EMERGING MARKETS FOR BIOFUELS

by

Soren Tyler Anderson

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2008

Doctoral Committee:

Assistant Professor Lucas William Davis, Co-Chair
Professor Stephen W. Salant, Co-Chair
Emeritus Professor Gary Rand Solon
Assistant Professor Meredith Lynn Fowlie

For Emily

ACKNOWLEDGEMENTS

For their wonderful guidance, I thank my committee members Lucas Davis, Meredith Fowlie, Stephen Salant, and Gary Solon. For his help in co-authoring the third chapter of this dissertation, I thank my classmate James Sallee. Together we thank Wei Fan for significant help with the vehicle transaction data. For helpful comments and suggestions, I thank Matias Busso, Meghan Busse, Brian Cadena, Steven Haider, Ben Keys, Brian Kovak, Michael Moore, Richard Newell, Richard Porter, Alex Resch, Adam Schubert, Mike Taylor, Roger von Haefen, Sarah West, and seminar participants at the University of Michigan, University of California Energy Institute, NBER Summer Workshop, Michigan State University, and TREE Seminar (Duke, NC State, and RTI International).

I thank the Minnesota Department of Commerce and the American Lung Association for providing ethanol price and quantity data, the Minnesota Department of Public Safety for providing vehicle registration data, and the U.S. Department of Transportation for providing fuel-economy data. I thank Alan Berkowitz and Peter Feather for answering numerous questions about CAFE standards. I thank the University of Michigan's Center for Local, State, and Urban Policy and Rackham Graduate School for research funding.

The research described in this dissertation has been funded wholly or in part by the U.S. Environmental Protection Agency (EPA) under the Science to Achieve Results (STAR) Graduate Fellowship Program. EPA has not officially endorsed this publication and the views expressed herein may not reflect the views of the EPA.

All remaining errors are my own.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Introduction

Policies designed to reduce U.S. petroleum consumption increasingly promote markets for ethanol and other biofuels through subsidies, mandates, and funding for research. Many policymakers argue that substituting toward biofuels will enhance national energy security, reduce carbon dioxide emissions, mitigate local air and water quality impacts associated with petroleum refining and consumption, and benefit domestic farmers. Many recent policies mandate, either explicitly or implicitly, a minimum market share for ethanol. A prime example is the federal renewable fