one for each vehicle-period. We perform this weighting procedure separately for vehicles produced by the same manufacturer and vehicles produced by competitors; the result is a set of empirical weights that we denote $\widetilde{\omega}^2_{jkt}$ and $\widetilde{\omega}^3_{jkt}$.[13]

The vehicle-time-specific constants represent the net price effects of the demand and marginal cost shifters. We specify these effects using vehicle fixed effects, time fixed effects, and controls for the number of weeks that each vehicle has been on the market. Denoting the number of weeks a vehicle has been on the market as $\lambda_{jt}$ and the weighted average number of weeks since the vehicles in the set $A$ have been on the market as $\bar{\lambda}_{A,t}$, the specification takes the form:

$$\gamma_{jt} = \delta_t + \kappa_j + f(\lambda_{jt}) + g(\bar{\lambda}_{k \notin \Im_j, t}) + h(\bar{\lambda}_{k \in \Im_j,\ k \neq j, t}) + \epsilon_{jt}$$

where $\delta_t$ and $\kappa_j$ are time and vehicle fixed effects, respectively, and functions $f$, $g$, and $h$ capture the net price effects of learning-by doing and predictable demand changes over the model-year.[14] In the main results, we specify these functions as third-order polynomials; the results are robust to the use of higher-order or lower-order polynomials. The error term $\epsilon_{jt}$ captures vehicle-time-specific cost and demand shocks.

Two final adjustments produce the main regression equation that we take to the data. First, we incorporate regional variation in manufacturer prices and gasoline prices and add a corresponding set of region fixed effects.[15] Second, we

[13]Thus, the weighting scheme is based on the inverse Euclidean distance between vehicle attributes among vehicles of the same type. There are four vehicle types in the data: cars, SUVs, trucks and vans. We use the following set of vehicle attributes in the initial weights: manufacturer suggested retail price (MSRP), miles-per-gallon, wheel base, horsepower, passenger capacity, and dummies for the vehicle type and segment. Although the initial weights are constant across time for any vehicle pair, the final weights may vary due to changes in the set of vehicles available on the market. The results are robust to the use of various alternative weighing schemes based on straight averages over all competitors, over competitors of the same type, and over competitors of the same segment; we provide details in an appendix.

[14]Copeland, Dunn and Hall (2005) document that vehicle prices fall approximately nine percent over the course of the model-year.

[15]Adding regional variation in prices does not complicate the weight calculations because

impose a homogeneity constraint that reduces the total number of parameters to be estimated; the constraint eliminates vehicle-time variation in the coefficients, so that $\phi^i_{jt} = \phi^i \; \forall \; j, t$. (In supplementary regressions, we permit the coefficients to vary across manufacturers and vehicle types.) The regression equation is:

$$\begin{aligned} p_{jtr} &= \beta^1 \frac{\mathrm{gp}_{tr}}{\mathrm{mpg}_j} + \beta^2 \sum_{k \notin \Im_j} \widetilde{\omega}^2_{jkt} \frac{\mathrm{gp}_{tr}}{\mathrm{mpg}_k} + \beta^3 \sum_{l \in \Im_j,\; l \neq j} \widetilde{\omega}^3_{jlt} \frac{\mathrm{gp}_{tr}}{\mathrm{mpg}_l} \\ &+ f(\lambda_{jt}) + g(\bar{\lambda}_{k \notin \Im_j, t}) + h(\bar{\lambda}_{k \in \Im_j,\; k \neq j, t}) + \delta_t + \kappa_j + \eta_r + \epsilon_{jt}, \qquad (3.6) \end{aligned}$$

where the fuel cost coefficients incorporate the discount factor, i.e., $\beta^i = \tau \phi^i$ for $i = 1, 2, 3$; for reasonable discount factors, these coefficients could be much larger than one in magnitude. Thus, we estimate the average response of a vehicle's price to changes in its fuel costs, changes in the weighted average fuel cost among vehicles produced by competitors, and changes in the weighted average fuel cost among other vehicles produced by the same manufacturer.

We estimate Equation 3.6 using ordinary least squares. We are able to identify the fuel cost coefficients in the presence of time, vehicle, and region fixed effects precisely because changes in the gasoline price affect manufacturer prices differentially across vehicles.[16] We argue that manufacturers price as if consumers respond to gasoline prices if the fuel cost coefficient is negative (i.e., $\beta^1 < 0$) and the competitor fuel cost coefficient is positive (i.e., $\beta^2 > 0$). The theoretical results suggest that the fuel cost coefficient should be larger in magnitude than the competitor fuel cost coefficient (i.e., $|\beta^1| > |\beta^2|$); more generally, the relative magnitude of these coefficients determines the extent to which average manufacturer prices fall in response to an adverse gasoline shock. We cluster the standard errors at the vehicle level, which accounts for arbitrary correlation patterns in the error terms.[17]

---

there is no regional variation in the vehicles available to consumers.

[16]The fixed effects help mitigate various endogeneity concerns. Consider two examples: First, the demand for new automobiles in the U.S. likely has a small effect on the global price of oil. Time fixed effects account for the overall effect, however, so only changes in the distribution of demand (e.g., greater demand for efficient vehicles) could create bias. Fuel costs are most obvious source of such relative demand changes, but their effect is unproblematic because fuel costs are included in the model. Second, manufacturers adjust the characteristics of their models in response to changes in the gasoline price. However, the inclusion of vehicle fixed effects restricts identification to changes in the gasoline price that occur within the model-year; and model characteristics are fixed within the model-year.

[17]The results are robust to the use of brand-level or segment-level clusters; brands and segments are finer gradations of the manufacturers and types, respectively. There are 21 brands and 15 vehicle segments in the data. Examples of brands include Chevrolet (GM), Dodge (Chrysler), Mercury (Ford), and Lexus (Toyota). Examples of segments include compact cars, luxury SUVs, and large pick-ups.

## 3.3 Data Sources and Regression Variables

### 3.3.1 Data sources

Our primary source of data is Autodata Solutions, a marketing research company that maintains a comprehensive database of manufacturer incentive programs. We have access to the programs offered by Toyota and the "Big Three" U.S. manufacturers – GM, Ford, and Chrysler – over the period 2003-2006.[18] There are just over 190,000 cash incentive-vehicle pairs in the data. Each lasts a fixed period of time, and provides cash to consumers ("consumer-cash") or dealerships ("dealer-cash") at the time of purchase. The incentive programs may be national, regional, or local in their geographic scope; we restrict our attention to the national and regional programs.[19] Thus, we are able to track how manufacturer incentives change over time and across regions for each vehicle in the data.

By "vehicle," we mean a particular model in a particular model-year. For example, the 2003 Ford Taurus is one vehicle in the data, and we consider it as distinct from the 2004 Ford Taurus. Overall, there are 681 vehicles in the data – 293 cars, 202 SUVs, 105 trucks, and 81 vans. The data have information on the attributes of each, including MSRP, miles-per-gallon, horsepower, wheel base, and passenger capacity.[20] We impute the period over which each vehicle is available to consumers as beginning with the start date of production, as listed in Ward's Automotive Yearbook, and ending after the last incentive program for that vehicle expires.[21] For each vehicle, we construct observations over the relevant period at the week-region level.

We combine the Autodata Solutions data with information from the Energy Information Agency (EIA) on weekly retail gasoline prices in each of five distinct geographic regions. The EIA surveys retail gasoline outlets every Monday for the per gallon pump price paid by consumers (inclusive of all taxes).[22] In addition

[18] The German manufacturer Daimler owned Chrysler over this period. We exclude Mercedes-Benz from this analysis since it is traditionally associated with Daimler rather than Chrysler.

[19] Because the gas price data from the Energy Information Agency is at the regional level, we consider an incentive to be regional if it is available across an entire Energy Information Agency region. We exclude incentives that are available in only a single city or state.

[20] Attributes sometimes differ for a given vehicle due to the existence of different option packages, also known as "trim." When more than one set of attributes exists for a vehicle, we use the attributes corresponding to the trim with the lowest MSRP.

[21] The start date of production is unavailable for some vehicles. For those cases, we set the start date at August 1 of the previous year. For example, we set the start date of the 2006 Civic Hybrid to be August 1, 2005. We impose a maximum period length of 24 months. In robustness checks, we used an 18 month maximum; the different period lengths do not affect the results.

[22] The survey methodology is detailed online at the EIA webpage. The regions include the

to the regional measures, the EIA calculates an average national price. Figure 3.1 plots these retail gasoline prices over 2003-2006 (in real 2006 dollars). A run-up in gasoline prices over the sample period is apparent. For example, the mean national gasoline price is 1.75 dollars-per-gallon in 2003 and 2.57 dollars-per-gallon in 2006. The sharp upward spike around September 2005 is due to Hurricane Katrina, which temporarily eliminated more than 25 percent of US crude oil production and 10-15 percent of the US refinery capacity (EIA 2006). Although gasoline prices tend to move together across regions, we are able to exploit limited geographic variation to strengthen identification.

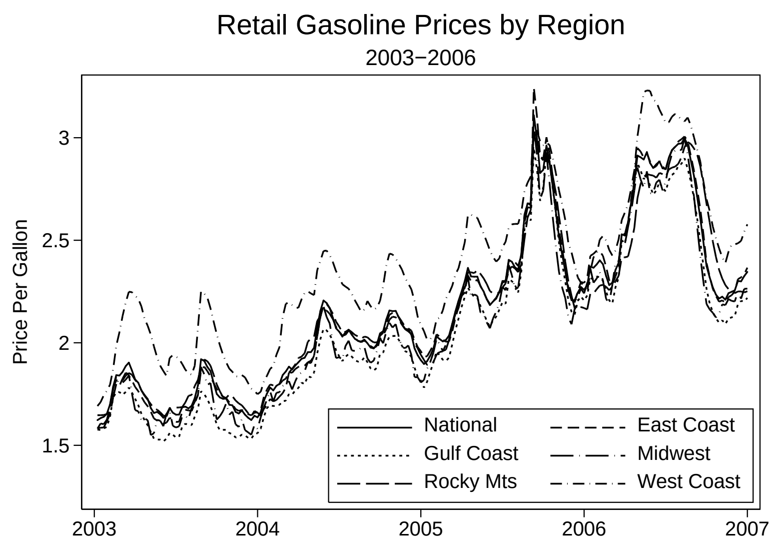


Figure 3.1: The weekly retail price of gasoline by region over 2003-2006, in real 2006 dollars.

We purge the gasoline prices of seasonality prior to their use in the analysis. Since automobile manufacturers adjust their prices cyclically over vehicle model-years (e.g., Copeland, Hall, and Dunn 2005), the presence of seasonality in gasoline prices is potentially confounding. Further, the use of time fixed effects alone may be insufficient in dealing with seasonality because gasoline prices affect the fuel costs of each vehicle differentially (e.g., Equation 3.6). We employ the X-12-ARIMA program, which is state-of-the-art and commonly employed elsewhere, for example by the Bureau of Labor Statistics to deseasonalize inputs to the consumer price index.[23] Figure 3.2 plots the resulting deseasonalized national

East Coast, the Gulf Coast, the Midwest, the Rocky Mountains, and the West Coast.

[23]We use data on gasoline prices over 1993-2008 to improve the estimation of seasonal factors, and adjust each national and regional time-series independently. We specify multiplicative

gasoline prices together with the seasonal adjustments. As shown, the program adjusts the gasoline price downward during the summer months and upwards during the winter months. The magnitude of the adjustments increases with gasoline prices.

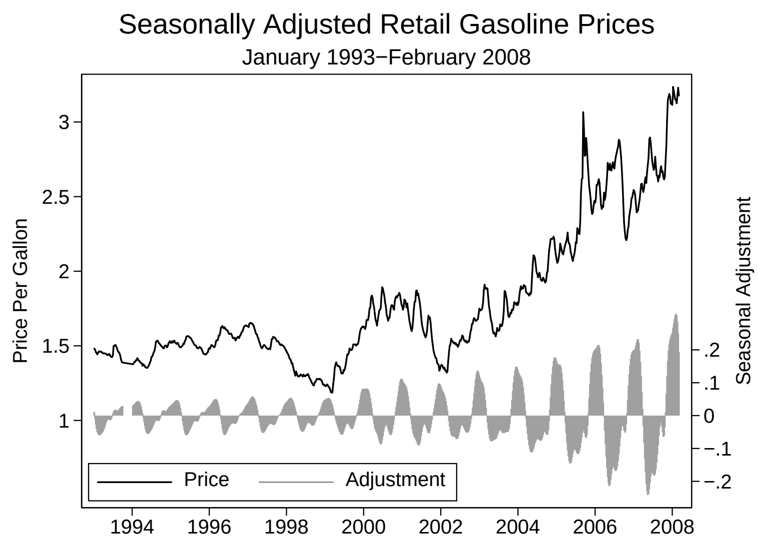


Figure 3.2: Seasonally adjusted retail gasoline prices at the national level over 1993-2008, in real 2006 dollars. Seasonal adjustments are calculated with the X-12-ARIMA program.

## 3.3.2 Regression variables

The two critical variables that enable regression analysis are manufacturer price and fuel cost. We discuss each in turn. To start, we measure the manufacturer price of each vehicle as MSRP minus the mean incentive available for the given week and region. We also show results in which the variable includes only regional incentives and only national incentives, respectively. From an econometric standpoint, the MSRP portion of the variable is irrelevant for estimation because the vehicle fixed effects are collinear (MSRP is constant for all observations on a given vehicle). It is the variation in manufacturer incentives across vehicles, weeks, and regions that identifies the regression coefficients.

decomposition, which allows the effect of seasonality to increase with the magnitude of the trend-cycle. The results are robust to log-additive and additive decompositions. For more details on the X-12-ARIMA, see Makridakis, Wheelwright and Hyndman (1998) and Miller and Williams (2004).

At least two important caveats apply to our manufacturer price variable. First, the variable does not capture any information about final transaction prices, which are negotiated between the consumers and the dealerships. Changes in negotiating behavior could dampen or accentuate the effect we estimate between gasoline prices and manufacturer prices. Second, although we observe the incentive programs, we do not observe the actual incentives selected. In some circumstances, it is possible that consumers may stack multiple incentives or choose between different incentives. To the extent that manufacturers are more lenient in allowing consumers to stack incentives when gasoline prices are high, our regression estimates are conservative relative to the true manufacturer response.[24]

We measure the fuel costs of each vehicle as the gasoline price divided by the miles-per-gallon of the vehicle. As discussed above, this has the interpretation of being the gasoline expense associated with a single mile of travel. Since the gasoline price varies at the week and region levels and miles-per-gallon varies at the vehicle level, fuel costs vary at the vehicle-week-region level. In an extension, we construct alternative fuel costs based on 1) the mean of the gasoline price over the previous four weeks and 2) the price of one-month futures contract for retail gasoline. The futures data are derived from the New York Mercantile Exchange (NYMEX) and are publicly available from the EIA. The alternative variables permit tests for whether consumers are backward-looking and forward-looking, respectively.[25]

Table 3.1 provides means and standard deviations for the manufacturer price and the gasoline price variables, as well as for five vehicle attributes used in the weighting scheme – MSRP, miles-per-gallon, horsepower, wheel base, and passenger capacity. The statistics are calculated from the 299,855 vehicle-region-week observations formed from the 681 vehicles, 208 weeks, and five regions in the data. As shown, the mean manufacturer price is 30.34 (in thousands). The mean fuel cost is 0.11, so that gasoline expenses average roughly eleven cents per mile. The means of MSRP, miles-per-gallon, horsepower, wheel base, and passenger capacity are 30.78, 21.56, 224.12, 115.19, and 4.91, respectively. As the standard deviations suggest, vehicles differ substantially in these observed characteristics; differences exist both within and across vehicle types. Of course, vehicles also differ along unobserved dimensions. We use vehicle fixed effects to control for all

[24] To check the sensitivity of the results, we construct a number of alternative variables that measure manufacturer prices: 1) MSRP minus the maximum incentive, 2) MSRP minus the mean consumer-cash incentive, 3) MSRP minus the mean dealer-cash incentive, and 4) MSRP minus the mean publicly available incentive. None of these alternative dependent variables substantially change the results.

[25] We use one-month futures contracts for reformulated regular gasoline at the New York harbor. In order to ensure that the regression coefficients are easily comparable, we normalize the futures price to have the same global mean over the period as the national retail gasoline price.

Table 3.1: Summary Statistics

| Variables | Definition | Mean | St. Dev. |
|---|---|---|---|
| Manufacturer price | $\mathrm{MSRP}_j - \overline{\mathrm{INC}}_{jrt}$ | 30.344 | 16.262 |
| Fuel cost | $\mathrm{gp}_{rt}/\mathrm{mpg}_j$ | 0.108 | 0.034 |
| MSRP | $\mathrm{MSRP}_j$ | 30.782 | 16.299 |
| Miles-per-gallon | $\mathrm{mpg}_j$ | 21.555 | 5.964 |
| Horsepower | | 224.123 | 71.451 |
| Wheel base | | 115.193 | 12.168 |
| Passenger capacity | | 4.911 | 1.633 |

Means and standard deviations based on 299,855 vehicle-region-week observations over the period 2003-2006. The manufacturer price is defined as MSRP minus the mean regional and national incentives (in thousands). The fuel cost is the gasoline price divided by miles-per-gallon, and captures the gasoline expense per mile. The manufacturer price, the fuel cost, and MSRP (in thousands) are in real 2006 dollars; wheel base is measured in inches.

of this heterogeneity – observed and unobserved – in our regression results.[26]

## 3.4 Empirical Results

### 3.4.1 Regression with the homogeneity constraint

We regress manufacturer prices on fuel costs, as specified in Equation 3.6. To start, we impose the full homogeneity constraint that all vehicles share the same fuel cost coefficients. The regression coefficients estimate the average response of manufacturer prices to fuel costs. Table 3.2 presents the results. In Column 1, the dependent variable is MSRP minus the mean of the regional and national incentives. In Columns 2 and 3, the dependent variables are MSRP minus the mean regional incentive and MSRP minus the mean national incentive, respectively. Although the first column may provide more meaningful coefficients, we believe that the second and third columns are interesting insofar as they examine whether manufacturers respond at the regional and national levels, respectively.

As shown, the fuel cost coefficients of -55.40, -56.96, and -63.75 are precisely

[26] The working paper provides summary statistics separately for each vehicle type. These statistics are consistent with the generalization that cars are smaller, more fuel efficient, and less powerful than SUVs, trucks, and vans.

Table 3.2: Manufacturer Prices and Fuel Costs

| | Incentive level: | | |
|---|---|---|---|
| | Regional+ National | Regional Only | National Only |
| Variables | (1) | (2) | (3) |
| Fuel cost | -55.40*** | -56.96*** | -63.75*** |
| | (7.73) | (7.86) | (8.77) |
| Average competitor fuel cost | 50.76*** | 50.16*** | 50.09*** |
| | (7.15) | (7.39) | (8.12) |
| Average same-firm fuel cost | 1.15 | 2.62 | 1.31 |
| | (2.29) | (1.78) | (2.30) |
| $R^2$ | 0.5260 | 0.6763 | 0.5289 |
| # of observations | 299,855 | 299,855 | 59,971 |
| # of vehicles | 681 | 681 | 681 |

Results from OLS regressions. The dependent variable is the manufacturer price, i.e., MSRP minus the mean regional and/or national incentives (in thousands). The units of observation in Columns 1 and 2 are at the vehicle-week-region level. The units of observation in Column 3 are at the vehicle-week level. All regressions include vehicle and time fixed effects, and Columns 1 and 2 include region fixed effects. The regressions also include third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

estimated and capture the intuition that manufacturers adjust their prices to offset changes in fuel costs. The competitor fuel cost coefficients of 50.76, 50.16, and 50.09 are also precisely estimated and capture the intuition that increases in competitors' fuel costs raise demand due to consumer substitution. The magnitudes of the fuel cost coefficients exceeds the magnitudes of the competitor fuel cost coefficients, which is suggestive that the first effect dominates for most vehicles. The same-firm fuel cost coefficients are close to zero, consistent with roughly symmetric demand. Finally, a comparison of coefficients across columns suggests that manufacturers adjust their prices similarly at the regional and national levels in response to changes in fuel costs.[27]

We explore the effect of retail gasoline prices on manufacturer prices in Figure 3.3. The gasoline price enters through the fuel costs, average competitor fuel costs, and average same-firm fuel costs. We calculate the net effect of a one dollar increase in the gasoline price for each vehicle-week-region observation:

$$\frac{\partial p_{jrt}}{\partial \mathrm{gp}_{rt}} = \frac{\widehat{\beta}_1}{\mathrm{mpg}_j} + \widehat{\beta}_2 \sum_{k \neq j} \frac{\widetilde{\omega}^2_{jkt}}{\mathrm{mpg}_k} + \widehat{\beta}_3 \sum_{k \neq j} \frac{\widetilde{\omega}^2_{jkt}}{\mathrm{mpg}_k}.$$

We plot these effects (in thousands) on the vertical axis against vehicle miles-per-gallon on the horizontal axis. We focus on the first dependent variable, i.e., MSRP minus the mean regional and national incentive.[28] The median manufacturer response to a one dollar increase in the gasoline price is a price reduction of \$171. The responses range from a price reduction of \$1,506 for the 2005 GM Montana SV6 to a price increase of \$998 for the 2006 Toyota Prius. And, although manufacturer prices fall for 83 percent of the vehicles, the prices of fuel efficient vehicles fall less and the prices of extremely fuel efficient vehicles actually increase.[29]

[27]Fuel costs explain about ten percent of the variance in manufacturer prices, based on comparisons to regressions that exclude fuel costs (not shown). The results do not seem to be driven by outliers; the coefficients are similar when we exclude the extremely fuel efficient or fuel inefficient vehicles from the sample. In an appendix, we provide tests for non-linearities, regressions that use alternative weighting schemes, sub-sample regressions for each region and for the cities of San Francisco and Houston, and sub-sample regressions for each model-year; the results are robust to each change.

[28]We plot each vehicle only once because the derivatives do not vary substantively over time or regions. Indeed, the only variation within vehicles is due to changes in the set of other vehicles available.

[29]A striking characteristic of the results is that the relationship between the price effect and vehicle miles-per-gallon is concave. The concavity is consistent with the intuition that gasoline prices are more relevant to fuel inefficient cars (e.g., moving from 14 to 15 miles-per-gallon has a larger absolute effect on fuel costs than moving from 30 to 31 miles-per-gallon). One might worry that the construction of the fuel cost variable imposes concavity artificially. In an appendix, we estimate more flexible regressions and demonstrate that the concavity is data

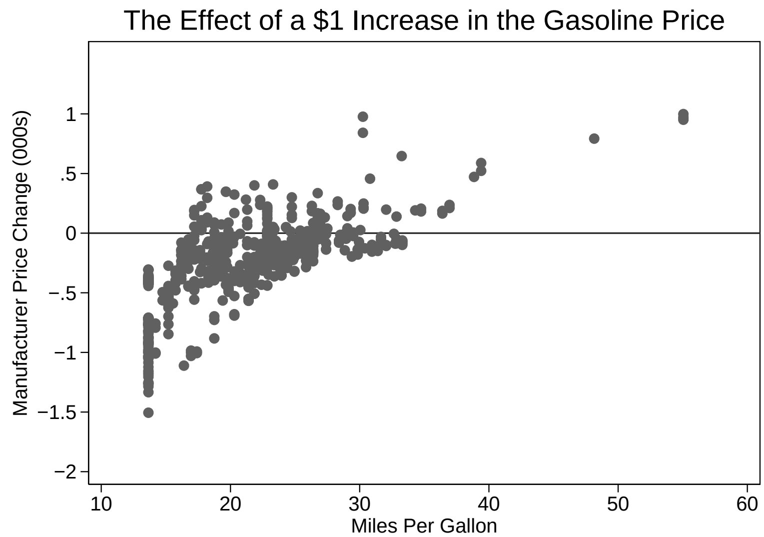


Figure 3.3: The estimated effects of a one dollar increase in the retail gasoline price on the manufacturer price, based on the regression results in Column 1 of Table 3.2. Each point represents the price effect for a single vehicle. See text for details.

### 3.4.2 Regression relaxing the homogeneity constraint

We use sub-sample regressions to relax the homogeneity constraint that all vehicles share the same fuel cost coefficients. In particular, we regress manufacturer prices on the fuel cost variables for each combination of vehicle type (cars, SUVs, trucks, and vans) and manufacturer (GM, Ford, Chrysler, and Toyota). The dependent variable in each case is MSRP minus the mean of the regional and national incentives. These regressions should be more accurate if, for example, some manufacturers have more elastic demand than others and/or car purchasers differ systematically from SUV purchasers.[30] The regression coefficients appear in Table 3.3. As expected, the fuel cost coefficients tend to be negative and the competitor fuel costs coefficients tend to be positive. We use figures to explore the results in detail.

driven. We take this as substantial support for the main specification.

[30]One might additionally suspect that responses to fuel costs changes over time. To test for such heterogeneity, we split the observations to form one sub-sample over the period 2003-2004 and another over the period 2005-2006; the results from each sub-sample are quite close. Similarly, we divide the sample between the 2003-2004 model-years and the 2005-2006 model-years without substantially changing the results. We conclude that the effects of any time-related heterogeneity are small.

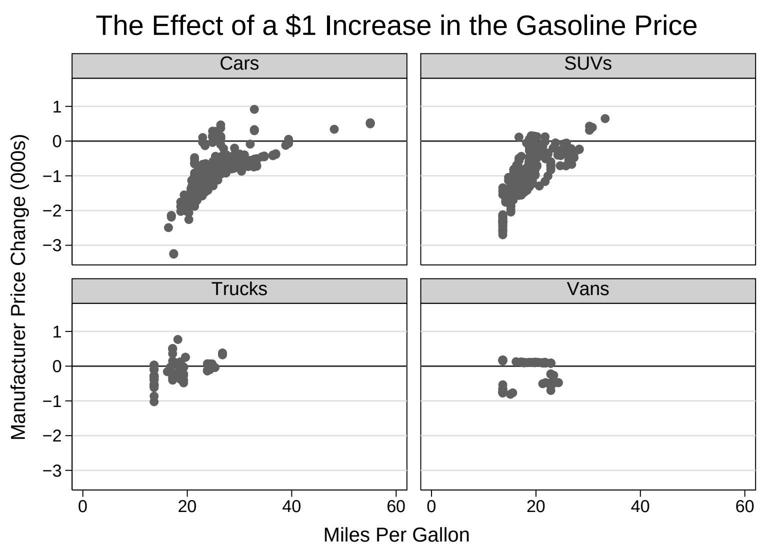


Figure 3.4: The estimated effects of a one dollar increase in the retail gasoline price on the manufacturer price, based on the regression results of Table 3.3. Each point represents the price effect for a single vehicle. See text for details.

Table 3.3: Manufacturer Prices by Vehicle Type and Manufacturer

| Vehicle Type: | Cars | | | | SUVs | | | |
|---|---|---|---|---|---|---|---|---|
| Manufacturer: | GM | Ford | Chrysler | Toyota | GM | Ford | Chrysler | Toyota |
| Fuel cost | -97.52*** | -146.80** | -152.09*** | -77.13*** | -75.98*** | -72.10*** | 45.87* | -62.11*** |
| | (17.85) | (71.68) | (30.81) | (21.55) | (23.26) | (23.50) | (24.92) | (17.82) |
| Average competitor fuel cost | 85.78*** | 80.61 | 159.99*** | 46.38*** | 64.08*** | 66.75*** | -29.56 | 44.25*** |
| | (18.25) | (57.25) | (41.91) | (18.52) | (21.67) | (23.79) | (19.01) | (16.41) |
| Average same-firm fuel cost | -6.47 | 41.10 | -2.56 | 6.12 | 8.19 | -1.51 | -17.88* | 1.73 |
| | (8.79) | (29.64) | (11.33) | (4.48) | (9.02) | (6.28) | (10.42) | (4.27) |
| $R^2$ | 0.6173 | 0.5254 | 0.5294 | 0.7282 | 0.7861 | 0.6758 | 0.7126 | 0.8352 |
| # of vehicles | 101 | 92 | 34 | 66 | 94 | 50 | 24 | 34 |
| Vehicle Type: | Trucks | | | | Vans | | | |
| Manufacturer: | GM | Ford | Chrysler | Toyota | GM | Ford | Chrysler | Toyota |
| Fuel cost | -43.10*** | -61.49*** | 26.07 | 2.70 | 2.26 | 4.02 | 8.60 | 30.47 |
| | (5.46) | (17.91) | (20.50) | (15.35) | (1.75) | (6.93) | (8.61) | (14.26) |
| Average competitor fuel cost | 37.77*** | 57.54*** | -30.63 | -0.68 | -1.12 | -1.51 | -4.40 | -28.52* |
| | (5.74) | (17.39) | (19.31) | (14.07) | (3.51) | (7.48) | (10.78) | (11.91) |
| Average same-firm fuel cost | 2.70* | 5.13 | -1.69 | -0.36 | 0.88 | -0.39 | -15.33*** | -5.33 |
| | (1.44) | (3.49) | (2.27) | (1.18) | (2.28) | (1.50) | (4.84) | (3.65) |
| $R^2$ | 0.8946 | 0.7959 | 0.8248 | 0.5659 | 0.9051 | 0.8610 | 0.7074 | 0.8769 |
| # of vehicles | 59 | 22 | 16 | 8 | 30 | 19 | 28 | 4 |

Results from OLS regressions. The dependent variable is the manufacturer price, i.e., MSRP minus the mean regional and national incentives (in thousands). The units of observation are at the vehicle-week-region level. All regressions include vehicle, time, and region fixed effects, as well as third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

Figure 3.4 plots the estimated manufacturer price responses to a one dollar increase in the gasoline price against vehicle miles-per-gallon, separately for each vehicle type. For cars and SUVs, the median responses are price reductions of $779 and $981, respectively. Among cars, the responses range from a price reduction of $3,255 for the 2006 Ford GT to a price increase of $915 for the 2003 Chrysler SRT4. Among SUVs, the responses range from a price reduction of $2,698 for the 2005 GMC Envoy to a price increase of $649 for the 2006 Ford Escape Hybrid.[31] Though manufacturer prices fall for nearly all cars and SUVs, the prices of fuel efficient vehicles fall less and the prices of extremely fuel efficient vehicles actually increase. Turning quickly to trucks and vans, the estimated manufacturer responses are smaller in magnitude, as is the strength of the relationship between the responses and vehicle fuel efficiency; we remain agnostic about the source of these differences.[32]

[31] Appendix Table 3.5 lists the largest positive and negative price effects for both cars and SUVs.

[32] We cannot reject the null hypothesis that the car fuel cost coefficients are equal to the SUV fuel cost coefficients for Ford and Toyota; we can only weakly reject the null for GM. Chrysler is an exception that we explore in more detail below. For each manufacturer, we can reject the null of equality between 1) the car fuel cost coefficients and the truck/van fuel costs coefficients, and 2) the SUV coefficients and the van coefficients. We can reject the null of equality between the SUV and truck coefficients only for Toyota.

Table 3.4: Fuel Efficient and Inefficient Vehicles

| Cars | Brand | mpg | $\widehat{\frac{\partial p}{\partial gp}}$ | SUVs | Brand | mpg | $\widehat{\frac{\partial p}{\partial gp}}$ |
|---|---|---|---|---|---|---|---|
| The Most Positive Manufacturer Price Responses | | | | | | | |
| 2003 SRT4 | Dodge | 32.85 | 0.9148 | 2006 Escape Hybrid | Ford | 30.25 | 0.6485 |
| 2004 Prius | Toyota | 55.05 | 0.5268 | 2006 RX 400h | Lexus | 30.25 | 0.4304 |
| 2006 Prius | Toyota | 55.05 | 0.5227 | 2006 Mariner Hybrid | Mercury | 30.80 | 0.3944 |
| 2005 Prius | Toyota | 55.05 | 0.4971 | 2006 Highlander Hybrid | Toyota | 30.25 | 0.3111 |
| 2005 SRT4 | Dodge | 26.40 | 0.4661 | 2003 Wrangler | Jeep | 19.10 | 0.1551 |
| 2004 SRT4 | Dodge | 26.40 | 0.3740 | 2005 Wrangler | Jeep | 19.65 | 0.1442 |
| 2003 Prius | Toyota | 48.15 | 0.3414 | 2006 Liberty | Jeep | 20.20 | 0.1284 |
| 2004 Neon | Dodge | 32.85 | 0.3305 | 2003 Liberty | Jeep | 21.75 | 0.1246 |
| 2003 Neon | Dodge | 32.85 | 0.3244 | 2003 Durango | Dodge | 16.75 | 0.1155 |
| 2005 Neon | Dodge | 32.85 | 0.2981 | 2006 Wrangler | Jeep | 19.65 | 0.0691 |
| The Most Negative Manufacturer Price Responses | | | | | | | |
| Cars | Brand | mpg | $\widehat{\frac{\partial p}{\partial gp}}$ | SUVs | Brand | mpg | $\widehat{\frac{\partial p}{\partial gp}}$ |
| 2003 XKR | Jaguar | 19.85 | -2.0168 | 2003 H2 | Hummer | 13.65 | -2.3293 |
| 2004 GTO | Pontiac | 18.75 | -2.0239 | 2006 H2 SUV | Hummer | 13.65 | -2.3618 |
| 2004 Marauder | Mercury | 20.30 | -2.0617 | 2004 H1 | Hummer | 13.65 | -2.3711 |
| 2005 Viper | Dodge | 16.95 | -2.1401 | 2003 9-7X | Saab | 13.65 | -2.4298 |
| 2003 Viper | Dodge | 16.95 | -2.1462 | 2003 H1 | Hummer | 13.65 | -2.4511 |
| 2004 Viper | Dodge | 16.95 | -2.1880 | 2003 Escalade | Cadillac | 13.65 | -2.5031 |
| 2003 Marauder | Mercury | 20.30 | -2.2581 | 2006 H2 SUN | Hummer | 13.65 | -2.5640 |
| 2006 Viper | Dodge | 16.40 | -2.4917 | 2005 H2 SUN | Hummer | 13.65 | -2.578 |
| 2005 GT | Ford | 17.40 | -3.2390 | 2006 H1 | Hummer | 13.65 | -2.6173 |
| 2006 GT | Ford | 17.40 | -3.2552 | 2005 Envoy XUV | GMC | 13.65 | -2.6979 |

Based on Appendix Table 3.3 and Figures 3.6 and 3.7.

In order to assess the economic magnitude of these results, we use back-of-the-envelope calculations to (roughly) estimate the extent to which manufacturers offset changes in consumers' cumulative gasoline expenses. We assume an annual discount rate of five percent, a vehicle lifespan of thirteen years, and a utilization rate of 11,154 miles per year (these figures are based upon Department of Transportation estimates). Under these assumptions, the cumulative gasoline expense associated with a one dollar increase in the gasoline prices ranges between \$1,972 and \$7,953 among the sample vehicles; the expense for the median vehicle is \$5,073. We divide the estimated manufacturer responses by the computed cumulative gasoline expense; this ratio provides the percent of cumulative gasoline expenses offset by changes in the manufacturer price.

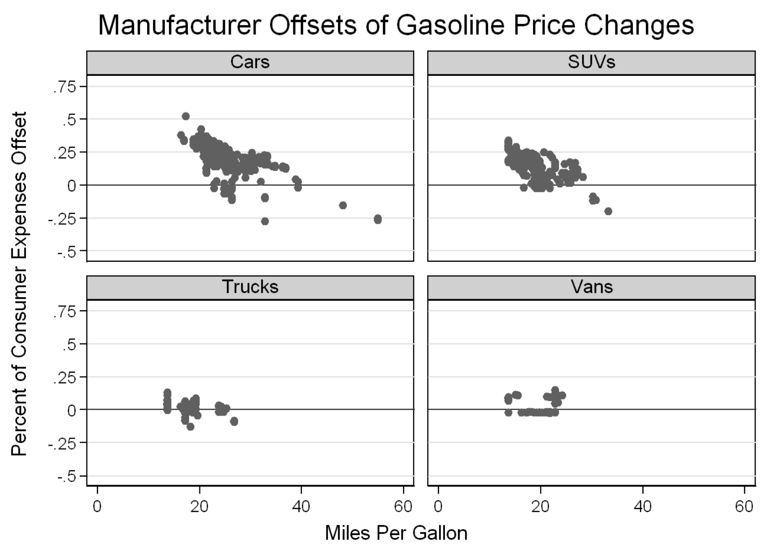


Figure 3.5: The percentages of consumer cumulative gasoline expenses, due to changes in the retail gasoline price, that are offset by changes in the manufacturer price. Each point represents the percentage for a single vehicle. Based on back-of-the-envelope calculations and the regression results of Table 3.3.

Figure 3.5 plots this "offset percentage" against vehicle miles-per-gallon, separately for each vehicle type. The median offset percentage is 18.17 and 15.27 for cars and SUVs, respectively, but climbs as high as 52.17 for cars (the 2006 Ford GT) and as high as 33.92 for SUVs (the 2004 GM Envoy XUV). These calculations are consistent with the notion that manufacturers offset a sizeable portion of the fuel costs that consumers expect to pay over the lifetimes of their vehicles. We wish to emphasize that these numbers should be interpreted with considerable caution. Alternative assumptions regarding the discount rate, the vehicle holding period, and the utilization rate could push the offset percentages higher or lower. Further, as previously discussed, the manufacturer price we use to estimate the regressions – MSRP minus the mean available incentive – could understate the manufacturer responses and the offset percentages if some consumers stack multiple incentives.

Finally, we examine the extent to which manufacturers differ in their pricing strategies. Figure 3.6 plots the estimated manufacturer price responses for cars, separately for each manufacturer. The prices of nearly all GM, Ford, and Toyota cars fall; the median responses for these manufacturers are price reductions of \$610, \$1,180, and \$758, respectively. By contrast, Chrysler lowers its prices for only a minority (38 percent) of its cars, and the median response is a price increase of \$107. Statistical tests provide weak support for the proposition that Chrysler follows a different pricing strategy than the other manufacturers.[33] There are a number of reasons why these differences might exist. For example, Chrysler could simply face distinct demand conditions. Chrysler could also adjust its prices through alternative mechanisms (e.g., through dealer negotiations) that are not observed in the data.

Figure 3.7 plots the estimated manufacturer price responses for SUVs, separately for each manufacturer. The prices of nearly all GM, Ford, and Toyota SUVs fall; the median responses for these manufacturers are price reductions of \$1,315, \$663, and \$754, respectively, and the responses are more negative for fuel inefficient SUVs. By contrast, Chrysler raises the price for a sizeable portion of its SUVs (29 percent), and these price increases occur for the more fuel *inefficient* vehicles. The unexpected pattern exists because Chrysler's fuel cost coefficient is positive and its competitor fuel cost coefficient is negative (see Table 3.3), inconsistent with the profit maximizing pricing rule derived in the theoretical framework.[34] It is difficult to make definitive statements about the optimality of Chryler's pricing strategy, however. For example, we cannot rule out the possibilities that consumers of Chrysler SUVs are distinctly unresponsive to fuel costs, and/or that Chrysler adjusts prices through mechanisms that are unobserved in the data.

[33] The tests are based on the null hypotheses that the various sub-sample regressions produce identical fuel costs, competitor fuel cost, and same-firm fuel cost coefficients. These tests for whether the Chrysler coefficients are identical to those of GM, Ford, and Toyota yield $p$-values of 0.2275, 0.1041, and 0.0506, respectively. The GM-Ford comparison yields a $p$-value of 0.4369, the GM-Toyota comparison yields a $p$-value of 0.1368, and the Ford-Toyota comparison yields a $p$-value of 0.6556.

[34] Statistical tests easily reject the null that Chrysler coefficients for SUVs are identical to those of the other manufacturers; the tests against GM, Ford, and Toyota yield $p$-values of 0.0023, 0.0066, and 0.0024, respectively. By contrast, the GM-Ford comparison yields a $p$-value of 0.9394, the GM-Toyota comparison yields a $p$-value of 0.2628, and the Ford-Toyota comparison yields a $p$-value of 0.9213.

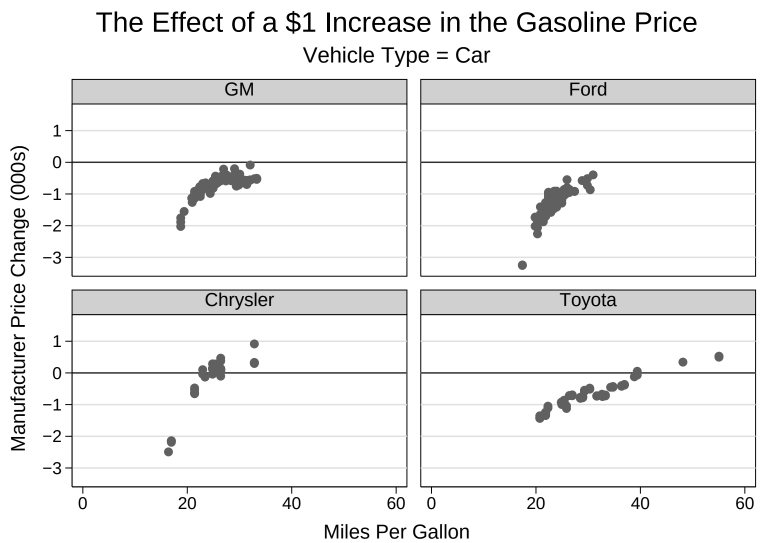


Figure 3.6: The estimated effects of a one dollar increase in the retail gasoline price on the manufacturer price, based on the regression results of Table 3.3. Each point represents the price effect for a single vehicle. See text for details.

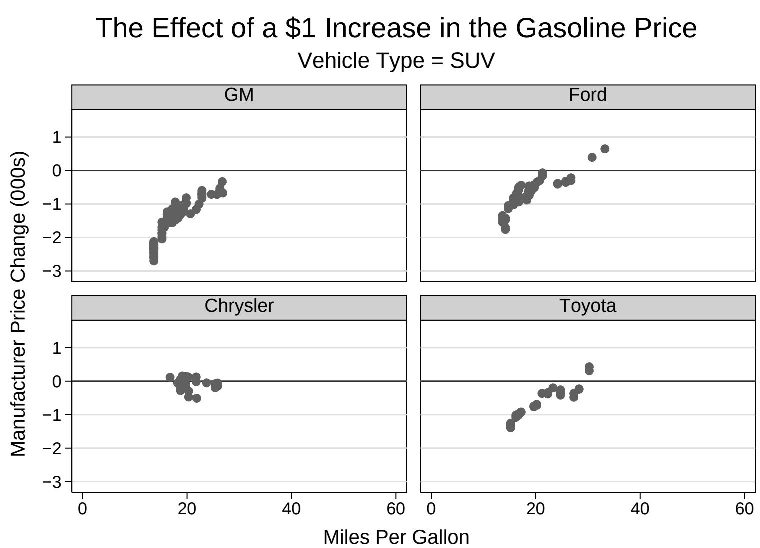


Figure 3.7: The estimated effects of a one dollar increase in the retail gasoline price on the manufacturer price, based on the regression results of Table 3.3. Each point represents the price effect for a single vehicle.

Table 3.5: Fuel Efficient and Inefficient Vehicles

| Cars | Brand | mpg | $\widehat{\frac{\partial p}{\partial gp}}$ | SUVs | Brand | mpg | $\widehat{\frac{\partial p}{\partial gp}}$ |
|---|---|---|---|---|---|---|---|
| The Most Positive Manufacturer Price Responses | | | | | | | |
| 2003 SRT4 | Dodge | 32.85 | 0.9148 | 2006 Escape Hybrid | Ford | 30.25 | 0.6485 |
| 2004 Prius | Toyota | 55.05 | 0.5268 | 2006 RX 400h | Lexus | 30.25 | 0.4304 |
| 2006 Prius | Toyota | 55.05 | 0.5227 | 2006 Mariner Hybrid | Mercury | 30.80 | 0.3944 |
| 2005 Prius | Toyota | 55.05 | 0.4971 | 2006 Highlander Hybrid | Toyota | 30.25 | 0.3111 |
| 2005 SRT4 | Dodge | 26.40 | 0.4661 | 2003 Wrangler | Jeep | 19.10 | 0.1551 |
| 2004 SRT4 | Dodge | 26.40 | 0.3740 | 2005 Wrangler | Jeep | 19.65 | 0.1442 |
| 2003 Prius | Toyota | 48.15 | 0.3414 | 2006 Liberty | Jeep | 20.20 | 0.1284 |
| 2004 Neon | Dodge | 32.85 | 0.3305 | 2003 Liberty | Jeep | 21.75 | 0.1246 |
| 2003 Neon | Dodge | 32.85 | 0.3244 | 2003 Durango | Dodge | 16.75 | 0.1155 |
| 2005 Neon | Dodge | 32.85 | 0.2981 | 2006 Wrangler | Jeep | 19.65 | 0.0691 |
| The Most Negative Manufacturer Price Responses | | | | | | | |
| Cars | Brand | mpg | $\widehat{\frac{\partial p}{\partial gp}}$ | SUVs | Brand | mpg | $\widehat{\frac{\partial p}{\partial gp}}$ |
| 2003 XKR | Jaguar | 19.85 | -2.0168 | 2003 H2 | Hummer | 13.65 | -2.3293 |
| 2004 GTO | Pontiac | 18.75 | -2.0239 | 2006 H2 SUV | Hummer | 13.65 | -2.3618 |
| 2004 Marauder | Mercury | 20.30 | -2.0617 | 2004 H1 | Hummer | 13.65 | -2.3711 |
| 2005 Viper | Dodge | 16.95 | -2.1401 | 2003 9-7X | Saab | 13.65 | -2.4298 |
| 2003 Viper | Dodge | 16.95 | -2.1462 | 2003 H1 | Hummer | 13.65 | -2.4511 |
| 2004 Viper | Dodge | 16.95 | -2.1880 | 2003 Escalade | Cadillac | 13.65 | -2.5031 |
| 2003 Marauder | Mercury | 20.30 | -2.2581 | 2006 H2 SUN | Hummer | 13.65 | -2.5640 |
| 2006 Viper | Dodge | 16.40 | -2.4917 | 2005 H2 SUN | Hummer | 13.65 | -2.578 |
| 2005 GT | Ford | 17.40 | -3.2390 | 2006 H1 | Hummer | 13.65 | -2.6173 |
| 2006 GT | Ford | 17.40 | -3.2552 | 2005 Envoy XUV | GMC | 13.65 | -2.6979 |

Based on Appendix Table 3.3 and Figures 3.6 and 3.7.

## 3.5 Extensions

### 3.5.1 Demand and cost factors

In the main regressions, we estimate a separate time fixed effect for each of the 208 weeks over 2003-2006. These fixed effects capture the combined influence of demand and cost factors that change over time through the sample period. In this section, we use a second-stage regression to decompose the fixed effects into contributions from specific time-varying demand and cost factors. We are particularly interested in whether the retail gasoline price affects manufacturer prices after having controlled for its impact on vehicle fuel costs. Such an effect could be present if higher gasoline prices increase manufacturer production costs or reduce consumer demand through an income effect.[35] One might expect these two channels to partially offset; we can identify only the net effect.

We regress the estimated time fixed effects on different the gasoline price, as well as the prime interest rate and the unemployment rate (demand factors) and deseasonalized price indices for electricity and steel (cost factors). As expected, the estimated time fixed effects exhibit substantial seasonality and peak in the winter weeks. We include 52 week dummies to remove this variation from the regression. We use the Newey and West (1987) variance matrix to account for first-order autocorrelation. The standard errors do not change substantially when we account for higher-order autocorrelation.[36]

Table 3.6 presents the results. Column 1 features only the gasoline price, Column 2 features the gasoline price and the demand factors, Column 3 features gasoline price and the cost factors, and Column 4 features all demand and cost factors. The coefficients are remarkably stable across specifications. In each column, the gasoline price coefficient is small and statistically indistinguishable from zero; gasoline prices appear to have little effect on manufacturer prices after controlling for vehicle fuel costs. The remaining coefficients take the expected signs. Based on the Column 4 regression, a one percentage point increase in prime interest rate reduces manufacturer prices by \$164 and a one percentage point increase in the unemployment rate reduces manufacturer prices by \$104 (though the latter effect is not statistically significant). Similarly, ten percentage point increases in the prices of electricity and steel raise manufacturer prices by \$283 and \$55, respectively.

[35]Gicheva, Hastings, and Villas-Boas (2007) identify an income effect of gasoline prices using scanner data on grocery purchases.

[36]To be clear, we estimate 52 week fixed effects using 208 weekly observations; equivalent weeks in each year are constrained to have the same fixed effect. Of course, the standard errors may inaccurate because the dependent variable is estimated in a prior stage.

Table 3.6: Demand and Cost Factors

| Variables | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| Gasoline Price | -0.015 | 0.011 | -0.102 | -0.096 |
| | (0.036) | (0.059) | (0.088) | (0.067) |
| Interest Rate | | -0.128*** | | -0.164*** |
| | | (0.027) | | (0.034) |
| Unemployment Rate | | -0.315*** | | -0.104 |
| | | (0.073) | | (0.091) |
| Electricity Price Index | | | 0.950* | 2.832*** |
| | | | (0.540) | (0.726) |
| Steel Price Index | | | 0.405*** | 0.549*** |
| | | | (0.113) | (0.152) |
| $R^2$ | 0.5160 | 0.6117 | 0.5829 | 0.6454 |

Results from OLS regressions. The data include 208 weekly observations over the period 2003-2006. The dependent variable is the time fixed effect estimated in Column 3 of Table 3.2. The regressions also include 52 week fixed effects; equivalent weeks in each year are constrained to have the same fixed effect. Standard errors are robust to the presence of heteroskedasticity and first-order autocorrelation. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

### 3.5.2 Lagged retail gasoline prices and gasoline futures

The main results are based on the premise that consumers form expectations about future retail gasoline prices based on current retail gasoline prices. Here, we explore whether consumers consider historical and futures prices when forming expectations about future gasoline prices. Interestingly, statistical tests based on Dicky and Fuller (1979) fail to reject the null that gasoline prices follow a random walk – for example, the $p$-statistic for the deseasonalized national time-series is 0.7035. These tests suggest that knowledge of the current gasoline price is sufficient to inform predictions over future gasoline prices.[37] If consumers form expectations efficiently, therefore, one would not expect historical and/or futures prices of gasoline to influence vehicle purchase decisions.

We construct two new sets of fuel cost variables. The first uses the mean retail gasoline price over the previous four weeks, and the second uses the one-month futures price for retail gasoline. To the extent that consumers are backward-looking and forward-looking, respectively, one should observe that manufacturers adjust vehicle prices with these new fuel cost variables. In conducting the test, we discard regional variation because futures prices are available only at the national level; the units of observation are at the vehicle-week level. The results are therefore comparable to Column 3 of Table 3.2.

Table 3.7 presents the regression results. Columns 1 and 2 include variables based on mean lagged gasoline prices and gasoline futures prices, respectively. The fuel cost coefficients are -64.55 and -47.66; the competitor fuel cost coefficients are 50.01 and 63.32. The coefficients are statistically significant and consistent with the theoretical model. Still, the more interesting question is whether these variables matter after controlling for the current price of retail gasoline. Columns 3 and 4 include variables based on mean lagged gasoline prices and gasoline futures prices, respectively, together with variables based on the current gasoline price. Each of the coefficients takes the expected sign and statistical significance is maintained for all but two coefficients. Finally, Column 5 includes variables based on mean lagged gasoline prices and variables based on gasoline futures prices. The coefficients are precisely estimated and again take the correct sign.[38]

[37] The result is consistent with the academic literature and statements of industry experts. For example, Alquist and Kilian (2008) find that the current spot price of crude oil outperforms sophisticated forecasting models as a predictor of future spot prices, and Peter Davies, the chief economist of British Petroleum, has stated that "we cannot forecast oil prices with any degree of accuracy over any period whether short or long..." (Davies 2007) See also Davis and Hamilton (2004) and Geman (2007).

[38] In the working paper version, we estimate an impulse response function based on ten lags of the fuel cost variables. The results are broadly consistent with those presented here. Finally, we note that the inclusion of *all* the fuel cost variables – i.e., those based on lagged,

Table 3.7: Gasoline Price Lags and Futures Prices

| Variables | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| Fuel cost | -64.55*** | | -36.51*** | | -30.08*** |
| | (8.77) | | (10.65) | | (8.42) |
| Average competitor fuel cost | 50.01*** | | 23.19** | | 30.24*** |
| | (8.16) | | (10.09) | | (9.93) |
| Fuel cost | | -47.66*** | | -35.52** | -31.69*** |
| | | (7.11) | | (16.42) | (9.39) |
| Average competitor fuel cost | | 63.32*** | | 19.87 | 27.73** |
| | | (10.44) | | (24.95) | (13.21) |
| Fuel cost | | | -29.70*** | -22.58 | |
| | | | (10.83) | (16.46) | |
| Average competitor fuel cost | | | 27.70*** | 33.38* | |
| | | | (8.14) | (18.87) | |
| $R^2$ | 0.5291 | 0.5286 | 0.5295 | 0.5295 | 0.5305 |

Results from OLS regressions. The dependent variable is the manufacturer price, i.e., MSRP minus the mean national incentive (in thousands). The sample includes 59,971 observations on 681 vehicles at the vehicle-week level. Fuel cost variables labeled “lagged retail” are constructed using the mean retail gasoline price over the previous four weeks. Fuel cost variables labeled “futures” are constructed using the one-month futures price of retail gasoline. Fuel cost variables labeled “retail” are constructed using the current retail gasoline price. All regressions include the appropriate average same-firm fuel cost variable(s). The regressions also include vehicle and time fixed effects, as well as third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

The finding that consumers use historical gasoline prices and gasoline futures prices to form expectations for gasoline prices is interesting, in part because both the empirical evidence and the conventional wisdom of industry experts suggest that gasoline prices follow a random walk (as we outline Section 3.3). One could argue that some consumers form inefficient expectations for future gasoline prices. Alternatively, some consumers may be imperfectly informed about the current gasoline price; these consumers could rationally turn to alternative sources of information, such as historical prices and/or futures prices. We are skeptical that our data can untangle these informal hypotheses and hope that future research better addresses the topic.

### 3.5.3 Vehicle inventories

In the theoretical model, we assume that manufacturers have full information about consumer demand. It is not clear whether the assumption is justifiable based solely on theoretical grounds. For example, manufacturers may receive only noisy signals about demand, and accurate information may be costly to obtain. In such an environment, one might expect manufacturers to set their prices based on their (easily) observed inventories. As a specification test, we estimate the empirical model controlling for inventories.

We collect data on automobile inventories from from Automotive News, a major trade publication. We measure inventory using "days supply," the current inventory divided by mean daily sales over the previous month. The measure should be high when demand is sluggish and low when demand is great. Unfortunately, we observe days supply at the at the month-model level. Thus, the data do not vary across weeks within a month, and lump all vehicles within a given model (e.g., the 2003 Dodge Neon and 2004 Dodge Neon). We map the data into the main regression sample by using cubic splines to interpolate weekly observations. We then apply the days supply to every vehicle in the model category. The procedure generates a regression sample of 500 vehicles and 41,822 vehicle-week observations.[39]

Table 3.8 presents the regression results. In Column 1, we re-estimate the same specification as in Table 3.2, Column 3 using only those observations for which we have information on inventories. The fuel cost and competitor fuel cost coefficients are -69.23 and 53.16, respectively.[40] We add the days supply

present, and futures gasoline prices – appears to over-tax the data. The coefficients produced are unreasonably large and imprecisely estimated.

[39]We have inventory data for 500 of the 589 domestic vehicles in the data. The Toyota data are insufficiently disaggregated to support analysis. The mean days supply among the 41,822 vehicle-week observations is 92.18, and the $25^{th}$, $50^{th}$, and $75^{th}$ percentiles are 62.26, 84.63, and 109.42, respectively.

[40]The fact that these coefficients are close to those produced by the full sample provides

Table 3.8: Manufacturer Prices, Fuel Costs, and Inventories

| Variables | (1) | (2) |
|---|---|---|
| Fuel cost | -69.23*** | -69.11*** |
| | (11.57) | (11.54) |
| Average competitor fuel cost | 53.16*** | 53.00*** |
| | (9.79) | (9.76) |
| Average same-firm fuel cost | 1.95 | 1.94 |
| | (3.36) | (3.36) |
| Vehicle inventory | | 0.0001 |
| | | (0.0001) |
| $R^2$ | 0.6202 | 0.6203 |

Results from OLS regressions. The dependent variable is the manufacturer price, i.e., MSRP minus the mean national incentive (in thousands). The sample includes 41,822 observations on 500 vehicles over the period 2003-2006, at the vehicle-week level. The regressions include vehicle and time fixed effects, as well as third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

measure to the specification in Column 2. The fuel cost and competitor fuel cost coefficients of -69.11 and 53.00 are virtually unchanged. The days supply coefficient is small and statistically indistinguishable from zero, though we are wary of interpreting this coefficient too strongly because inventories may be correlated with vehicle-time specific cost and demand shocks. Overall, the results could suggest that manufacturers respond to changes in demand conditions before these changes affect inventories; one might infer that manufacturers are well informed about consumer preferences, consistent with our theoretical framework.[41]

some comfort that the smaller inventory sample does not introduce sample selection problems or other complexities.

[41] One could alternatively attribute the results to the poor quality of the inventories data.

## 3.6 Conclusion

We provide empirical evidence that automobile manufacturers adjust vehicle prices in response to changes in the price of retail gasoline. In particular, we show that the vehicle prices tend to decrease in their own fuel costs and increase in the fuel costs of their competitors. The net effect is such that adverse gasoline price shocks reduce the price of most vehicles but raise the price of particularly fuel efficient vehicles. We argue, based on theoretical micro foundations, that these empirical results are consistent with the notion that automobile manufacturers set prices as if consumers value (low) fuel costs. The results suggest that market-based policy instruments such as cap-and-trade regulation or carbon taxes would likely prove effective in mitigating the negative externalities associated with gasoline combustion in automobiles. The results do not speak, however, to the optimal magnitude of any policy responses; we leave that important matter to future research.

# Chapter 4

# Toward a Comprehensive Assessment of Road Pricing Accounting for Land Use

### Coauthored with Clifford Winston, Brookings Institution[1]

## 4.1 Introduction

Congestion on U.S. highways is a well-known social and economic problem, which becomes progressively worse every year.[2] Travel delays impose large costs on motorists, truckers, and shippers that currently approach some $40 billion annually (Winston and Langer (2006)). Economists have repeatedly attributed the problem to policymakers' failure to implement marginal cost congestion tolls to charge road users efficiently for their contribution to delays.

By undercharging vehicles for using the nation's roadways, policymakers have also reduced the per-mile cost of commuting (including out-of-pocket and travel time costs) for most motorists and distorted the development of metropolitan areas by inducing households to live in more-distant and lower-density locations, which contributes to urban sprawl. Precise definitions of sprawl and estimates of its costs are elusive, because it is difficult to characterize an optimal pattern of land use.[3] At the same time, it is likely that households' residen-

---

[1] This chapter of my dissertation was published in the 2008 Brookings-Wharton Papers on Urban Affairs under the same title.

[2] We are reminded of this fact by media summaries of the Texas Transportation Institute's latest Urban Mobility Report (mobility.tamu.edu/ums/).

[3] As a working definition, Nechyba and Walsh (2004) define sprawl as a tendency toward lower city densities as city footprints expand. Optimal land use requires efficient pricing of rural land and city services. Determining efficient prices for both is difficult.

tial location decisions—while maximizing their utility—have resulted in socially inefficient outcomes because they reduce economies of agglomeration.

For instance, according to the U.S. Census, between 1970 and 2000 the metropolitan population in the United States grew approximately 60 percent. We would therefore expect in a given city that a representative central city neighborhood with 10,000 residents in 1970 would house an additional 6,000 people in 2000. These people would live in new and converted housing, pay taxes, and consume city services. Neighborhood schools might add a new wing for more classrooms; police and fire departments might hire additional employees; and so forth. But, in fact, such expectations have not been realized; because according to the U.S. Census, central city density declined roughly 35 percent between 1970 and 2000—that is, we would observe 3,500 people moving out of the neighborhood. In effect, 9500 people (6,000 plus 3,500), nearly the entire original population of the neighborhood, chose to live in newly constructed homes farther from the urban center on land that may not have even been part of the metropolitan area in 1970.

Of course, these residents are able to buy more house per dollar in the suburbs; but they also incur the costs of longer commutes and other trips that sharply increase per capita vehicle miles traveled within the city; they live in residential densities where each household requires more feet of utility lines, more miles of school bus routes, and longer police and fire department response times than are called for in dense neighborhoods closer to the urban center; yet the schools, fire, and police services in the more centrally located neighborhoods might now have excess capacity because their population has decreased. (If the communities close to and the communities far from the urban center are located in different municipalities, which is almost always the case in U.S. metropolitan areas, then extra resources are not likely to be reallocated.)

In sum, although the residents' location choices reflect their self-interest, the city's economy would be more efficient if both current residents remained and new residents relocated in higher density urban communities or sub-centers, instead of locating in lower density suburban communities. The divergence between residents' choices and land use efficiency can be explained by public investments in limited access highways (Baum-Snow (2007)) and the undercharging of highway travel during congested periods. (In addition, zoning laws and other land use controls may also induce households to make choices that conflict with the public interest.)

It is well-known that congestion pricing can reduce travel delays and smooth the flow of highway traffic throughout the day, but its effect on land use has received little empirical attention. This paper presents rough estimates of the costs and benefits of congestion pricing, accounting for its effects on land use that could help reduce inefficient urban sprawl. Quantifying the full effects of road pricing is important because policymakers are giving it unprecedented consideration as a way to reduce congestion and provide stable, long-term financing for the nation's highways without unduly affecting households' welfare.[4]

Because we expect road pricing's effects on road users' travel times and out-of-pocket expenses to be capitalized in property values, we develop a hedonic model of housing prices that includes travel delays and unpriced congestion (that is, the benefits of not implementing road pricing) as influences. Housing prices are also influenced by elements of land use, such as city wide density and entropy (i.e., the spatial variation in density), which are also treated as simultaneously determined by travel delays, unpriced congestion, and housing prices. Finally, travel delays and unpriced congestion are determined by characteristics of the metropolitan area.

Our model allows residents to increase their welfare by moving in response to the adoption of congestion pricing or by remaining in their present location and, in most cases, benefit from improvements in land use such as greater density. Either response will increase the social net benefits of road pricing and reduce its adverse distributional effects. Policymakers have generally opposed road pricing because it imposes direct losses on most travelers; but by accounting for changes in land use, we show that policymakers can substantially reduce these undesirable effects by returning some of the congestion toll revenues to households through lower local taxes and still have sufficient revenues to finance a large portion of the road system's maintenance and expansion.

Based on a sample of the 98 largest Metropolitan Statistical Areas (MSAs) in the nation, we find that efficient road pricing would generate $120 billion in annual revenues (2000 dollars), while reducing the value of the annual flow of services from housing $80 billion (2000) dollars, thus generating an annual net benefit of $40 billion. Our estimate of the benefits of congestion pricing is considerably greater than previous estimates that do not account for adjustments in land use and represents a first step toward accounting fully for road pricing's benefits. We conclude that policymakers should recognize that road pricing mitigates

[4]Congress established the National Surface Transportation Policy and Revenue Commission and the National Surface Transportation Infrastructure Financing Commission to consider policies to relieve traffic congestion and to meet short-term and long-term highway revenue shortfalls among other challenges. The U.S. Department of Transportation has encouraged the nation's cities to submit proposals for reducing congestion that it would help fund. The Department has indicated that it would help Mayor Michael Bloomberg finance his plan to reduce traffic in Manhattan by charging tolls to drivers entering the busiest parts of the borough. It also indicated that it would help San Francisco pay for construction on Doyle Drive, which approaches the Golden Gate Bridge and handles some 90,000 vehicles a day, if local officials agreed to charge a congestion toll for the road. The economic effects of and political obstacles to road pricing are well documented in, for example, Small, Winston, and Evans (1989), Mohring (1999), Santos (2004), and Lindsey (2006).

congestion and improves the quality in life in a metropolitan area by improving land use.

## 4.2 Conceptual Framework

The standard conceptual framework for motivating and assessing the economic effects of congestion tolls has been presented so often that it is referred to as the conventional diagram (Lindsey (2006)). Our discussion at this point proceeds without the diagram, which assumes residents' locations are fixed, and accounts for road pricing's standard effects as well as its potential effects on land use.

When there is a low volume of traffic on a road, every vehicle is able to travel at free flow speed and each driver incurs the private cost of a trip, which includes vehicle operating costs and the value of the driver's travel time. As traffic volume increases, drivers must reduce their speed and each driver's private cost diverges from the social cost of his or her trip because the social cost includes the driver's contribution to congestion as indicated by the cost of the delay incurred by other drivers. An efficient congestion toll applied to all drivers on the congested road bridges the gap between the private average cost of drivers' trips and the marginal social cost of their trips by making them pay for the delays they impose on other drivers; hence, scarce road capacity is used efficiently by drivers whose marginal benefit of driving is equal to or above the marginal social cost of their trips.

By affecting drivers' behavior, the toll reduces congestion on the road and raises travel speeds. In the short run, when road users' residences and workplaces are fixed, motorists may respond differently to congestion tolls because their values of travel time differ. Behavioral responses include the choice by some motorists to use the next best alternative to peak-period travel on the road, which may be traveling on it at a time when it is less congested, using a less congested but undoubtedly slower route, using another mode, or not traveling at all. In either case, these drivers are clearly worse off from the toll. Other motorists will stay on the road because their next best alternative is worse than continuing to use the road; but on balance they are worse off because the out-of-pocket costs of the toll exceed their value of the travel time savings. Still other motorists will stay on the road and are better off, because their value of the time savings exceeds the out-of-pocket costs of the toll. In fact, other motorists with high values of time who were deterred from using the road in congested conditions will now find that they are also better off using the tolled road. But, on average, travelers' welfare will be reduced by the toll because the initial full price of travel, including the cost of travel time, was below the marginal social cost of travel. On net, the toll results in a welfare gain, but only because the toll revenues to the government exceed the net loss to motorists.[5]

In the long run, motorists can change where they live and work in response to a toll, while continuing to live in the same metropolitan area or by moving to a new metropolitan area. In this paper, we confine our long-run analysis to motorists' changes of where they live within the same metropolitan area. We discuss the likely effects of the other long-run responses on our findings in the conclusion. We also discuss our findings in light of how policymakers' can allocate the revenues raised by congestion pricing to ameliorate distributional concerns.

From a theoretical perspective, motorists' changes in their residential location in response to road pricing and the effect on land use can be determined from the relationship between transportation costs and household location decisions analyzed by, for example, Alonso (1964), Mills (1967), and Muth (1969). We draw on the theoretical discussion presented by Pickrell (1999).

Households locate where the costs of commuting to work exactly balance the savings in housing costs that accrue from living in a more distant location. Formally, this result is derived under the assumption that a household chooses a combination of housing, h, and other goods, g, to maximize a utility function, $U(h, g)$, subject to a budget constraint given by $Y = pgg+ph(d)h+T(d, v)$, where $Y$ is income, $pg$ denotes the composite price of the nonhousing good, $ph(d)$ is the price per unit of housing, which is a function of distance $d$ from the workplace, and $T(d, v)$ denotes transportation costs for commuting to and from work, which depend on commuting distance and the value of travel time, $v$, which itself is a function of income, $Y$.

Assuming for simplicity that households have identical preferences for identical units of housing, the relevant first-order condition for a constrained utility maximum is $-h(\partial ph/\partial d) = \partial T/\partial d$ (see Pickrell for the complete derivation). The condition states that at the household's equilibrium location, the change in its housing costs from moving slightly closer to or farther from the workplace exactly offsets the resulting change in commuting costs. We can rewrite the first-order condition to obtain the household's bid-rent function, $(\partial ph/\partial d) = -(\partial T/\partial d)/h$, which indicates that the price the household is willing to pay for housing declines with distance from its workplace in proportion to the rate of increase in transportation costs.

If the assumptions of identical housing preferences and housing units are relaxed, households will consume different quantities of housing; in particular, they will respond to the decline in housing prices with distance from the city center by demanding more housing services at more distant locations. Thus larger households and others with preferences for more residential space will tend to seek more distant locations because they can realize significant savings in

[5] Lindsey (2006) discusses in detail the simplifying assumptions of the standard framework for analyzing congestion pricing and its findings.

housing costs. Home builders will respond to declining land prices at increasing distances from the city center by substituting progressively more land for capital; that is, by constructing lower density housing. The result is that the density of residential development will decline as the distance from the city's central business district increases, and households who live in lower density developments will incur greater commute distances.

Given this framework, we can assess how households will adjust their locations in response to the adoption of road pricing and determine the impact on land use. Congestion tolls will increase household members' out-of-pocket expenses and reduce their travel time. But as noted, households on average will face higher per-mile transportation costs and a higher rate at which their commuting costs rise with increasing distance from their workplaces. Because households seek to locate where the savings in land and housing costs from distant locations offset the increase in commuting costs, the increase in per-mile commuting costs will induce some households to seek closer and higher-density residential locations. Those households who make this adjustment increase utility by reducing the out of pocket cost of the toll and travel time costs—savings in transport costs that presumably exceed the increase in land and housing costs. In the process of moving closer to their workplaces, such households also reduce the cost of social services by increasing citywide density.

Wheaton (1998) shows that congestion tolls in a monocentric city should increase density, with the largest increase at the city center and increases in other parts of the city decreasing with distance from the central business district. Lee (1992) assesses congestion pricing in a polycentric city and argues that it should increase density in the central city as well as in suburban sub-centers. He also suggests that densities should increase more in the part of the suburban sub-center that is closest to the central city than they should in other locations, thereby decreasing the variation in density because residents would take account of both their distance from the central city and the sub-center to reduce total transport costs.

We illustrate these ideas in Figure 4.1 with a two dimensional cut of a stylized city that consists of a central business district, two sub-centers, and suburbs. The figure shows that congestion pricing causes the urban density function to have higher peaks and fewer neighborhoods in low-density areas, indicating greater densities near the city center and sub-centers, as well as less variation in density.

Figure 4.1: Urban Density Gradients with and without Congestion Pricing

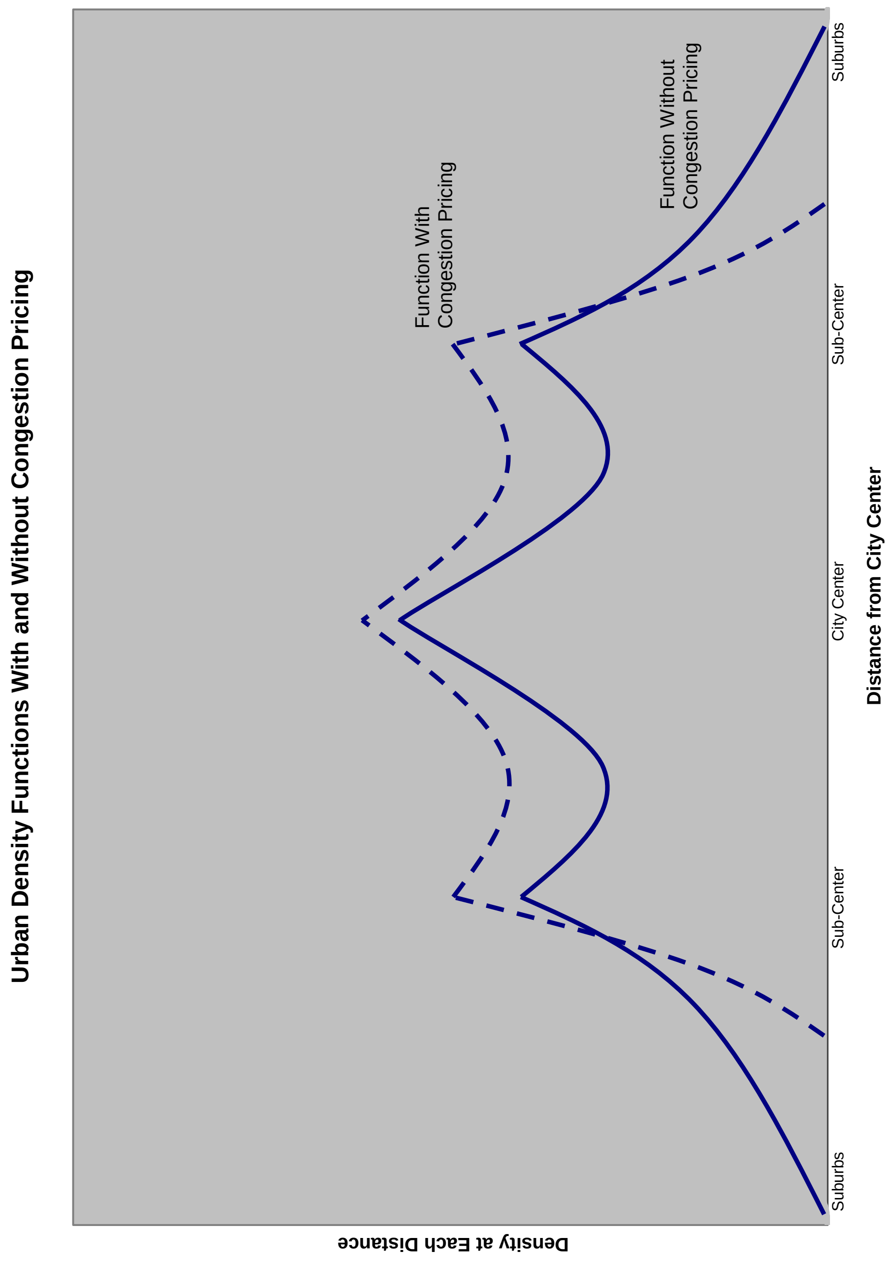

## 4.3 Econometric Approach

How can one estimate the economic effects of road pricing accounting for its impact on land use? A disaggregate approach for a metropolitan area would model the determinants of a commuter's choice of mode, departure time, destination, route, and residential location and simulate how these choices change in response to an efficient congestion toll. Given the aggregate change in land use measures implied by the changes in residential location, it would be possible, in principle, to estimate how households' responses affect the cost of city services and other costs related to sprawl. The change in social welfare from road pricing would be obtained by summing the costs and benefits to residents and the government from the changes in peak-period travel conditions (namely, the toll and travel time) and the resulting changes in land use.

Unfortunately, the data and modeling requirements of a disaggregate approach—especially determining a commuter's residential location alternatives and their attributes—are formidable.[6] In addition, as noted, little quantitative evidence exists on the costs of sprawl, so it would be difficult to estimate directly how changes in land use would affect these costs. Finally, even if one could estimate and simulate a disaggregate model for one metropolitan area, its findings would not necessarily generalize to other areas.

As an alternative and more tractable approach, we draw on the idea—long recognized by economists—that housing prices reflect many factors including residents' accessibility to work and recreational activities. Accordingly, we estimate a basic hedonic model of housing prices across U.S. metropolitan areas including, among other influences, highway congestion variables. We then simulate how annualized housing values would change if efficient prices were set to internalize congestion costs. An advantage of our approach is that we can extend the specification to include measures of land use as endogenous determinants of housing prices and allow highway congestion variables to affect land use. Thus, road pricing is modeled as having a direct effect on housing prices and an indirect effect through its impact on land use. A disadvantage of our approach is that the current state of economic theory enables us to identify the model only through exclusion restrictions.

*Extending a Hedonic Model of Housing Prices.* We wish to extend a basic

[6] Discrete choice models with random parameters are specified to capture unobserved preference heterogeneity. In practice, researchers have often found that it is necessary to include more than one observation for each individual in the sample to obtain precise estimates of unobserved deviations from mean tastes. In the case of a disaggregate residential location choice model, this could be done by estimating a household's ranking of alternative residential locations. But such data are not publicly available, so it is likely that a researcher would have to conduct a new and expensive survey of residential location choices to obtain satisfactory estimates of preference heterogeneity.

hedonic model of housing prices (Song and Knaap (2003) is a recent example), which is typically specified as a function of attributes of the housing stock and characteristics of the metropolitan area, to capture salient features of highway congestion and land use. A useful starting point is the monocentric city model, which suggests that commuting costs should affect home prices. In our case, the time costs of travel delays caused by congestion should decrease home prices for two reasons. First, residents incur costs from longer commutes whether by auto or surface transit, especially during peak-periods, and from longer non-work trips, some of which may be taken in congested conditions. Second, residents incur costs because they have to wait longer for people who provide them with services such as deliveries or repairs. These costs could become quite large if emergency police, fire, or medical services are delayed.

The cost of delays must be balanced against the benefit that is realized by residents because they and the people who provide them with services do not face out-of-pocket costs, besides vehicle operating costs, for driving whenever and wherever they choose, regardless of the social costs.[7] We capture the effect with a metropolitan-wide measure of unpriced congestion—that is, the difference between private average costs and the social marginal costs of driving—which should increase home prices because residents can spend more money on housing if their auto transportation and services are subsidized. On balance, we expect that the net effect of delays and unpriced congestion is to increase home prices because, as noted, the average resident benefits from the absence of congestion pricing. We treat delays and unpriced congestion as endogenous because both capture the economic vitality of a metropolitan area and could be correlated with unobserved metropolitan area characteristics that affect housing prices.

Land use may affect home prices as consumers balance the benefits from greater proximity to social, cultural, and economic opportunities with the cost of crowding, noise, and a higher likelihood of crime. The basic variable for characterizing urban residential land use is citywide density, or population per unit of land area. Given that individuals choose a combination of distance from employment and lot size (land per person) to maximize utility, density is simply the citywide aggregation of individual households' lot size decisions. We expect density to have a positive effect on home prices because the economies of agglomeration are likely to outweigh the diseconomies of crime and noise. Of course, density must be treated as endogenous to verify this effect empirically.[8]

It is also important to characterize the spatial variation in density within a

[7]Similarly, housing prices may be higher in the absence of HOV lanes and on-street parking regulations that restrict driving. However, these disincentives are difficult to measure in our context.

[8]Our specification will also allow the effect of density to vary with commute length, enabling its net benefits to fluctuate throughout the city.

city because although two cities may have identical overall densities, their costs of providing city services may differ if one is characterized by an extremely dense urban center surrounded by low density suburbs and the other by a series of fairly dense sub-centers. Let $x_i$ be the density of land area $i$, which is smaller than the entire city, and $N$ be the number of land areas in the sample. A measure of entropy, which describes the extent of spatial variation, is given by:

$$Entropy = \left(\frac{1}{ln(N)}\right)\left\{\sum_{i=1}^{N}\left(\frac{x_i}{\sum_{i=1}^{N} x_i}\right)\left(ln\frac{\sum_{i=1}^{N} x_i}{x_i}\right)\right\}$$

Entropy ranges from 0 to 1, with higher values implying more uniform density and greater sprawl and smaller values implying more variation in density and less sprawl.[9] For example, if every census tract in a city has exactly the same density, then the city would have entropy of 1. But if a city had a mix of densities, then the city's entropy would be lower.

Although it is possible to measure entropy, its a priori effect on home prices is not clear. Consider an increase in entropy, which means that density will be spread more evenly across neighborhoods in the city. Some suburban residents would benefit from this change because moderately dense neighborhoods, with their cultural and economic attractions, would become more accessible. But other residents may prefer to live in a low density neighborhood and would find it more difficult to do so, while others may find that the benefits from very dense urban corridors have been diluted. Thus the net effect of greater entropy on home prices, simultaneously increasing accessibility to certain attractions but limiting the extent to which preference heterogeneity is accommodated, must be resolved empirically.[10]

In Figure 4.2, we present some stylized density functions that characterize cities in our sample to illustrate how density and entropy interact to generate varied urban forms.[11] For example, Los Angeles, California and Jersey City, New Jersey exhibit high densities and high entropies because they are moderately—but

[9] The measure is similar to a Gini Coefficient but has the advantage that it is independent of the number of observations. The Gini Coefficient allows for observations with zero density, but such observations do not arise here.

[10] It may be useful to also identify spatial variation that arises between cities with one dense core and those with multiple small sub-centers. A measure of centrality, the Moran coefficient, can be used to characterize how close together the city's population is located spatially. In our empirical work, we found that we could not use the Moran coefficient because it exhibited little variation and it did not have a statistically significant effect on home prices. Other land use measures that we explored in our empirical work, but that did not perform as well as density and entropy, were maximum density, density at the 90th percentile of census tract density, and a Geary coefficient (an alternative measure of centrality).

[11] Citywide density and entropy in the examples are not precisely derived from actual U.S. Census data.

uniformly—dense throughout their metropolitan areas. Their high densities are important because they indicate that these cities sprawl substantially less than a city such as Phoenix, Arizona, which is characterized by low density and high entropy, and is commonly thought of as a sprawling city with few checks on development and no dense residential centers. We tended to find notable variation in density (low entropy) across census tracts in small cities with low average density, such as Little Rock, Arkansas and Albany, New York, rather than in larger, more densely populated cities. In fact, our sample contains very few cities that are non-sprawling, as defined by very high density and low entropy. Boston, Massachusetts is probably the best example of a non-sprawling, multi-centric city, although it is not one of the 10 lowest entropy cities in our sample.

*Land Use and Highway Congestion Models.* Density and entropy cannot be treated as exogenous in our framework because both are likely to be affected by house prices (for example, the monocentric model predicts that a change in the bid-rent function influences population density) and correlated with unobserved influences on housing.

We are not aware of previous econometric models that seek to estimate the determinants of citywide density and entropy, but a plausible specification is that metropolitan land use measures are influenced by average home prices, highway congestion variables, and metropolitan area characteristics. An increase in all home prices throughout a metropolitan area is likely to reduce density but increase entropy as some households move further away from employment centers to reduce housing costs. We expect an increase in travel delays to increase density and decrease entropy as households move closer to work to avoid the higher costs of commute time and possibly reduce the delays to and wait times for emergency and non-emergency services. Conversely, unpriced congestion reduces the total cost of driving for most drivers, which encourages people to live further from work; thus, it decreases density and increases entropy.

Land use variables are likely to be influenced by metropolitan area characteristics originating from historical development patterns, which have imposed an underlying form on the city without necessarily affecting current home values. For example, certain features of a metropolitan area's climate affect the suitability of land for farming, which is likely to affect density, while the number of municipalities within the MSA reflects the historical development of sub-centers that may affect population dispersion and entropy.

As noted, it is appropriate to treat the highway congestion variables—delay and unpriced congestion—as endogenous. Both variables are both strongly related to the (traffic)volume/(highway)capacity ratio—but they are distinct from each other—and are therefore a function of metropolitan area characteristics, such as income and natural limits on road building, which affect this ratio.

*Summary and Final Form of the Model.* The model of housing prices, land

Figure 4.2: Representative Density Functions

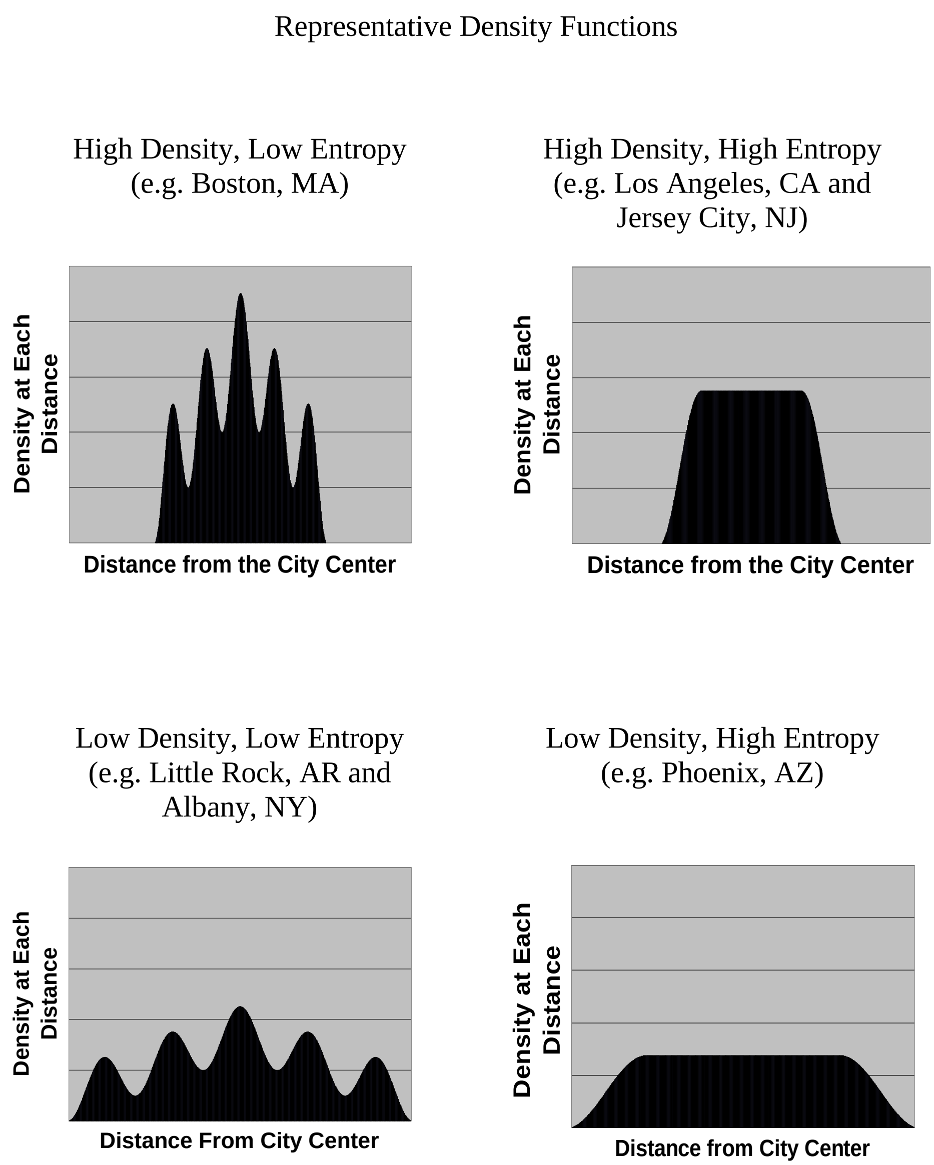

use, and highway congestion variables that we have proposed can be summarized as:

*Housing Prices = f(highway congestion variables, land use measures, housing stock attributes, metropolitan area characteristics)*
*Land Use = g(highway congestion variables, housing prices, metropolitan area characteristics)*
*Highway Congestion variables = h(metropolitan area characteristics affecting the volume/capacity ratio).*

Theory and empirical evidence suggest that commuters self-select into residential locations that can be classified by their proximity to places of employment. For example, Calfee and Winston (1998, 2001) argue that in a given metropolitan area, commuters with the highest values of travel time live close to their workplaces (say, within a 5 minute commute) and those with lower values of travel time tend to live farther from their workplaces. When a house is sold, it is likely that the new homeowner's commute time will be similar to the previous owner's commute time if the cost of commuting has not changed. We therefore specify average housing price equations for each commuting time block within a metropolitan area (as indicated below, the time blocks are less than 5 minutes from the workplace, 5-19 minutes from the workplace, 20-45 minutes from the workplace, and greater than 45 minutes from the workplace), while we specify equations for density, entropy, delay, and unpriced congestion at the metropolitan level. Given this specification, the effect of road pricing on residential location is captured by residents' shifts to new commuting time blocks.

We stress that the commute time blocks are not proxies for a commuter's distance from the city center; thus, a resident with a commute of less than 5 minutes from the workplace lives close to work but does not necessarily live close to the central business district. In addition, by using commute time rather than distance from the central city, we are able to generalize from the monocentric city model and allow for more realistic commuting patterns (including commuting between sub-centers). We are also able to use the Census Public Use Microdata, which includes information on households' commute time, but does not include information on the exact location of each household and workplace. Unfortunately, data are not available to enable us to disaggregate density, entropy, and the volume-capacity ratio by time group.

Specifying commute time blocks does not affect the expected signs of the highway congestion variables, as average delay should have a negative effect and unpriced congestion should have a positive effect on housing prices for all time blocks. For the land use variables, density should have a positive effect on housing prices for all time blocks while entropy's effect is unclear and may vary across time blocks. To understand why, consider an increase in entropy. Residents in the farthest time blocks will benefit most from greater accessibility to certain attractions in dense neighborhoods, while residents in the closest time block may not be affected much because they live in dense neighborhoods close to employment centers regardless of any change in city-wide entropy.

The effects of housing prices on the land use variables will also vary to some extent by commute time block. Households may be influenced to move closer to or further from employment centers when housing prices in their commuting time block increase. We expect households to move farther from places of employment if the savings in housing prices exceed the increase in (total) commuting costs. We expect households to move closer to places of employment if the increase in housing prices is less than the savings in (total) commuting costs.

Given these considerations, we can analyze how changes in home prices within each commuting time block would be expected to affect density and entropy, ceteris paribus. We expect density to decrease if housing prices for the shortest commuting time block increase because households will move to less dense areas farther from employment centers. The effect of a change in housing prices in the middle two commute time blocks (5 to 19 minutes and 20 to 45 minutes) on density will depend on the volume of households who shift to time blocks closer to and further from employment centers. If housing prices for the longest commuting time block increase, we expect density to increase because households from that time block will move closer to employment centers.

Entropy measures the relative variation in density; thus, if low-density areas gain proportionately more households than high-density areas gain, the variation in density will decrease and entropy will increase. This occurs when housing prices in the farthest commute time block increase because households move marginally closer to employment centers, thereby increasing the (below average) density in the 20 to 45 minute time block. We therefore expect an increase in housing prices in the farthest commute time block to increase entropy.

An increase in housing prices in the closest commute time block will cause households who have the longest commute in that block to move because they are more likely than other households to find that the savings in housing costs exceed the increase in commuting costs. The population density where such households lived is undoubtedly lower than the population densities in other parts of the commute time block because it is furthest from the employment center, so households' relocation increases the variation in density in the time block (decreases entropy). Entropy in the rest of the city is likely to change very little because density will increase the most in those areas that have the shortest commute lengths in each commute block, which is where density is relatively high. In sum, we expect an increase in housing prices in the closest commute time block to decrease citywide entropy by decreasing entropy in that block and leaving entropy unchanged elsewhere in the city.

The effect of a change in housing prices in the middle two blocks on entropy is ambiguous, but we can see from the preceding cases that the variation in citywide density tends to decline (entropy increases) as housing prices in commute blocks further from employment centers increase. Thus, we would expect that compared with an increase in housing prices in the 5 to 19 minute commute time block, an increase in housing prices in the 20 minute to 45 minute commute time block will cause citywide entropy to decline by a smaller amount (or increase entropy by a greater amount).

Accounting for the inclusion of commuting time blocks, the important expected signs of the model can be summarized as:

| | Expected Signs Effect On: | | |
|---|---|---|---|
| | Housing Prices | Density | Entropy |
| Average Delays | - | + | - |
| Unpriced Congestion | + | - | + |
| Density | + | — | — |
| Entropy | ? | — | — |
| Housing Prices (closest in) | — | - | - |
| Housing Prices (middle blocks) | — | ? | ? |
| Housing Prices (farthest out) | — | + | + |

*Identification.* The land use and highway congestion equations can be identified by certain attributes of the housing stock that affect housing prices but have no theoretical reason to affect land use, delays, and unpriced congestion. It appears difficult to identify the housing price equations on purely theoretical grounds because they are likely to be influenced by the same type of variables that influence the land use measures. But, as noted, it is reasonable to expect that some of the metropolitan area characteristics that help explain land use may not affect housing prices and vice versa. In any case, the exclusion restrictions that ultimately identify the housing price equations are based on statistical (in)significance rather than an unambiguous theoretical exclusion requirement.

## 4.4 Sample and Variables

We identified the 100 largest Metropolitan Statistical Areas (MSAs) in the nation and constructed a sample based on the 98 largest MSAs in the 48 contiguous states (excluding Honolulu, Hawaii and San Juan, Puerto Rico) for the year 2000. The MSAs account for a large fraction of U.S. highway congestion and exhibit a wide range of land use patterns.[12]

*Endogenous Variables.* Data on average reported housing prices for owner-occupied households are from the 2000 U.S. Decennial Census. As noted, we divide residents into average commuting time blocks of less than 5 minutes, 5-19 minutes, 20-45 minutes, and greater than 45 minutes.[13] We obtained less satisfactory statistical fits using other commuting time blocks. Roughly 24 percent of owner-occupied households did not have any members who worked outside the home. The Census PUMS data does not allow us to place these households in a commute time group. However, the value of their homes will be affected by changes in highway congestion variables because home prices are determined by demand and supply in the entire metropolitan area housing market. As discussed later, our simulations account for changes in the prices of homes owned by people who do not work outside the home.

To ease the interpretation of our parameter estimates and simulation results, we annualize housing prices by drawing on Himmelberg, Mayer, and Sinai's (2005) estimates of the annual "user cost" of housing. The authors estimate the annual cost of owning a home—that is, the percent of housing that is consumed each year—in cities across the United States accounting for the opportunity cost of owning a home, the change in house prices, and federal income and local property tax rates. We apply the authors' user cost estimates for 46 specific cities to annualize home values for the same cities in our sample and use the average annual user cost in their sample, 5.85 percent, to annualize home values for the remaining cities in our sample.

Table 4.1 presents the current or, more accurately, present value and the annualized value of houses in our sample by commute distance block. Households pay a substantial premium to live close to their workplaces, and the large standard deviation indicates that some homes in this time block are quite expensive. As expected, average housing prices decline as households live further from their workplaces. Note, however, that the averages reported in the table control only for commute time and not other attributes such as lot and house size; thus, the bid-rent function may be steeper than implied by our summary data.

Citywide density, measured as population per unit of land, and entropy, as specified in equation (1), are constructed from U.S. Census data on the popu-

[12]The Texas Transportation Institute limits their annual assessment of congestion in the United States to 85 major MSAs because they find that congestion falls sharply as city size declines. Thus, our omission of MSAs outside of the top 100 should have little effect on our findings. Although some MSAs include rural areas of predominantly urban counties, some residents in rural areas that are not included do commute into the city. Such problems are unavoidable when using MSAs to define cities in a national analysis.

[13]Average commuting times were based on the average commute time of all household members who work outside the home. Commuters who walk to work are placed in the closest time group regardless of the time it takes them to get to work.

Table 4.1: Population Weighted Housing Values in the Sample by Commute Time Block

| | Present Value (thousands of 2000 dollars) | | Annualized Value (2000 dollars) | |
|---|---|---|---|---|
| Commute Time Block | Average | Std Dev | Average | Std Dev |
| Less than 5 minutes | 231.86 | 95.06 | 12,687 | 4,020 |
| 5-19 minutes | 195.72 | 80.47 | 10,731 | 3,425 |
| 20-45 minutes | 188.93 | 75.81 | 10,366 | 3,203 |
| Greater than 45 minutes | 185.76 | 74.32 | 10,195 | 3,163 |

lation and land area of each census tract in the city. We use Census tracts to determine urban subunits in the entropy formula. Census tracts have similar populations because they are defined such that they include roughly 4,000 people living contiguously, although their land area can vary greatly. Las Vegas, Bakersfield, and Tucson have the lowest densities in our sample, Jersey City, New York City, and Orange County (CA) the highest. Harrisburg, Ann Arbor, and Syracuse have the lowest entropy (i.e., density that varies greatly) because they have pockets of densely populated land surrounded by areas with very few people. Fort Lauderdale, Orange County, and San Jose have the highest entropy because they have moderately high density over a broad area. Density and entropy are clearly capturing distinct aspects of land use because their correlation is 0.28.

The highway congestion variables that we include in the model are travel delays and the benefits of unpriced congestion. We measure the delay per mile on highways in the city during the peak travel hours based on the average peak-hour volume-capacity ratio reported in the Federal Highway Administration, *Highway Statistics*. We estimate the difference between actual and free flow speeds using a speed-flow curve developed by the Bureau of Public Roads (a derivation is provided in the appendix).[14] The benefit to road users of unpriced congestion is measured by the difference between the average cost per mile (including the monetary value of travel time) and the marginal social cost per mile (a derivation is provided in the appendix). In our sample, Orange County, CA has the greatest average delay per mile, 41 seconds, and thus the largest benefit of unpriced congestion per mile, 36 cents, while Scranton, PA has the lowest average delay per mile, 0.5 seconds, and smallest benefit of unpriced congestion per mile, 0.34 cents.

We specify delay per mile in minutes instead of multiplying it by an assumed

[14] Following the Texas Transportation Institute, we assume a free-flow speed of 60 miles per hour for urban highways. We obtained similar results using alternatives to the Bureau of Public Roads' speed-flow curve, such as the Metropolitan Transportation Commission's Bay Area speed-flow curve.

value of time and expressing it as a cost because we wish to estimate the implicit value of travel delays based on the effect of delays on housing prices However, unpriced congestion is measured as the difference between social marginal costs and private average costs and must include a value of time. This variable is also used to set the efficient congestion toll faced by all road users. We assume that the value of time is half of the average wage in the city (following Small (1992)) and later discuss how our main findings would change based on alternative assumptions.

Finally, although we use highway delay and unpriced congestion in our model, we are not assuming that all road users travel entirely on freeways. Rather, we are assuming that the variation in highway delay and unpriced highway congestion across MSAs is a good indicator of the variation in delay and unpriced congestion on all major thoroughfares servicing MSAs.

*Exogenous variables.* The housing stock attributes in the annualized housing price equations for each commuting block include the percent of homes with a cellar and the percent of homes with missing tiles or other damage to the roof.[15] Metropolitan area characteristics include the office vacancy rate, average annual household income, mean number of days per year below 32 degrees Fahrenheit (obtained from the National Oceanographic and Atmospheric Administration), and the percent of the state that is classified as urban.[16] We expect that home prices will increase if houses have a cellar and if they are located in cities with affluent residents. We expect home prices to fall if houses have a damaged roof and if they are located in cities that experience a lot of cold weather, have vacant office space, and as indicated by the state's urbanization compete with other cities in the state to attract residents. Finally, we include state-level fixed effects to capture variation in state taxes and government services that may influence

[15]Data for the percent of homes with a cellar and with missing tiles or other damage to the roof are from the American Housing Survey. We assign to these variables the average value in the sample for those cities in which the Census did not conduct a housing survey between 1997 and 2003. Variables such as the number of bedrooms and the number of rooms in the house are likely to be endogenous and their exclusion had little effect on our main findings. We also tried including several additional variables from the American Housing Survey but they were statistically insignificant. The variables were the percent of houses with a "major problem," porch, fireplace, washer and drier, rodents, garage, and structural damage such as a crumbling foundation, cracks in the wall larger than a dime, or sloping walls.

[16]Household income and the percent of the state that is classified as urban are from the Decennial Census. The office vacancy rate is from CB Richard Ellis. Other metropolitan area characteristics we tried to include in the model but found to be statistically insignificant were the cost of living in the city, number of high air pollution days, annual precipitation, and the average number of extremely warm (¿ 90 degrees Fahrenheit) days per year. Variables such as crime rates, foreign born residents, and the like were not included because they are likely to be endogenous.

home prices.[17]

Metropolitan area characteristics that affect the availability of land outside the city to facilitate expansion are appropriate to include in the density equation. Thus we specify a coast dummy to indicate whether development was limited by a large body of water such as an ocean or a gulf. We also specify a dummy to indicate whether the MSA contains a traffic bottleneck (e.g., a bridge over a body of water) that might encourage (discourage) development before (beyond) a major point of congestion. Finally, we include average annual precipitation in the MSA (obtained from the National Oceanographic and Atmospheric Administration) as a proxy for the historical attraction and development of an MSA, which was determined to a certain extent by whether the climate was conducive to local agriculture.[18]

We include the number of municipalities (obtained from the Office of Management and Budget) as a metropolitan area characteristic in the entropy equation. We expect cities with a large number of municipalities to have dispersed populations and greater entropy than cities with fewer municipalities. We also include the coast dummy and a dummy that indicates whether the MSA has an interstate running through it, both of which we expect to increase entropy.[19] And we include the range of elevation in the MSA as a topographical characteristic. We expect a greater range of elevation to decrease the variation in population density—that is, have a negative effect on entropy—because most residents will tend to live in flat areas where it is cheaper to build housing, and to live close to sea level in coastal cities. We measure the range of elevation in the half degree longitude-latitude square including the city.[20]

[17]Taxes and government services tend to vary at the state and municipality level. Because an MSA contains many municipalities and a state may contain more than one MSA, state fixed effects are more appropriate than MSA fixed effects to control for the variation in taxes and government services. For those states that have only one major city, state and MSA fixed effects are equivalent.

[18]In the density equation, we also explored other ways of capturing limits on a city's development including a CMSA dummy to indicate whether the city is part of a larger metropolitan area with less land but it was statistically insignificant. We also tried geographical variables including the mean and standard deviation of a city's elevation and slope, but they were also insignificant.

[19]In the entropy equation, we found that the interstate dummy fit better than the bottleneck dummy we include in the density equation. One possible explanation is that transportation is less expensive along the interstate causing dense areas to be dispersed along the interstate rather than concentrated in one area, which would increase the variation in density across the metropolitan area.

[20]To capture the range of elevation over the area covered by the MSA, not just the central city, we needed a measure of the range of elevation surrounding the city. The best publicly available data for this measure are from the International Satellite Land-Surface Climatology Project (ISLSCP), which provides the range of elevation for the half degree longitude-latitude square including the city. The actual data were from ISLSCP's Initiative II Data Archive.

We include the state-level fixed effects in the density and entropy equations. We include a dummy variable for New York City in the density equation because that city is distinguished from others in the sample by its island geography. We also include a "major city" dummy variable in the entropy equation that denotes cities that are located in a state that has more than one city in our sample and that are the most populous.[21] This dummy captures the likelihood that large cities in a given state share similarities in population entropy that they are not likely to share with other cities in the state.

Delays and unpriced congestion should be influenced by the MSA's population and economic vitality, as indicated by the income of its residents, and by road network characteristics that may exacerbate congestion (e.g., an interstate highway carrying through traffic) and that may limit capacity expansion (e.g., close proximity to a large body of water). We also include the standard deviation of elevation. Although the range of elevation captures whether there are particularly high peaks and low valleys near an MSA, the standard deviation of elevation indicates whether there are many of these peaks and valleys or just a few. A high standard deviation of elevation could call for additional road capacity to enhance vehicle safety, thereby reducing delays and unpriced congestion. Finally, we include the state-level fixed effects and the "major city" dummy variables. Major cities in the same state may share congested-related influences that they do not share with other cities in the state.

## 4.5 Estimation Results

We specified a linear functional form for the annualized housing price, density, entropy, delay, and unpriced congestion equations and jointly estimated the system by three-stage least squares to account for the endogenous influences and contemporaneous correlation of the errors.[22] The estimation results presented in Table 4.2 indicate that most of the variables are estimated with good precision and have the expected sign.

[21] Major cities are New York City, Philadelphia, Boston, Detroit, Dallas, Houston, and Los Angeles.

[22] We also estimated models using log-linear functional forms for all the equations and log-linear functional forms for the housing price and land use equations and semi-log linear functional forms for the highway congestion equations. These functional forms did not fit the data as well as the linear functional forms fit the data.

Table 4.2: Three Stage Least Squares Parameter Estimates

| | Dependent Variables | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Annualized owner-occupied housing value by commute time group (in thousands of 2000 dollars) | | | | Average MSA Population Density | MSA Population Entropy* 1,000 | Delay (minutes per mile) | Un-priced Congestion (dollars per mile) |
| Independent Variables | Less than 5 minutes | 5-19 minutes | 20-45 minutes | Greater than 45 minutes | (1000s of people/mi$^2$) | | | |
| Average Delay (minutes per mile) | -30.38 | -17.58 | -15.79 | -20.75 | 33.41 | – | – | – |
| | (6.11) | (3.66) | (3.22) | (3.55) | (5.74) | | | |
| Highway Congestion Distortion: | 83.83 | 40.56 | 35.69 | 47.04 | -67.31 | – | – | – |
| marginal cost per mile minus average cost per mile (dollars per mile) | (12.20) | (7.24) | (6.36) | (6.99) | (12.94) | | | |
| Average MSA Population Density (thousands of people per square mile) | 1.03 | 0.70 | 0.65 | 0.47 | – | – | – | – |
| | (0.21) | (0.13) | (0.11) | (0.12) | | | | |
| MSA Population Entropy | -3.74 | 10.28 | 8.46 | 11.38 | – | – | -0.53 | -0.45 |
| | (12.46) | (7.29) | (6.25) | (7.06) | | | (0.54) | (0.26) |
| Annualized housing value for owner-occupied working households with commutes less than 5 minutes (thousands of 2000 dollars) | – | – | – | – | -0.98 | -8.31 | – | – |
| | | | | | (0.24) | (2.62) | | |
| Annualized housing value for owner-occupied working households with commutes between 5 and 19 minutes (thousands of 2000 dollars) | – | – | – | – | 0.59 | -7.72 | – | – |
| | | | | | (0.46) | (5.53) | | |

Continued Next Page

Three Stage Least Squares Parameter Estimates (Continued)

| | Dependent Variables | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Annualized owner-occupied housing value by commute time group (in thousands of 2000 dollars) | | | | Average MSA Population Density (1000s of people/mi$^2$) | MSA Population Entropy* 1,000 | Delay (minutes per mile) | Un-priced Congestion (dollars per mile) |
| Independent Variables | Less than 5 minutes | 5-19 minutes | 20-45 minutes | Greater than 45 minutes | | | | |
| Annualized housing value for owner-occupied working households with commutes between 20 and 45 minutes (thousands of 2000 dollars) | – | – | – | – | 1.21 | 13.3 | – | – |
| | | | | | (0.49) | (5.70) | | |
| Annualized housing value for owner-occupied working households with commutes over 45 minutes (thousands of 2000 dollars) | – | – | – | – | -0.10 | 5.24 | – | – |
| | | | | | (0.24) | (2.57) | | |
| Percent of homes in MSA with a cellar | 1.44 | 0.88 | 1.13 | 1.41 | – | – | – | – |
| | (0.55) | (0.39) | (0.34) | (0.37) | | | | |
| Percent of homes in MSA with tiles etc. missing from roof | -44.66 | -27.48 | -29.84 | -25.81 | – | – | – | – |
| | (18.00) | (12.21) | (10.60) | (11.91) | | | | |
| Mean number of days per year with minimum temperature under 32 degrees Fahrenheit | -0.042 | -0.026 | -0.020 | -0.022 | – | – | – | – |
| | (0.009) | (0.006) | (0.005) | (0.006) | | | | |
| Percent of state classified as "urban" | -0.62 | -0.34 | -0.30 | -0.25 | – | – | – | – |
| | (0.11) | (0.07) | (0.06) | (0.07) | | | | |

Continued Next Page

Three Stage Least Squares Parameter Estimates (Continued)

| | Dependent Variables | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Annualized owner-occupied housing value by commute time group (in thousands of 2000 dollars) | | | | Average MSA Population Density | MSA Population Entropy* 1,000 | Delay (minutes per mile) | Un-priced Congestion (dollars per mile) |
| Independent Variables | Less than 5 minutes | 5-19 minutes | 20-45 minutes | Greater than 45 minutes | (1000s of people/mi$^2$) | | | |
| Average annual household income for owner-occupied, working households in each time group (thousands of dollars) | 0.053 | 0.115 | 0.106 | 0.100 | – | – | – | – |
| | (0.012) | (0.009) | (0.007) | (0.006) | | | | |
| MSA office vacancy rate | -0.041 | -0.032 | -0.042 | -0.046 | – | – | – | – |
| | (0.016) | (0.012) | (0.01) | (0.011) | | | | |
| Bottleneck dummy variable (1 if city contains a bottleneck, 0 otherwise) | – | – | – | – | 0.81 | – | – | – |
| | | | | | (0.34) | | | |
| Coast Dummy ( 1 if MSA is located on an ocean or the Gulf of Mexico, 0 otherwise) | – | – | – | – | 1.00 | 17.50 | 0.06 | 0.03 |
| | | | | | (0.40) | (3.80) | (0.02) | -0.007 |
| Average Annual Precipitation in the MSA (hundreds of inches) | – | – | – | – | 4.19 | – | – | – |
| | | | | | (2.53) | | | |
| Interstate Dummy (1 if MSA has an interstate highway ending in 0 or 5, 0 otherwise) | – | – | – | – | – | 5.94 | 0.0379 | 0.0208 |
| | | | | | | (3.17) | (0.0149) | (0.0065) |
| Number of municipalities in the MSA | – | – | – | – | – | 0.23 | – | – |
| | | | | | | (0.04) | | |

Continued Next Page

Three Stage Least Squares Parameter Estimates (Continued)

| | Dependent Variables | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Annualized owner-occupied housing value by commute time group (in thousands of 2000 dollars) | | | | Average MSA Population Density | MSA Population Entropy* 1,000 | Delay (minutes per mile) | Un-priced Congestion (dollars per mile) |
| Independent Variables | Less than 5 minutes | 5-19 minutes | 20-45 minutes | Greater than 45 minutes | (1000s of people/mi$^2$) | | | |
| Range of Elevation within a half-degree longitude-latitude square containing the MSA (feet) | – | – | – | – | – | -0.012 | – | – |
| | | | | | | (0.004) | | |
| Average annual household income for owner-occupied working households in the MSA (thousands of dollars) | – | – | – | – | – | – | 0.0034 | 0.0029 |
| | | | | | | | (0.0006) | (0.0003) |
| Standard Deviation of elevation within a half-degree longitude-latitude square containing the MSA | – | – | – | – | – | – | -1.81 E-4 | -8.30 E-5 |
| | | | | | | | (7.37 E-5) | (3.17 E-5) |
| New York City Dummy Variable | No | No | No | No | Yes | No | No | No |
| Major City Dummy Variables | No | No | No | No | No | Yes | Yes | Yes |
| State Dummy Variables | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| R2 | 0.8 | 0.92 | 0.93 | 0.93 | 0.47 | 0.74 | 0.79 | 0.82 |

*Sample is the 98 of the 100 largest Metropolitan Statistical Areas in 2000 that are located within the continuous 48 states.

*Housing price equations.* The effects of average delay and unpriced congestion on annualized home values are central to our analysis. As noted, these variables capture all aspects of highway travel including commuting, non-work trips, and emergency and non-emergency services. In all likelihood, residents are likely to place the greatest weight on changes in highway travel that affect their commute.

Interestingly, we find that the pattern of coefficients exhibits a U-shape (in absolute value) as the prices of homes where residents have the shortest and longest commuting times are more responsive to average delay and unpriced congestion than are the prices of homes occupied by residents who have commuting times between these extremes.[23] We expect households who have self-selected into residential locations where they are close to work to have a high value of travel time and attach a large disutility to delay. We also expect households who have self-selected into residential locations where they have a long commute to have much lower values of travel time. But given that average delay is measured on a per-mile basis, we find that such households attach a high disutility to an increase in average delay in the metropolitan area because it will have the greatest cumulative effect on the duration of their commute and on other vehicle travel. Given their disutility from the cost of delay, households who have the shortest and longest commutes experience the largest decrease in home values for a given increase in delay, and thus have the highest willingness to pay to reduce delay. These households also experience the largest increase in home values from a given increase in the benefit of not having to pay out-of-pocket for the congestion they cause. Households who live close to work avoid being charged for delaying a large flow of motorists, some of whom have a high value of time, while households who live far from work avoid being charged for contributing to congestion during a lengthy commute.

Additional perspective on the delay coefficients is typically provided in mode choice studies by using them to calculate the implied value of travel time. However, it is difficult for us to do that here for two reasons. First, as noted, the coefficients reflect delay that the household experiences from its work and non-work trips and from highway travel by service providers who affect the household's welfare. We do not know the proportion of delay costs that are accounted for by each trip purpose. Second, the coefficients capture the effect of delay on the *household* not on a single commuter, and it is not clear how to properly apportion the costs of delay to various household members.

The coefficients for unpriced congestion indicate the total annual miles sub-

[23]We tried interacting the highway congestion variables with indicators of a city's population to see if their effects varied for large and small cites and with indicators of high and low volume-capacity ratios to see if their effects varied by traffic density. We did not find any notable changes in the estimates.

ject to some congestion that are experienced directly and indirectly by households who live in a given time block.[24] As expected, these significantly exceed a household's annual vehicle miles traveled, especially for households who live close to employment centers and tend to "consume" a disproportionate share of the city's amenities.[25]

Turning to the other coefficients, the effect of MSA density on annualized housing prices becomes stronger as homes are located closer to employment centers. This finding appears to reflect preference heterogeneity: residents who choose to enjoy the benefits of living close to employment centers are also more likely to place a relatively higher value on these benefits and a lower value on the negative aspects of density such as crime and noise. Residents who live farther away from employment centers may place a somewhat lower value than other residents place on the positive aspects of density and a higher value on the negative aspects. The effect of entropy on annualized home prices is highly insignificant for the closest time block and has a positive effect that does not vary much across the other time blocks. The closest time block is likely to be densely populated in most MSAs, thus home values are likely to be insensitive to changes in entropy over the entire MSA. The positive effect for the remaining time blocks reflects those residents' preferences for accessibility to areas with the average density. Although the a priori effect of entropy involves a trade-off between accessibility and accommodating preference heterogeneity, our empirical findings indicate that residents place a higher value on greater accessibility to the attractions of higher density neighborhoods.

The magnitudes of the positive coefficient for a cellar and the negative coefficient for a damaged roof do not vary greatly across commute time blocks. But the negative effect on housing prices of cold weather is noticeably greater for the shortest commuting time block than for other commuting time blocks, possibly because more people who live in the shortest commuting time block may walk to work and to perform errands and therefore experience greater disutility than

[24]Note that the units of annualized housing costs are dollars and units of unpriced congestion are dollars per mile.

[25]If we ignore the complications of trying to apportion the costs of delay to various household members, we can obtain a rough estimate of the value of travel time by dividing the delay coefficients by the total annual miles for which the household is exposed to some congestion (note that annualized housing costs are expressed in dollars and delay is expressed in minutes per mile). This procedure yields values of time that cluster around $20 per hour, which are aligned with estimates for households in Los Angeles (Bajari and Kahn (2000) and Small, Winston, and Yan (2006)) but may be above the value that might be expected for the nation as a whole based on work trips. On the other hand, our estimate is based on working-households who own a home and have incomes that generally exceed the incomes of commuters in general. Average income in our sample is just under $40 per hour, which would imply a value of time that is roughly equal to half of the wage.

other residents experience from cold weather. We also find the negative effect of competition from other urban areas in the state is somewhat greater for the shortest commuting time block while the positive effect of income is smaller.[26] An increase in office vacancy rates reduces home prices fairly uniformly across commuting time blocks.

*Land use equations.* We find, as expected, that average delay increases density and that the failure to price congestion decreases density. (The highway congestion variables had insignificant effects on entropy.) *Thus, holding home prices constant, the failure to price highway congestion contributes to sprawl, while the existence of the resulting delays reduces sprawl.* And given that the increase in density from pricing congestion is greater than the expected decrease in density from reducing delay, inefficient highway policy results in an increase in sprawl.[27]

Annualized home prices have varying effects on density and entropy that are plausible. An increase in home prices for the shortest commuting time block reduces density and entropy as residents move to time blocks that are further away from employment centers. An increase in home prices for the middle two commute blocks increases density as residents move to time blocks that are closer to employment centers. In the case of the 5 to 19 minutes time block, this reduces entropy but in the case of the 20 to 45 minute time block entropy increases because areas of the city near the average density gain the most population. Finally, an increase in home prices in the farthest commuting block has an insignificant effect on average MSA density but a positive effect on entropy as households move to a time block whose density is closer to the MSA's average.

Density increases if an MSA has a bottleneck, is located along a coast, and experiences a relatively high level of precipitation. Entropy increases if an MSA is located along a coast, has an interstate running through it, and has a large number of municipalities, while it decreases if its elevation varies greatly (i.e., the MSA is in a area of large, flat places with mountains or valleys) because the bulk of the residents will live in flat places.

The highway congestion variables are explained by parsimonious specifica-

[26] As noted in our conceptual framework, we do not account for the long run response of households to move to new MSAs in response to road pricing. Average household income is therefore taken as exogenous here, although we acknowledge that it could be correlated with unobserved attributes of the city such as cultural offerings that may attract certain residents and affect home prices. However, our simulations, which keep the estimated parameters and value of income fixed during predictions of the base case without road pricing and predictions that capture the economic effects of road pricing will net out any possible bias from the correlation.

[27] At the average levels of the congestion distortion and delay in our sample, (marginal cost) congestion pricing would increase density over 800 people per square mile. To offset this effect, delays would have to decrease more than 70 percent, which seems unlikely given the magnitude of optimal congestion tolls.

tions that include the same influences. Delays and the cost of unpriced congestion are greater if an MSA is located along a coast, has an interstate running through it, and if its residents earn relatively high incomes.[28] An interesting finding is that the congestion variables were more affected by population entropy than by average population density. Specifically, delays and the costs of unpriced congestion are higher as density becomes more varied throughout the MSA (i.e., as entropy decreases). Policymakers may be better able to increase highway capacity to match average density over the metropolitan area than to adjust capacity to ameliorate congestion related to extremely high density in certain parts of the MSA. Finally, a high standard deviation of elevation (i.e., the MSA consists of a series of hills and valleys) reduces delays and the costs of unpriced congestion, possibly because additional road capacity is likely to be constructed to accommodate slower traffic and allow vehicles to pass each other safely.

We have acknowledged that our model is identified by exclusion restrictions. But we can report that the key parameter estimates capturing the effect of the transportation and land use variables on housing prices tended to be robust to alternative specifications that were estimated to determine statistically significant influences on the endogenous variables in our model. In the final specification, certain metropolitan area characteristics identified the housing price equations including a bottleneck dummy, coast dummy, interstate dummy, annual precipitation, and number of municipalities. Given our aggregate approach and the fact that these variables are determined by an MSA's geography or history, they are plausible instruments.

## 4.6 Simulating the Effects of Road Pricing

The central finding of our model is that policymakers' failure to charge motorists for the congestion they cause has raised all home prices in metropolitan areas, while also contributing to sprawl. We use the model to simulate the welfare effects of instituting marginal cost congestion tolls on the nation's urban highways to capture two major effects. The first is that the tolls will generate toll revenues while causing home prices (and property tax revenues) to decline because, on average, residents' higher out-of-pocket highway costs will exceed their value of the reduction in travel time.[29] The increase in toll revenues is a

[28] As indicated in footnote 24, we treat household income as exogenous.

[29] Glazer and Van Dender (2002) develop a theoretical model that predicts that introducing congestion tolls without allowing residents to relocate will reduce property values. Home values could increase if the government uses the toll revenues to increase government services or reduce property taxes. Our model includes state fixed effects to control for the variation in government services. These effects do not change in the simulation. Thus, the initial change in home prices caused by congestion tolls is unambiguously attributable to the change in out-of-pocket costs

welfare gain, assuming the revenues are used for socially desirable purposes, but the decline in home prices represents a welfare loss. The reason is because, on net, congestion pricing reduces the attractiveness of homes and lowers their price by decreasing consumer demand. If the price of housing dropped because the supply of homes increased, the price decline would be associated with a welfare gain from an increase in the housing stock.

The second effect is that congestion pricing will cause certain residents to move, thereby increasing metropolitan area density and partially offsetting the initial reduction in home prices as home prices rise in response to the decreased cost of city services and to residents' greater access to urban amenities. At the same time, because the prices of homes in the shortest commute time groups will fall more than the prices of homes that are farther from employment centers will fall, entropy will increase, which feeds back to further increase home prices, especially those in the farthest time groups. We point out that our model does not capture other benefits of congestion pricing associated with reducing sprawl, such as preserving the natural habitat, discouraging wasteful suburban expansion of rail transit, and weakening restrictive land use regulations. We discuss those and other effects later.

We make the following assumptions to perform a base case simulation and then conduct sensitivity analysis or discuss their likely effects in the conclusion. First, tolls will reduce the average delay per mile in accordance with motorists' long-run elasticity of vehicle miles traveled with respect to commuting costs. We assume that this elasticity is -0.3 and that the pre-toll private cost of driving is \$0.40 per mile.[30] Although optimal congestion tolls vary both across different highways and by time of day, our aggregate approach enables us to calculate one congestion toll per city. Thus, the congestion toll calculations can be interpreted as representing the average congestion toll paid by a city's motorists, and our results reflect the toll's average effect on home prices in each commute time block and on overall urban land use. Second, zoning regulations and physical constraints on the provision of housing are likely to limit the extent that density can increase in response to tolls. We assume that a city's density can increase no more than 50 percent or exceed the density of the Chicago, IL Metropolitan Area, whichever is greatest.[31] The Chicago Metro Area is a mix of a high density center and very low density suburbs, making it an attainable density limit. At

---

and travel time.

[30] The elasticity is consistent with Mannering and Winston's (1985) estimate of the long-run elasticity of vehicle utilization with respect to operating costs. The private cost of driving, including gasoline and vehicle capital costs, is slightly above the IRS tax deduction for driving a personal vehicle for business use.

[31] Ten percent of cities in our sample have an initial density that is greater than the density of Chicago.

the same time, we do not think it is likely and therefore do not want small, sparsely populated cities to become denser than New York City or Northern New Jersey (the two densest MSAs in the sample). In fact, in our analysis, New York City and Jersey City are the only MSAs that are able to increase their density beyond the level that is currently observed in our sample. Third, our simulation assumes that residents do not change their place of work. Finally, our simulation assumes that residents do not move to another metropolitan area—in terms of the monocentric city model, we are employing a "closed-city" instead of an "open-city" model.

The welfare effects of congestion tolls are obtained by calculating an iterated equilibrium. First, we determine the optimal congestion toll in each MSA in our sample. To do so, we calculate the current full cost of driving, which includes private costs of $0.40 per mile and average travel time costs (see the appendix). We then calculate the marginal cost congestion toll based on current VMT (and road capacity) and recalculate VMT given the introduction of the toll and an elasticity of -0.3. The process is repeated until the change in VMT and in the congestion toll is very small. At which point we obtain the toll revenues for each MSA by multiplying the optimal toll per mile times the exposure of each household to congested vehicle miles.[32]

We then calculate the changes in home values and the land use variables by predicting home values in each time group assuming the optimal congestion toll is implemented but keeping the land use variables at their current values. The predicted home values and the optimal congestion toll are then used to predict new values of the land use variables. The process is repeated until the change in the predicted home prices and land use variables is small, at which point we calculate for each MSA the change in annualized home prices, density, entropy, annual toll revenues, and annual property tax revenues.[33]

Table 4.3 presents the (predicted) average changes in annualized housing values, density, and entropy that result from efficient road pricing. Annualized housing values decline for homes in all commute blocks, especially for those closest to and farthest from places of employment. The average decline in housing values over all MSAs in the sample is 18 percent.

It is useful to recall our previous discussion of a household's constrained utility maximization problem to understand our findings. This problem can be

[32]As noted, the exposure of households by commute distance block to congested *passenger* miles can be obtained from the coefficients capturing the effect of the highway congestion distortion on home prices. Dividing this value by average vehicle occupancy, 1.25 people (following the Texas Transportation Institute), generates the household's congested *vehicle* miles.

[33]Average implied property tax rates for each city are generated by using the self-reported property tax payments in the Census Public Use Microdata. These rates are applied to the pre and post simulation home values to generate expected property tax revenues.

Table 4.3: Population Weighted Predicted Average Annualized Value of Housing and Land Use Variables Before and After Congestion Pricing

| | Before Congestion Pricing | After Congestion Pricing |
|---|---|---|
| Annualized Housing Value (2000 dollars per year) | | |
| <5 minutes | 12,636 | 7,515 |
| 5-19 minutes | 10,697 | 8,878 |
| 20-45 minutes | 10,352 | 8,790 |
| >45 minutes | 10,169 | 7,668 |
| Average Annualized Housing Value in the MSA (2000 dollars per year) | 10,545 | 8,656 |
| Average Population Density (People per square mile) | 1552 | 2527 |
| Population Density Entropy | 0.929 | 0.945 |

reformulated as a household maximizing an indirect utility function, where the relevant "prices" for our discussion are the price per unit of housing, travel time costs, and the out of pocket costs of driving including the toll. Note that the price per unit of housing is negatively related to the commute distance from the workplace and that the travel time and out of pocket costs are positively related to the commute distance.

Congestion pricing reduces the relative attractiveness of homes that require less than a 5 minute commute because some commuters who tend to have the highest values of time choose to live somewhat farther from their workplace, consume more house per dollar, and experience little, if any, increase in their commute time without having to pay excessive out-of-pocket costs. Recall from table 1 that these households have been paying a premium in housing prices to live extraordinarily close to work—and in all likelihood for the benefits of a dense living environment—and they no longer have to pay this premium under road pricing. That is, road pricing has *lowered* the rate at which their (total) commute costs rise with increasing distance from their workplaces enabling them to seek more distant and relatively less expensive residential locations.

Residents who live far from their workplace are able to get to work faster but the relative attractiveness of homes that require more than a 45 minute commute is significantly reduced for some of these residents because they have to pay a very high toll per trip and they tend to place a lower value than other residents place on the travel time savings. Thus road pricing has *increased* the rate at which their (total) commute costs rise with increasing distance from their workplaces causing them to seek closer residential locations even if they are more expensive. Note, however, that it is unlikely that these households are now willing to pay a housing premium to live in very dense neighborhoods. Given these adjustments, it is plausible that congestion pricing would cause the values of homes located in the middle two commuting time blocks to decrease by smaller amounts than the values of homes located in the extreme commuting blocks decrease, and that the value of these homes exceeds the value of homes located in the extreme commute time blocks.

To put the finding somewhat differently, households who used to experience higher housing costs to live extremely close to work are now able to live in a neighborhood with less expensive housing (relative to the cost of their original home) that is farther from work but that suits their tastes for urban land use. The movement of such households from the shortest commute time block is not offset by the movement of households who do not like paying higher out of pocket costs but who also are not willing to pay a premium for living in a dense neighborhood close to an employment center.

Our findings contrast with conventional theory based on the monocentric model, which predicts that congestion tolls will cause the bid-rent function to become steeper throughout the metropolitan area (Segal and Steinmeier (1980)). The reason that this does not occur here is because we allow for *preference heterogeneity* whereby residents with high values of time sort into housing close to employment centers and residents with low values of time choose to live farther from work. Preference heterogeneity implies that residents in different commuting time blocks will react differently to the introduction of congestion pricing.

The dramatic change in the structure of home prices caused by road pricing is associated with a significant change in land use. Table 3 reports that cities would become much denser and that density is somewhat more uniformly distributed across the urban area. The former result is consistent with conventional theory. Overall, road pricing has encouraged residents to move closer to places of employment including city sub-centers. In terms of Figure 4.2, the change in land use has reduced the low-density, leap-frog development that characterizes urban sprawl by transforming low density-high entropy functions to high density-high entropy functions. In the process, residents' losses from road pricing because of lower housing values are mitigated to a certain extent by changes in land use that reduce the cost of sprawl.

Before presenting our empirical estimates of the welfare effects of road pricing that account for the changes in land use, it is useful to provide some perspective on why, in theory, the effects differ from previous analyses of road pricing that do not account for changes in land use. Figure 4.3 reproduces the conventional diagram presented in Lindsey (2006). In this framework, the average cost, $AC$,

Figure 4.3: The Conventional Diagram

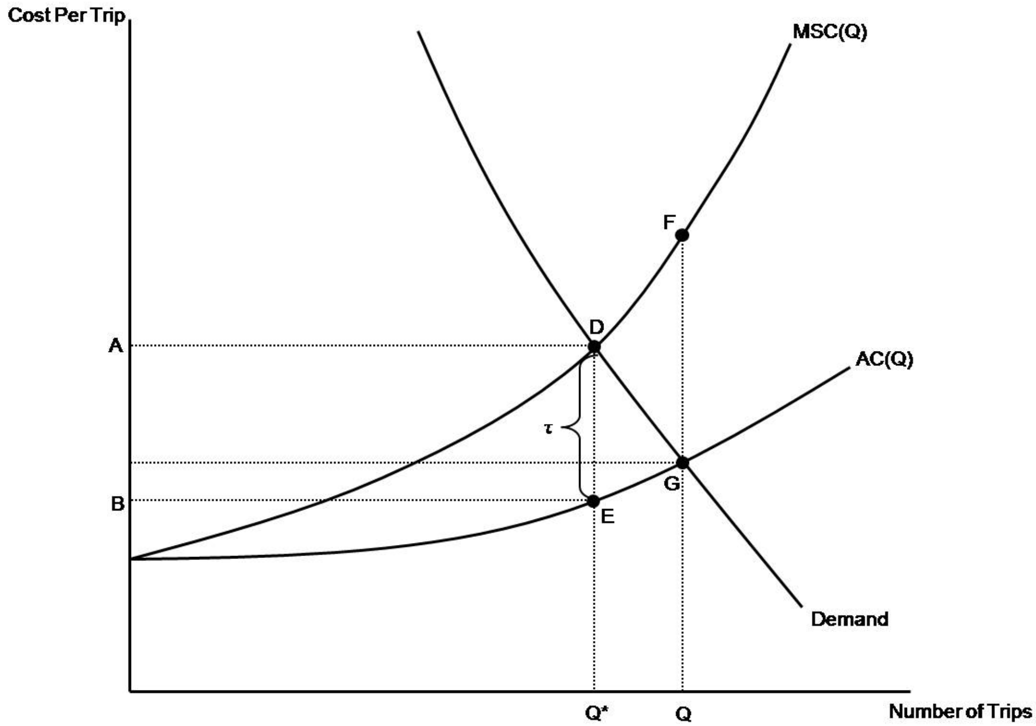


of drivers' trips is less than the marginal social cost, $MSC$, of these trips because drivers do not pay for their contribution to congestion and delays. An optimal congestion toll, $\tau$, reduces travel from its inefficient level, $Q$, to the optimal level $Q*$. In response to the toll, some motorists no longer use the road during peak periods while others continue to use the road and pay the toll. The loss to both groups is given by the area ADGEB. The toll raises revenue equal to ADEB. The toll also reduces but does not eliminate the social cost of delay. This gain is given by the area DFGE. Comparing the areas yields a welfare gain of FDG. That gain is also obtained as the integral of the marginal social cost minus the marginal social benefit of those trips that would have been taken without congestion pricing but are not taken after the toll is imposed. Note that the direct loss in consumer surplus is relatively large and exceeds the gain in revenues, and that the welfare gain is small relative to the loss in consumer surplus and the gain in revenues.

An important effect of accounting for changes in land use is that motorists' demand for travel will become more elastic because households have an additional response—namely, changing their residence—to the introduction of the toll. This is shown in Figure 4.4 by modifying the conventional diagram to include a more elastic demand curve that intersects the average cost curve at point G and yields the pre-toll equilibrium number of trips $Q$. The introduction of a congestion toll

Figure 4.4: The Conventional Diagram Accounting for Land Use

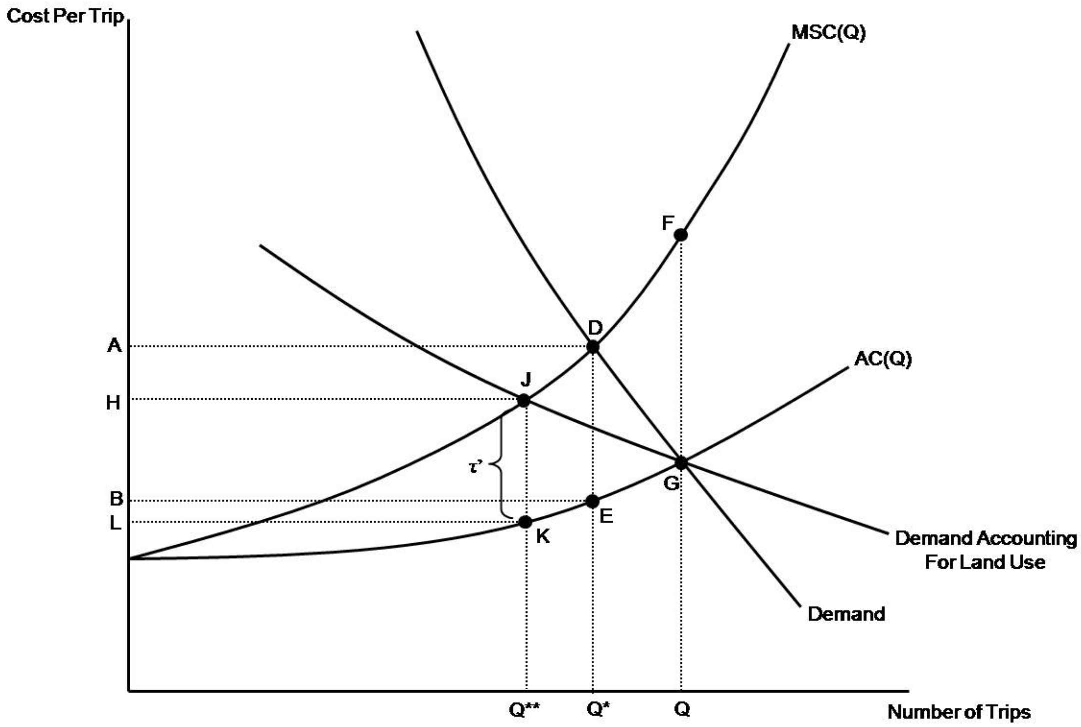


now generates a welfare gain, FGJ, which exceeds the original welfare gain and results in less redistribution. Note that consumer surplus losses, HJGKL, and revenue gains, HJKL, are smaller than they were without land use adjustments, and that both the optimal toll $\tau'$, and the optimal quantity of trips, $Q^{**}$, are smaller than the optimal toll and quantity of trips when housing locations are held fixed. However, the modified graph accounts only for changes in transportation costs and does not include the benefits from more efficient land use, which will further increase the welfare gain of the toll and reduce adverse redistribution.

The welfare effects of the changes induced by road pricing on residents and the government are shown in Table 4.4. At the national level, the annualized value of owner-occupied housing declines \$56.6 billion (2000 dollars)). Because we do not know the value of renter-occupied housing, we assume that the annualized value of renter-occupied housing is equal to twelve times the monthly rent as reported in the U.S. Census. We then assume that the value of housing occupied by working renters declines at the same percentage as the value of housing occupied by working owners in the same commute time group. Thus, we find that the value of renter-occupied housing declines \$18.6 billion (2000 dollars). Finally, as noted, we do not know where non-working (i.e., non-commuting) households live. But given that these households do not make peak-hour commutes, it is likely

Table 4.4: Annual Effect of Congestion Tolling (Billions of 2000 dollars)

| | Change |
|---|---|
| **Owner-Occupied, Working households** | |
| Decrease in Annualized Housing Value: | 56.6 |
| Decrease in Tax Revenue: | 0.7 |
| **Renter-Occupied, Working Households** | |
| Decrease in Annualized Housing Value: | 18.6 |
| Decrease in Tax Revenue: | 0.2 |
| **All Non-working Households:** | |
| Decrease in Annualized Housing Value: | 3.5 |
| Decrease in Tax Revenue: | 0.04 |
| **Total Costs** | 79.6 |
| **Owner-Occupied, Working Households** | |
| Toll Payments | 63.6 |
| **Renter-Occupied, Working Households** | |
| Toll Payments | 40.2 |
| **All Non-working Households** | |
| Toll Payments | 16.6 |
| **Total Toll Revenues** | 120.4 |
| **Annual Social Benefits** | 40.8 |

that compared with working households they select residential locations whose housing values are less affected by congestion tolling. So, for example, they tend to live in the middle two commute blocks instead of in the closest block. We therefore assume that annualized housing values for non-working households decrease by half the percentage decrease of the housing values of working households, which results in an additional $3.5 billion loss in the annualized value of the housing stock.[34]

Thus, the reduction in the annual value of the entire housing stock from the introduction of road pricing is $78.7 billion. The loss in property tax revenue associated with the decline in the value of the housing stock is $0.9 billion, which raises the annual cost from road pricing to $79.6 billion. At the same time, toll revenues amount to $120.4 billion, indicating that implementing congestion

[34] As an upper bound, if we assume that these homes lose the same percentage as working households' homes lose, the loss to non-working households is $7 billion. The additional loss does not alter our basic findings very much.

pricing in the nation's congested metropolitan areas would yield an annual welfare gain of $40.8 billion (2000 dollars).[35]

Consistent with the preceding theoretical discussion, our estimate of the net benefits of adopting road pricing nationwide is much greater than previous estimates that do not account for changes in land use. For example, based on Lee's (1982) estimates, Small, Winston, and Evans (1989) report a nationwide policy of congestion pricing would yield annual revenues of $54 billion (1981 dollars) and, accounting for road users' out-of-pocket losses and travel time savings, an annual welfare gain of $6 billion (1981 dollars). Winston and Shirley (1998) estimate that the annual welfare gain from nationwide congestion pricing amounts to $3.2 billion (1990 dollars), but their estimate includes the losses from additional transit subsidies generated by auto users who shift to bus or rail. (Our analysis does not account for how road pricing may affect transit finances.) As indicated in Figure 4.3, the welfare gains from congestion pricing in these studies are small relative to the redistributive effects.

To be sure, our welfare gain exceeds previous estimates partly because of the growth in traffic congestion during the past two decades. But as indicated by figure 3b and the discussion of additional benefits, our welfare gain is also greater because our model allows residents to change their residential locations, which increases the net benefits of road pricing. Although the benefits from road pricing still entail considerable redistribution, they are so large that the government could retain a sizable amount of toll revenues to maintain and, where appropriate on cost-benefit grounds, to expand the road system, and use part of the revenues to offset residents' losses by, for example, reducing property taxes and/or supplementing reduced tax revenues to the metropolitan area. Thus, in addition to reducing congestion, policymakers would have a stable long-term source of funding to prevent the nation's road system from deteriorating.

As noted, our simulations are based on certain assumptions. Table 4.5 shows the sensitivity of our findings to various assumptions about the elasticity of vehicle miles traveled with respect to congestion tolls and limits on the increase in density in our sample. We allow the elasticity of vehicle miles traveled to range from -0.1 to -0.5 and constrain density of each city to be no greater than the

[35] Some plausibility checks on our estimates are as follows. Owner-occupied, working households in our sample account for over $300 billion in annualized housing value and $56.6 billion in losses—or an 18.8 percent drop in annualized housing value—from congestion pricing. Turning to toll revenues, the cities in our sample have nearly 3 trillion annual vehicle miles traveled on freeways. Our analysis indicates that owner-occupied housing accounts for slightly more than half of total toll revenues and pays congestion tolls on approximately 1.2 trillion vehicle miles annually. Therefore, all housing pays congestion tolls on 2.4 trillion vehicle miles annually or 80 percent of freeway vehicle miles. Alternatively, the $120 billion in toll revenues that we estimate could account for less than 80 percent for freeway vehicle miles and some percentage of congested arterial miles.

80th percentile in the sample (the density of Milwaukee, WI), or the 95th percentile (the density of Los Angeles, CA), or 1.5 times the city's current density, whichever is greater. Even under the least favorable assumptions, congestion pricing generates an annual welfare gain of $28.0 billion.[36] Finally, it should be kept in mind that we have estimated an average congestion toll for each city that is used in the simulations. If we allowed congestion tolls to vary both by roadway and time of day, commuters would be charged more precisely for the marginal cost of their trips, which would increase the annual welfare gain from congestion pricing.

[36] An additional area of sensitivity is the assumed value of time used to measure the unpriced congestion variable and the optimal tolls (and thus the toll revenues). We assumed an average value of time equal to half of the average wage. If we assumed a higher value, the net welfare gains would be larger. This direction would be justified because we do not account for improvements in travel time reliability that would result from road pricing (Small, Winston, and Yan (2006)). Of course, there is a *distribution* of the value of travel time that is likely to reveal differences in behavior within a given commute time group. However, the preceding conclusion based on the average value of time would still hold if our analysis were based on such a distribution.

Table 4.5: Results of Sensitivity Analysis* (billions of 2000 dollars)

| VMT Elasticity | **-0.3** | -0.3 | -0.3 | -0.5 | -0.5 | -0.5 | -0.1 | -0.1 | -0.1 |
|---|---|---|---|---|---|---|---|---|---|
| Density Limit (percentile) | **0.9** | 0.8 | 0.95 | 0.9 | 0.8 | 0.95 | 0.9 | 0.8 | 0.95 |
| Billions of 2000 Dollars | | | | | | | | | |
| **Owner-Occupied** | | | | | | | | | |
| Decrease in Annualized Home Value | **56.6** | 63.1 | 47.7 | 47.4 | 53.9 | 38.5 | 70.0 | 76.5 | 61.0 |
| Decrease in Tax Revenue | **0.7** | 0.8 | 0.6 | 0.6 | 0.7 | 0.5 | 0.8 | 0.9 | 0.7 |
| **Renter-Occupied** | | | | | | | | | |
| Decrease in Annualized Home Value | **18.6** | 20.1 | 16.6 | 14.9 | 16.4 | 12.9 | 24.1 | 25.7 | 22.1 |
| Decrease in Tax Revenue | **0.2** | 0.2 | 0.2 | 0.2 | 0.2 | 0.1 | 0.3 | 0.3 | 0.2 |
| **Non-workers** | | | | | | | | | |
| Decrease in Annualized Home value | **3.5** | 3.9 | 3.0 | 2.8 | 3.2 | 2.3 | 4.5 | 4.8 | 3.9 |
| Decrease in Tax Revenue | **0.04** | 0.04 | 0.03 | 0.03 | 0.04 | 0.03 | 0.05 | 0.05 | 0.04 |
| **Annual Decrease** | | | | | | | | | |
| Tax Revenue | **0.94** | 1.04 | 0.83 | 0.83 | 0.94 | 0.63 | 1.15 | 1.25 | 0.94 |
| Home Value | **78.7** | 87.1 | 67.3 | 65.1 | 73.5 | 53.7 | 98.6 | 107.0 | 87.0 |
| Total Costs | **79.6** | 88.1 | 68.1 | 65.9 | 74.4 | 54.3 | 99.8 | 108.3 | 87.9 |
| Annual Toll Revenue | **120.4** | 120.4 | 120.4 | 102.4 | 102.4 | 102.4 | 146.2 | 146.2 | 146.2 |
| Annual Social Benefit | **40.8** | 32.3 | 52.3 | 36.5 | 28.0 | 48.1 | 46.4 | 37.9 | 58.3 |

Base case results are in **bold.** Minor discrepancies are due to rounding.

## 4.7 Qualifications and Discussion

Many authors who have written about road pricing have asserted that it may have important effects on land use. Almost independently, a literature has recently developed suggesting—but rarely quantifying—that sprawl causes significant social costs. Surprisingly no one, to the best of our knowledge, has attempted to analyze empirically how road pricing may lower the costs of sprawl by improving land use.

We have applied a methodology to account for the effects of road pricing on land use and found that road pricing's net benefits are substantial, in part because of improvements in land use. Thus the government obtains an efficiently generated source of funding for the road system and is better able to address distributional concerns that have long been identified as a political obstacle to adopting road pricing. Note that certain renters would directly gain from the policy because lower housing prices would be reflected in lower rents. In the process, more affordable housing would be available for renters, whose incomes tend to be lower than homeowners' incomes.

Our findings should be qualified because certain households would incur (lump sum) transactions costs from selling their homes and moving into existing or newly constructed housing.[37] At the same time, our findings are understated because although we assume in our simulations that households are able to move to locations within their own MSA, we assume they are not able to take a different job in the MSA or relocate to a different MSA. Accounting for job and intercity mobility would enable households to optimize their response to road pricing even further. From a distributional perspective, road pricing is likely to cause congested cities to gain population and increase property values at the expense of less congested cities. But as pointed out in Winston and Langer (2006), the growth in delays during the past twenty years in the United States is to a notable extent accounted for by cities that experienced little congestion in the early 1980s but now experience measurable congestion. Hence, many small, low-density areas are likely to benefit in the future from adopting road pricing.

We also do not explicitly account for how the private sector would respond to and be affected by road pricing.[38] Our revenue estimates do include higher out-of-pocket costs incurred by truckers, which are largely passed on to consumers. On the other hand, these losses are likely to be offset because by reducing delays congestion pricing reduces truckers' operating costs and enables firms to hold

[37]The transactions costs are transfers to realtors, movers of household goods, and the construction industry.

[38]An issue that has arisen in London's road pricing experiment is its effect on retail sales. Quddus, Carmel, and Bell (2007) found that the congestion charge raised the sales for a specific store located in the priced zone but that it did not affect overall retail sales in central London.

lower inventories (Shirley and Winston (2004))—savings that in large part are passed on to consumers.

Finally, we have indicated that the government could use the toll revenues to soften the distributional effects of road pricing. In addition, if part of the revenues were used to finance efficient infrastructure investments (Deakin (1994)), such as expanding highway capacity in a dense corridor, another round of land use adjustments would result and produce additional welfare gains. We must also acknowledge that like a notable share of recent transportation expenditures, the additional revenues generated by tolls could also result in wasteful spending (Winston (2006)).

In the final analysis, we appear to obtain plausible benefits from improvements in land use caused by road pricing. For example, Burchell and others (1998) concluded that compared with "sprawling" development, "compact" development roadway infrastructure costs are 25 percent lower, utility costs are 20 percent lower, and school infrastructure costs are 5 percent lower. Such figures suggest that the annual cost savings from reduced infrastructure costs that are internalized into home prices could easily amount to $15 billion.

We also point out that by improving land use, road pricing may produce additional social benefits that we have not been able to quantify. First, increasing density and decreasing entropy could promote social interactions and strengthen the bonds that underpin a healthy society (Bruckner 2000)). In particular, changes in land use could reduce the distance between poor and affluent residents and make it harder for the wealthy to ignore the problems of those less well off (Kahn (2006)). Recall, that in response to congestion pricing we found that households who live in the most expensive homes in an MSA (and in many cases, who have the highest incomes) move from homes in the central city and sub-centers that are basically within walking distance of work to neighborhoods that are between 5 and 45 minutes from work, while households who live more than 45 minutes from work move to neighborhoods that are less than 45 minutes from work. Overall, the city will become denser, indicating that poor and rich people will be living closer together.

Second, reducing sprawl could enhance the protection of natural habitat at the urban boundary. Third, increasing density could encourage the use of vehicles that are more fuel efficient and that emit fewer emissions than vehicles used in less dense metropolitan environments (Fang (2006)). Finally, reducing congestion and sprawl weakens the ostensible rationale for policymakers to use inefficient policies to address these problems, such as zoning laws (Glaeser and Gyourko (2003)), urban growth boundaries (Anas and Rhee (2006), Brueckner (2007)), transit oriented development (Winston and Maheshri (2007)), and various taxes and fees that are intended to raise money for transportation improvements.[39] Similarly, with the recent interest in reducing carbon emissions in the United States, congestion pricing would reduce vehicle miles traveled for most households, thus decreasing the nation's vehicle emissions of all pollutants. By efficiently raising the price, on average, of urban travel, policymakers could potentially reduce the size of any future carbon taxes that might inefficiently seek to tax travel instead of taxing carbon emissions directly. Hopefully, policymakers would be less inclined to pursue these approaches.

If subsequent work confirms that road pricing's appeal extends far beyond congestion mitigation, then it would appear that the policy community has substantially underestimated pricing's social benefits and may have exaggerated its undesirable distributional features. Eventually, opposition to road pricing may wear thin.

[39]For example, Virginia lawmakers recently gave the Northern Virginia Transportation Authority the power to impose new taxes and increase existing vehicle registration and safety inspection fees with the expectation of raising some $325 million for roads and transit. More than half of the money will be accounted for by a "congestion relief fee," which is actually a real estate seller's tax of 40 cents per $100 of assessed valuation on the sale price of a house. The constitutionality of the new levies is currently being challenged before the Virginia Supreme Court.

# Bibliography


[1] Aizcorbe, Ana, Benjamin Bridgman, and Jeremy Nalewaik 2009 "Heterogeneous car buyers: a stylized fact" *Finance and Economics Discussion Series* 2009-12, Board of Governors of the Federal Reserve System

[2] Akerlof, George A., and Rachel E. Kranton 2000 "Economics and identity" Quarterly Journal of Economics 115(3): 715-753

[3] Alonso, William (1964), Location and Land Use: Toward A General Theory of Land Rent, Harvard University Press, Cambridge, Massachusetts.

[4] Altonji, Joseph G. and Charles R. Pierret 2001 "Employer learning and statistical discrimination" *Quarterly Journal of Economics* CXVI: 293-312

[5] Alquist, Ron and Lutz Kilian. 2008. "What Do We Learn from the Price of Crude Oil Futures? CEPR Discussion Paper 6548.

[6] Anas, Alex and Hyok-Joo Rhee (2006), "Curbing Excess Sprawl with Congestion Tolls and Urban Boundaries," Regional Science and Urban Economics, vol. 36, July, pp. 510-541.

[7] Arrow, Kenneth 1973 "The theory of discrimination" in Discrimination in Labor Markets, O. Ashenfelter and A. Rees, eds (Princeton, NJ: Princeton University Press)

[8] Ayres, Ian. 1991 "Fair Driving: Gender and Race Discrimination in Retail Car Negotiations" 104 *Harvard Law Review* 817.

[9] Ayres, Ian and Peter Siegelman 1995 "Race and Gender Discrimination in Bargaining for a New Car" *American Economic Review* 85(3): 304-21

[10] Azar, Omar 2003 "Can price discrimination be bad for firms and good for all consumers? A theoretical analysis of cross-market price constraints with entry and price discrimination" *Topics in Economic Analysis & Policy*, 3(1)

[11] Babcock, Linda, and Sara Laschever 2003 *Women Don't Ask: Negotiation and the Gender Divide* (Princeton, NJ: Princeton University Press)

[12] Bajari, Patrick and Matthew E. Kahn (2004), "The Private and Social Costs of Urban Sprawl: The Lot Size Versus Commuting Tradeoff," working paper, July.

[13] Baum-Snow, Nathaniel (2007), "Did Highways Cause Suburbanization?," Quarterly Journal of Economics, vol. 122, pp. 775-805.

[14] Bento, Antonio, Lawrence Goulder, Mark Jacobsen, and Roger von Haefen. 2009. "Distributional and Efficiency Impacts of Increased U.S. Gasoline Taxes." *American Economic Review*, 99(3): 667-669.

[15] Beresteanu, Arie and Shanjun Li. Forthcoming. "Gasoline Prices, Government Support, and Demand for Hybrid Vehicles in the U.S." *International Economic Review*.

[16] Berry, Steve, 1994 "Estimating discrete choice models of product differentiation" *RAND Journal of Economics* 25:242-262

[17] Berry, Steve, Jim Levinsohn, and Ariel Pakes. 2004. "Differentiated product demand systems from a combination of micro and macro data: The new car market." *Journal of Political Economy*, 112(1): 68-105.

[18] Berry, Steve, Jim Levinsohn, and Ariel Pakes. 1995 "Automobile prices in market equilibrium" *Econometrica*, 63(4): 841-890

[19] Borenstein, Severin 1985 "Price discrimination in free-entry markets" *Rand Journal of Economics* 16(3): 380-397

[20] Borenstein, Severin 1991 "Selling costs and switching costs: explaining retail gasoline margins" *RAND Journal of Economics* 22(3): 354-369

[21] Bresnahan, Timothy F. 1981 "Departures from marginal-cost pricing in the American automobile industry: Estimates for 1977-1978" *Journal of Econometrics*, 17(2): 201-227

[22] Bresnahan, Timothy F. 1987 "Competition and collusion in the American automobile industry: The 1955 price war" *Journal of Industrial Economics* 35(4): 457-482

[23] Bresnahan, Timothy F., and Peter C. Reiss. 1985 "Dealer and manufacturer margins" *The RAND Journal of Economics*, 16(2): 253-268

[24] Brueckner, Jan (2007), "Urban Growth Boundaries: An Effective Second-Best Remedy for Unpriced Traffic Congestion?," Journal of Housing Economics, forthcoming.

[25] Brueckner, Jan (2000), "Urban Sprawl: Diagnosis and Remedies," International Regional Science Review, vol. 23, pp. 160-171.

[26] Burchell, Robert W. , Naveed A. Shad, David Listokin, Hilary Phillips, Anthony Downs, Samuel Seskin, Judy S. Davis, Terry Moore, David Helton, and Michelle Gall (1998), "The Costs of Sprawl—Revisited," Transit Cooperative Research Program Report 39, Transportation Research Board, Washington, DC

[27] Busse, Megan R., Christopher R. Knittel, and Florian Zettelmeyer. 2009. "Pain at the Pump: How Gasoline Prices Affect Automobile Purchasing." http://www.econ.ucdavis.edu/faculty/knittel/gaspaper_latest.pdf

[28] Calfee, John and Clifford Winston (1998), "The Value of Automobile Travel Time: Implications for Congestion Policy," Journal of Public Economics, July, vol. 69. pp. 83-102.

[29] Calfee, John, Clifford Winston, and Randolph Stempski (2001), "Econometric Issues in Estimating Consumer Preferences from Stated Preference Data: A Case Study of the Value of Automobile Travel Time," Review of Economics and Statistics, November, vol. 83, pp. 699-707.

[30] Charles, Kerwin Kofi, and Jonathan Guryan 2008 "Prejudice and wages: An empirical assessment of Becker's *The Economics of Discrimination*" *Journal of Political Economy* 116(5): 773-809

[31] Charles, Kerwin Kofi, Jonathan Guryan, and Jessica Pan "Sexism and women's labor market outcomes" Working Paper

[32] Charles, Kerwin Kofi, Erik Hurst, and Melvin Stephens Jr. 2008 "Rates for vehicle loans: race and loan source" *American Economic Review, Papers and Proceedings* 315-320

[33] Copeland, Adam, Wendy Dunn, and George Hall. 2005. Prices, Production and Inventories over the Automotive Model Year." NBER Working Paper 11257.

[34] Corrado, Carol, Wendy Dunn, and Maria Otoo. 2006. "Incentives and Prices for Motor Vehicles: What Has Been Happening in Recent Years?" FEDS Working Paper 2006-09.

[35] Corts, Kenneth S. "Third-degree price discrimination in oligopoly: All-out competition and strategic commitment" RAND Journal of Economics 29(2): 306-323

[36] Davies, Peter. 2007. "What's the Value of an Energy Economist?" Speech presented at the International Association of Energy Economics, Wellington, New Zealand.

[37] Davis, Michael C. and James D. Hamilton. 2004. "Why Are Prices Sticky? The Dynamics of Wholesalse Gasoline Prices." *Journal of Money, Credit and Banking*, 36 (1): 17-37.

[38] Deakin, Elizabeth (1994), "Urban Transportation Congestion Pricing: Effects on Urban Form," in Curbing Gridlock: Peak-Period Fees to Relieve Traffic Congestion, Volume 2, Commissioned Papers, Transportation Research Board, Washington, DC.

[39] Dickey, D. A. and W. A. Fuller. 1979. "Distribution of the Estimators for Autoregressive Time Series with a Unit Root." *Journal of the American Statistical Association*, 74: 427Ű431.

[40] Energy Information Agency. 2006. "A Primer on Gasoline Prices." http://www.eia.doe.gov/bookshelf/brochures/gasolinepricesprimer/

[41] Fang, Hao Audrey (2006), "How Urban Sprawl Matters: A Joint Modeling of Households' Vehicle Choice and Vehicle Usage," UC Irvine working paper.

[42] Geman, Hélyette. 2007. "Mean Reversion Versus Random Walk in Oil and Natural Gas Prices." In *Advances in Mathematical Finance*, ed. Michael C. Fu, Robert A. Jarrow, Ju-Yi Yen, and Robert J. Elliott, 219-228. Boston: Birkhauser.

[43] Gicheva, Dora, Justine Hastings, and Sofia Villas-Boas. 2007. "Revisiting the Income Effect: Gasoline Prices and Grocery Purchases." CUDARE Working Paper 1044.

[44] Glaeser, Edward L. and Joseph Gyourko (2003), "The Impact of Zoning on Housing Affordability," Economic Policy Review, June, vol. 9, pp. 21-39.

[45] Glazer, Amihai and Kurt Van Dender (2002), "How Congestion Pricing Reduces Property Values," working paper 01-02-08, UC Irvine.

[46] Gneezy, Uri, Kenneth L. Leonard, and John A List 2009 "Gender differences in competition: Evidence from a matrilineal and a patriarchal society" *Econometrica* 77(5): 1637-1664

[47] Goldberg, Pinelopi. 1998. "The effects of the corporate average fuel economy standards in the automobile industry." *Journal of Industrial Economics*, 46: 1-33.

[48] Goldberg, Pinelopi. 1996 "Price discrimination in new car purchases: evidence from the consumer expenditure survey" *Journal of Political Economy*: 622-54

[49] Graddy, Kathryn 1995 "Testing for imperfect competition at the Fulton fish market" *RAND Journal of Economics* 26(1):75-92

[50] Graddy, Kathryn 1997 "Do fast-food chains price discriminate on the race and income characteristics of an area?" *Journal of Business and Economic Statistics* 15(4): 391-401

[51] Graddy, Kathryn, and George Hall 2009 "A dynamic model of price discrimination and inventory management at the Fulton fish market" *NBER Working Paper* 15019

[52] Gramlich, Jacob. 2008. "Gas Prices and Endogenous Product Selection in the U.S. Automobile Industry." http://pantheon.yale.edu/~jpg39/JMP_GramlichJacob.pdf

[53] Harless, David W. and George E Hoffer (2002) "Do women pay more for new vehicles? Evidence from transaction price data" *American Economic Review* 92(1): 270-279

[54] Heckman, James J., and Peter Siegelman 1993 "The Urban Institute audit studies: Their methods and findings" in *Clear and Convincing Evidence: Measurement of Discrimination in America*, M. Fix and R. Struyk, eds. (Washington, DC: The Urban Institute Press)

[55] Himmelberg, Charles, Christopher Mayer, and Todd Sinai, "Assessing High House Prices: Bubbles, Fundamentals, and Misperceptions," Journal of Economic Perspectives, Fall 2005, vol. 19, pp. 67-92.

[56] Holmes, Thomas J. "The effects of third-degree price discrimination in oligopoly" *American Economic Review* 79(1): 244-250

[57] Ireland, Norman J. "On the welfare effects of regulating price discrimination" *Journal of Industrial Economics* 40(3): 237-248

[58] Jacobsen, Mark. 2008. “Evaluating U.S. Fuel Economy Standards in a Model with Producer and Household Heterogeneity.” http://www.econ.ucsd.edu/∼m3jacobs/papers.html

[59] Kahn, Matthew E. (2006), Green Cities: Urban Growth and the Environment, Brookings Institution Press, Washington DC.

[60] Klier, Thomas and Joshua Linn. 2008. “The Price of Gasoline and the Demand for Fuel Efficiency: Evidence from Monthly New Vehicles Sales Data.” http://tigger.uic.edu/∼jlinn/vehicles.pdf

[61] Lee, Douglas B. (1982), “Net Benefits from Efficient Highway User Charges,” Transportation Research Record, no. 858, pp. 14-20.

[62] Lee, Douglas B. (1992) “A Market-Oriented Transportation and Land Use System: How Different Would It Be?,” US Department of Transportation working paper

[63] Li, Shanjun, Christopher Timmins, and Roger H. von Haefen. Forthcoming. “How Do Gasoline Prices Affect Fleet Fuel Economy?” *American Economic Journal: Economic Policy.*

[64] Lindsey, Robin (2006), “Do Economists Reach A Conclusion on Road Pricing? The Intellectual History of an Idea,” Econ Journal Watch, May, vol. 3, pp. 292-379.

[65] List, John. 2004 “The nature and extent of discrimination in the marketplace: Evidence from the field.” *Quarterly Journal of Economics* 119(1): 49-89

[66] Lott, John R. Jr and Russell D. Roberts “A guide to the pitfalls of identifying price discrimination” *Economic Inquiry* 29: 14-23

[67] Makridakis, Spyros, Steven C. Wheelwright, and Rob J. Hyndman. 1998. *Forecasting Methods and Applications, 3rd ed.* New York: Wiley.

[68] Mannering, Fred and Clifford Winston (1985), “A Dynamic Empirical Analysis of Household Vehicle Ownership and Utilization,” Rand Journal of Economics, Summer, vol. 16, pp. 215-236.

[69] McFadden, Daniel. 1974 “Conditional logit analysis of qualitative choice behavior” in P. Zarembka (ed.) *Frontiers in Econometrics*: 105-142

[70] Miller, Don M. and Dan Williams. 2004. “Damping Seasonal Factors: Shrinkage Estimators for the X-12-ARIMA Program. *International Journal of Forecasting*, 20 (4): 529-549.

[71] Mills, Edwin S. (1967), “An Aggregative Model of Resource Allocation in Metropolitan Areas,” American Economic Review, May, vol. 57, pp. 197-210.

[72] Mohring, Herbert (1999), “Congestion,” in Essays in Transportation Economics and Policy: A Handbook in Honor of John R. Meyer, J. Gomez-Ibanez, W. Tye, and C. Winston, editors, Brookings, Washington DC.

[73] Muth, Richard F. (1969), Cities and Housing, University of Chicago Press, Chicago, Illinois.

[74] Nechyba, Thomas and Randall Walsh (2004), “Urban Sprawl,” Journal of Economic Perspectives, Fall, vol. 18, pp. 177-200.

[75] Newey, Whitney K. and Kenneth D. West. 1987. “A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix.” *Econometrica*, 55 (3): 703-708.

[76] Parker, Steve. 2009 “Bottom Line: 2008 Car/Truck Sales Probably Worst Since 1974” *The Huffington Post* January 3, 2009

[77] Parry, Ian W. H., Winston Harrington, and Margaret Walls. 2007. “Automobile Externalities and Policies.” *Journal of Economic Literature*, 45, 373-400.

[78] Phelps, Edmund S. 1972 “The statistical theory of racism and sexism” *American Economic Review* 62: 659-661

[79] Pickrell, Don (1999), “Transportation and Land Use,” in Essays in Transportation Economics and Policy: A Handbook in Honor of John R. Meyer, J. Gomez- Ibanez, W. Tye, and C. Winston, editors, Brookings, Washington DC.

[80] Quddus, Mohammed A., Alon Carmel, and Michael G. H. Bell (2007), “The Impact of the Congestion Charge on Retail: The London Experience,” Journal of Transport Economics and Policy, January, vol. 41, pp. 113-133.

[81] Revelt, D. and K. Train. 1998 “Mixed logit with repeated choices: Households’ choices of appliance efficiency level” *Review of Economics and Statistics* LXXX(4): 647-657

[82] RTI International 2009 “Automobile industry retail price equivalent and indirect cost multipliers” prepared for US Environmental Protection Agency. Project No. EPA-420-R-09-003 February 2009

[83] Salop, Steven, and Joseph Stiglitz. 1977 “Bargains and ripoffs: A model of monopolistically competitive price dispersion,” *Review of Economic Studies* XLIV: 493-510

[84] Santos, Georgina, ed. (2004), Road Pricing: Theory and Evidence. Elsevier Press, Oxford, United Kingdom. Segal, David and Thomas L. Steinmeier (1980), “The Incidence of Congestion and Congestion Tolls,” Journal of Urban Economics, January, vol. 7, pp. 42-62.

[85] Schmalensee, Richard 1981 “Output and welfare implications of monopolistic third-degree price discrimination” *American Economic Review* 71: 242-247

[86] Scott Morton, Fiona M., Florian Zettelmeyer, and Jorge Silva-Risso 2003 “Consumer information and discrimination: Does the internet affect the pricing of new cars to women and minorities?” *Quantitative Marketing and Economics* 1: 65-92

[87] Shepard, Andrea 1991 “Price discrimination and retail configuration” *Journal of Political Economy* 99(1): 30-53

[88] Shirley, Chad and Clifford Winston (2004), “Firm Inventory Behavior and the Returns from Highway Infrastructure Investments,” Journal of Urban Economics, March, vol. 55, pp. 398-415.

[89] Small, Kenneth A. (1992), Urban Transportation Economics, Harwood Academic Publishers, Chur, Switzerland.

[90] Small, Kenneth A. and Kurt Van Dender. 2007. “Fuel Efficiency and Motor Vehicle Travel: The Declining Rebound Effect.” *Energy Journal*, 28 (1), 25-51.

[91] Small, Kenneth A., Clifford Winston, and Carol A. Evans (1989), Road Work: A New Highway Pricing and Investment Policy, Brookings, Washington DC.

[92] Small, Kenneth A., Clifford Winston, and Jia Yan (2006), “Differentiated Road Pricing, Express Lanes, and Carpools: Exploiting Heterogeneous Preferences in Policy Design,” Brookings Wharton Papers on Urban Affairs, pp. 53-96.

[93] Song, Yan and Gerrit-Jan Knaap (2003), “New Urbanism and Housing Values: A Disaggregate Assessment, Journal of Urban Economics, September, vol. 54, pp. 218-238.

[94] Villas-Boas, Sofia Berto "An empirical investigation of the welfare effects of banning wholesale price discrimination" *RAND Journal of Economics* 40(1): 20-46

[95] Train, Kenneth 2003 *Discrete Choice Methods with Simulation* (Cambrigde, MA: Cambridge University Press)

[96] Wheaton, William C. (1998), "Land Use and Density in Cities with Congestion," Journal of Urban Economics, March, vol. 43, pp. 258-272.

[97] Train, Kenneth and Clifford Winston. 2007. "Vehicle choice behavior and the declining market share of U.S. automakers." *International Economic Review*, 48: 1469-1496.

[98] US Department of Transportation Federal Highway Administration 2008. *Highway Statistics* Washington, DC.

[99] US Environmental Protection Agency 2010. *2010 U.S. Greenhouse Gas Inventory Report: Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2008* U.S. EPA #430-R-10-006 Washington, DC.

[100] Winston, Clifford (2006), Government Failure versus Market Failure: Microeconomics Policy Research and Government Performance, AEI-Brookings Institution, Washington DC, (www.aei-brookings.org/publications/abstract.php?pid=1117).

[101] Winston, Clifford and Ashley Langer (2006), "The Effect of Government Highway Spending on Road Users' Congestion Costs," Journal of Urban Economics, November, vol. 60, pp. 463-483.

[102] Winston, Clifford and Vikram Maheshri (2007), "On the Social Desirability of Urban Rail Transit Systems, Journal of Urban Economics, September, vol. 62, pp. 362-382.

[103] Winston, Clifford and Chad Shirley (1998), Alternate Route: Toward Efficient Urban Transportation, Brookings, Washington DC.

[104] Yinger, John 1998 "Evidence on discrimination in consumer markets" *Journal of Economic Perspectives* 12(2): 23-40

# Chapter 5

# Appendix to Chapter 2

## 5.1 Elasticity Bias

In our introductory remarks, we argued informally that structural estimation can understate consumer responsiveness to fuel costs if it fails to account for manufacturer price responses. We formalize our argument here in the context of logit demand. The relevant regression equation is:

$$\log(s_{jt}) - \log(s_{0t}) = \psi(p_{jt} + x_{jt}) + \kappa_j + \nu_{jt}, \tag{5.1}$$

where $s_{jt}$ and $s_{0t}$ are the market shares of vehicle $j$ and the outside good, respectively, $p_{jt}$ is the vehicle price, $x_{jt}$ captures the expected lifetime fuel costs, $\kappa_j$ is vehicle "quality," and $\nu_{jt}$ is an error term that captures demand shocks. Assuming away endogeneity issues, one can use OLS with vehicle fixed effects to obtain consistent estimates of $\psi$, the parameter of interest. However, suppose that one observes the mean price of each vehicle, $\overline{p}_j$, rather than the true time-specific price. The regression equation becomes:

$$\log(s_{jt}) - \log(s_{0t}) = \psi x_{jt} + \kappa_j^* + \nu_{jt}^*, \tag{5.2}$$

where $\kappa_j^* = \kappa_j + \psi\overline{p}_j$ and $\nu_{jt}^* = \nu_{jt} + \psi(p_{jt} - \overline{p}_j)$. The problem is now apparent. Gasoline price shocks affect not only $x_{jt}$ but also the composite error term $\nu_{jt}^*$ through the manufacturer response, and the regression coefficient is inconsistent:

$$\widehat{\psi} \xrightarrow{p} \psi\left(1 + \frac{Cov(x_{jt}, p_{jt})}{Var(x_{jt})}\right). \tag{5.3}$$

The OLS estimate of $\psi$ is biased downwards because fuel costs and vehicle prices are negatively correlated. In the working paper, we discuss how one might estimate bias and demonstrate that estimation is sensitive to the assumed discount rate. Although Equation 5.3 is specific to logit demand, our intuition is that other demand systems also generate bias; the unresolved problem remains negative correlation between (observed) fuel costs and (unobserved) vehicle price responses.

## 5.2 Additional Robustness Checks

### 5.2.1 Alternative weighting schemes

The baseline weighting scheme meets the criterion of the theoretical model that "closer" competitors receive greater weight. However, it also requires the comparison of vehicles over a set of (potentially) arbitrary vehicle characteristics. We develop four alternative weighting schemes in this appendix, and show that the results are broadly robust. Each weighting scheme sidesteps the selection of vehicle characteristics but fits the theoretical model less well relative to the baseline weighting scheme.

The results are shown in Table 5.1. In Column 1, we place equal weight on vehicles in the same segment and zero weight on vehicles in a different segment. In Column 2, we place equal weight on vehicles of the same type and zero weight on other vehicles. In Column 3, we place equal weight on all vehicles. Finally, in Column 4, we decompose the influence of competitor fuel costs into the effects of same-segment competitors, same-type competitors, and other competitors.[1] Across the four columns, the fuel cost coefficients are negative, the competitor fuel cost coefficients are positive, and the same-firm fuel cost coefficients (not shown) are small – consistent with both the baseline results and the theoretical model.

In terms of economic magnitudes, the regressions predict that a \$1 increase in the gasoline price would change the median manufacturer price by -\$422, \$145, \$64, and -\$120, respectively. The baseline regression produces a median effect of -\$171, so that the most flexible specification (Column 4) best matches the baseline results. Overall, we conclude that the basic intuition of the model – that manufacturer prices should decrease in fuel costs and increase in competitor fuel costs – is quite robust to the choice of the weighting scheme. More crude/restrictive weighing schemes, however, may produce less reasonable estimates of the net effect.

We also find it telling that, in Column 4, the coefficient on fuel costs of same-segment competitors is of greater magnitude than the coefficient on fuel costs of same-type competitors, which itself is of greater magnitude than the coefficient on

[1] Compact cars and luxury SUVs are examples of segments, and the vehicle types are cars, SUVs, trucks, and vans.

Table 5.1: Alternative Weighting Schemes

| | Weighting Approach: | | | |
|---|---|---|---|---|
| | Within Segment | Within Type | Overall Average | Within Group |
| Variables | (1) | (2) | (3) | (4) |
| Fuel cost | -33.50*** | -40.07*** | -23.90*** | -35.52*** |
| | (5.52) | (4.50) | (3.42) | (6.26) |
| Average competitor fuel cost | 24.18*** | 44.60*** | 22.27*** | |
| | (5.03) | (4.80) | (5.19) | |
| Average competitor fuel cost, same segment | | | | 19.54*** |
| | | | | (5.27) |
| Average competitor fuel cost, same type – different segment | | | | 10.90*** |
| | | | | (2.75) |
| Average competitor fuel cost, different type | | | | 3.56 |
| | | | | (4.62) |
| $R^2$ | 0.5227 | 0.5285 | 0.5222 | 0.5379 |

Results from OLS regressions. The dependent variable is the manufacturer price, i.e., MSRP minus the mean regional and national incentives (in thousands). The units of observation are at the vehicle-week-region level. The sample used in Columns 1, 2 and 3 includes 299,855 observations on 681 vehicles; the sample used in Column 4 includes 292,500 observations on 680 vehicles. All regressions include the appropriate average same-firm fuel cost variable(s). The regressions also include vehicle and time fixed effects, as well as third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

fuel costs of other competitors. This pattern underscores the notion, developed in the theoretical model and implemented in the baseline weighting scheme, that the fuel costs of close competitors are more relevant that those of more distant competitors.

### 5.2.2 Non-linear regression specifications

A striking characteristic of the main results is that the relationship between gasoline prices and manufacturer prices is concave (see Figures 3.4, 3.6, and 3.7). The concavity is consistent with the intuition that the effect of gasoline prices on fuel costs should diminish in vehicle fuel efficiency (e.g., gasoline prices are irrelevant for infinitely fuel-efficient vehicles). In this appendix, we provide some evidence that the estimated concavity is data-driven rather than an artifact of the regression specification.

We start with a simple plot of the residuals estimated from the baseline regression (Table 3.2 from the paper, Column 1) against vehicle miles-per-gallon. The plot appears in Figure 5.1. By construction, the residuals are mean zero and uncorrelated with fuel costs and the other regressors. Nonetheless, any misspecification of functional form should create a non-linear relationship between the residuals and fuel efficiency (e.g., a $\bigcup$-shape or $\bigcap$-shape relationship). No such non-linear relationship is evident, and a regression of the residuals on a second-order polynomial in fuel efficiency produces tiny and statistically insignificant coefficients.[2]

We also run regressions that alternatively include 1) squared fuel cost variables and 2) interactions between the fuel cost variables; these results appear in Table 5.2.[3] As shown, the standard linear fuel cost and competitor fuel costs coefficients operate similarly to those from the baseline results. By way of contrast, the new non-linear terms are of substantially smaller magnitude and are quite imprecisely measured.[4] Figure 5.2 plots the estimated effects of a one dollar increase in the gasoline price against vehicle miles-per-gallon, based on the specification with squared fuel cost variables. The concavity of the relationship is apparent. (Notably, Figure 5.2 is extremely close to Figure 3 in the text; the

[2]We average the residuals of each vehicle for graphical clarity. The procedure reduces the magnitude of the residuals but does not affect inference regarding the linearity of the relationship between gasoline prices and fuel efficiency.

[3]The more general specification with both squared terms and interactions overtaxes identification.

[4]In comparing these magnitudes, it is useful to keep in mind that the coefficient on the squared terms and interactions terms must be roughly five times the magnitude of the linear terms to produce the same economic effect. Mathematically, $\frac{\partial p}{\partial x} = \beta_1 + 2\beta_2 x$, where $\beta_1$ is the fuel cost coefficient, $\beta_2$ is the fuel cost squared coefficient, and a typical value of $x$ is around 0.10.

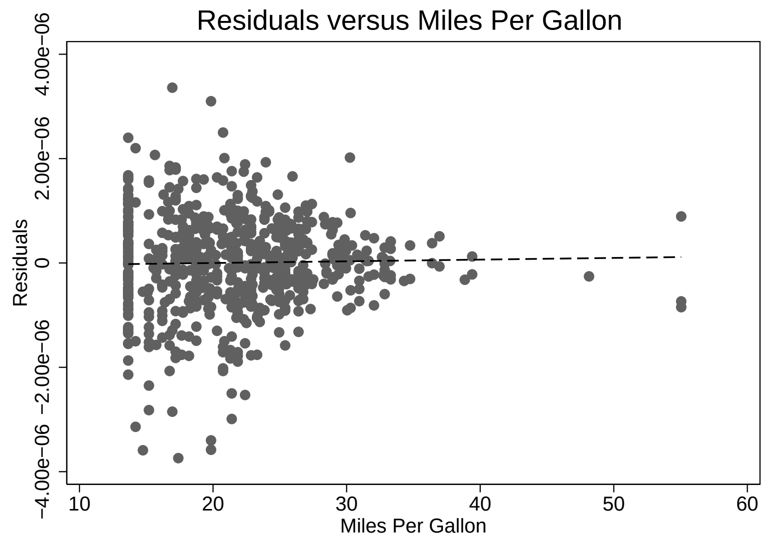


Figure 5.1: The estimated residuals from the baseline results estimated in Table 2 (Column 1) of the paper. Each point represents the mean residual for a single vehicle.

correlation coefficient between the two price effects is 0.9936.)

We interpret these results as substantial support for both the theoretical model and the linear regression equation we derive from it.

### 5.2.3 Model-year and region subsamples

The coefficients we estimate are reduced-form combinations of many structural parameters. To motivate the baseline regression specification, we assume that thee underlying structural parameters do not change over time nor over geographic areas. We examine that assumption in this appendix. In particular, we show that several sub-samples produce results similar to those of the main sample, and interpret the robustness of the baseline results as evidence that the structural parameters are roughly homogenous across time and regions.

First, we estimate subsample regressions by model-year so that the coefficients are flexible over time. The results appear in Table 5.3. As shown, the fuel cost coefficients for the 2004, 2005, and 2006 model-years are quite similar to those of the baseline results but the coefficients from the 2003 model-year are of smaller magnitude. Although we cannot be sure why the 2003 model-year coefficients are smaller, we suspect that the differences are due to sample construction – we do not observe the 2003 model-years until January 1, 2003 (most are introduced in Summer/Fall 2002) and we also do not observe some of the relevant

Table 5.2: Tests for Non-Linearity

| Variables | (1) | (2) | (3) |
|---|---|---|---|
| Fuel cost | -72.51*** | -57.47*** | -52.73*** |
| | (14.25) | (9.02) | (18.50) |
| Fuel cost$^2$ | 42.01 | | |
| | (39.56) | | |
| Fuel cost $*$ Gas price | | | -0.99 |
| | | | (7.72) |
| Average competitor fuel cost | 67.62*** | 50.04*** | 63.80*** |
| | (15.72) | (6.86) | (19.19) |
| Average competitor fuel cost$^2$ | -45.03 | | |
| | (47.79) | | |
| Fuel cost $*$ Average competitor fuel cost | | 8.17 | |
| | | (12.05) | |
| Average competitor fuel cost$*$ Gas Price | | | -5.79 |
| | | | (8.01) |
| $R^2$ | 0.5262 | 0.5260 | 0.5262 |

Results from OLS regressions. The dependent variable is the manufacturer price, i.e., MSRP minus the mean regional and national incentives (in thousands). The sample includes 299,855 observations on 681 vehicles at the vehicle-week-region level. All regressions include the average same-firm fuel cost variable. All regressions also include vehicle, time, and region fixed effects, as well as third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

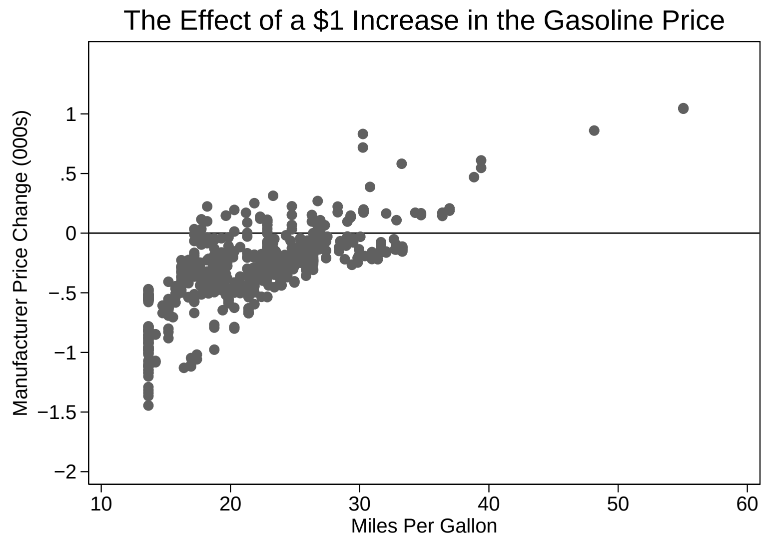


Figure 5.2: The estimated effects of a one dollar increase in the retail gasoline price on the manufacturer price. Based on Table 5.2 (Column 1), which includes squared fuel costs and competitor fuel costs terms. Each point represents the price effect for a single vehicle.

competitors (the unobserved 2002 model-years compete into 2003). Overall, the robustness of the results across model-years is notable.

Second, we estimate subsample regressions by region so that the coefficients are flexible over geographic space. The results appear in Table 5.4. Again, the fuel cost coefficients are quite similar to those of the baseline results. As a final test, we construct two specific *city-level* samples. We chose two cities that are arguable be at the extremes of the consumer demand for fuel efficiency – San Francisco and Houston (e.g., see Li, Timmins, and von Haefen forthcoming). The quality of the incentives data are lower at the city-level because the applicable area for some incentives is listed as “Select Counties in CA” or “Various Counties in TX” and we cannot assign these incentives to any particular city/cities. (High quality gasoline price data for both cities are available from the EIA.) As shown, the results mimic the those generated from the main sample even despite the poorer data quality.

Table 5.3: Model Year Subsamples

| | Model-Year | | | |
|---|---|---|---|---|
| Variables | 2003 | 2004 | 2005 | 2006 |
| Fuel cost | -21.58* | -58.98*** | -63.47*** | -57.57*** |
| | (12.31) | (18.70) | (11.77) | (13.80) |
| Average competitor fuel cost | 26.13* | 36.06*** | 54.17*** | 55.27*** |
| | (14.50) | (14.06) | (10.33) | (13.09) |
| Average same-firm fuel cost | -9.86*** | 8.90 | 5.01 | 2.38 |
| | (3.28) | (7.85) | (4.96) | (2.95) |
| $R^2$ | 0.2408 | 0.6812 | 0.5530 | 0.6199 |
| # of observations | 62,105 | 85,885 | 88,550 | 47,805 |
| # of vehicles | 163 | 170 | 176 | 172 |

Results from OLS regressions. The dependent variable is the manufacturer price, i.e., MSRP minus the mean regional and national incentives (in thousands). The units of observation are at the vehicle-week-region level. All regressions include vehicle, time, and region fixed effects, as well as third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

Table 5.4: Manufacturer Prices and Fuel Costs

| | Region: | | | | | City: | |
|---|---|---|---|---|---|---|---|
| Variables | East Coast | Gulf Coast | Midwest | Mountain West | West Coast | San Francisco | Houston |
| Fuel cost | -60.05*** | -56.39*** | -59.43*** | -67.00*** | -43.43*** | -59.86*** | -50.76*** |
| | (8.52) | (8.02) | (8.33) | (8.83) | (6.03) | (6.85) | (7.12) |
| Average competitor fuel cost | 49.07*** | 48.09*** | 48.95*** | 52.94*** | 40.16*** | 39.58*** | 31.85*** |
| | (7.91) | (7.58) | (7.84) | (8.20) | (5.77) | (7.25) | (7.54) |
| $R^2$ | 0.5280 | 0.5275 | 0.5252 | 0.5290 | 0.5239 | 0.4465 | 0.4099 |
| # of observations | 59,971 | 59,971 | 59,971 | 59,971 | 59,971 | 59,971 | 59,971 |
| # of vehicles | 681 | 681 | 681 | 681 | 681 | 681 | 681 |

Results from OLS regressions. The dependent variable is the manufacturer price, i.e., MSRP minus the mean regional and national incentives (in thousands). The units of observation are at the vehicle-week-region level. All regressions include vehicle, time, and region fixed effects, as well as third-order polynomials in the vehicle age (i.e., weeks since the date of initial production), the average age of vehicles produced by different manufacturers, and the average age of other vehicles produced by the same manufacturer. Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

### 5.2.4 Understanding Manufacturer versus Dealer Risk

In order to understand how vehicle manufacturers and dealers share the risk associated with changing gas prices, we estimated a version of the Busse, Knittel and Zettelmeyer (2009) regressions. Because their regressions look at vehicle transaction price responses to changes in gasoline prices, if our results are similar to theirs, then we can conclude that the majority of vehicle price risk is absorbed by manufacturers rather than dealers. We regressed manufacturer price on gasoline prices, interacted with a series of dummy variables based on the vehicle's fuel efficiency relative to others of the same type (we place vehicles into quartiles), as well as the standard controls. The specification should mitigate any bias due to endogenous fuel efficiencies because models should not change type or efficiency quartile over time.[5]

Table 5.5 shows the results. For cars in the least efficient quartile, a $1 dollar increase in the gasoline price is associated with a fall in the manufacturer price of $1,660; for cars in the second and third quartiles the fall is $490 and $200, respectively, and the prices of the most fuel efficient cars increase by $320. The SUV coefficients show a similar pattern, though prices fall even for the most efficient quartile. The effects are less pronounced for trucks and fuel efficiency does not appear to matter for vans. Overall, comparing these results with those in Busse, Knittel, and Zettelmeyer (2009) supports the idea that manufacturers are absorbing the majority of the revenue risk resulting from changes in gas prices.

[5]The specification is also a nice robustness check to the main results because it side-steps the weighting scheme entirely and is based on an entirely different functional form.

Table 5.5: Busse, Knittel, Zettelmeyer (2009) Specification

Matrix of Coefficients and Standard Errors

| | Cars | SUVs | Trucks | Vans |
|---|---|---|---|---|
| Gas price | -1.66*** | -1.63*** | | |
| ∗ Quartile 1 dummy | (0.24) | (0.19) | | |
| Gas price | -0.49*** | -0.79*** | | |
| ∗ Quartile 2 dummy | (0.10) | (0.11) | | |
| Gas price | -0.20** | -0.28*** | | |
| ∗ Quartile 3 dummy | (0.09) | (0.07) | | |
| Gas price | 0.32*** | -0.20*** | | |
| ∗ Quartile 4 dummy | (0.07) | (0.09) | | |
| Gas price | | | -0.43*** | -0.17*** |
| ∗ Half 1 dummy | | | (0.07) | (0.07) |
| Gas price | | | 0.02 | -0.18*** |
| ∗ Half 2 dummy | | | (0.09) | (0.06) |

Results from a single OLS regression. The dependent variable is the manufacturer price, i.e., MSRP minus the mean regional and national incentives (in thousands). The sample includes 299,855 observations on 681 vehicles at the vehicle-week-region level. All regressions include vehicle and region fixed effects, as well as a third-order polynomial in the vehicle age (i.e., weeks since the date of initial production). Standard errors are clustered at the vehicle level and shown in parenthesis. Statistical significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

# Chapter 6

# Appendix to Chapter 3

This appendix derives the measures of delay per mile and unpriced congestion. We use the Bureau of Public Roads (BPR) speed flow curve:

$$CongestedSpeed = \frac{FreeFlowSpeed}{1 + .15 * (V/C)^4}$$

where $V$ is volume of traffic, $C$ is capacity, and we assume a free flow speed of 60 miles per hour. Using this equation, delay per mile is equal to the inverse of congested speed (CS) minus free flow speed (FFS).

Unpriced congestion is the difference between the social marginal cost of delay and the average cost of delay. The average cost of delay (in dollars per mile) to a motorist is equal to the average value of travel time (in dollars per hour) times the delay per mile:

$$\begin{aligned} ValueOfTime * \left(\tfrac{1}{CS} - \tfrac{1}{FFS}\right) &= ValueOfTime * \left(\tfrac{1+.15*(V/C)^4}{FFS} - \tfrac{1}{FFS}\right) \\ &= ValueOfTime * \left(\tfrac{.15*(V/C)^4}{FFS}\right) \end{aligned} \tag{6.1}$$

Thus, the total cost of all driving on a given mile of road is:

$$ValueOfTime * \left(\frac{.15 * (V/C)^4}{FFS}\right) * V = ValueOfTime * \left(\frac{.15}{FFS * C^4}\right) * V^5.$$

The marginal cost of driving is:

$$\frac{\partial TC}{\partial V} = MC = 5 * \left(\frac{ValueOfTime * .15}{FFS}\right) * \left(\frac{V}{C}\right)^4. \tag{6.2}$$

Thus, the distortion from unpriced congestion is obtained by subtracting the average cost given in equation 6.1 from the marginal cost in equation 6.2 and plugging in FFS = 60:

$$Distortion = 4 * \left(\frac{ValueOfTime * .15}{60}\right) * \left(\frac{V}{C}\right)^4 = ValueOfTime * .01 * \left(\frac{V}{C}\right)^4.$$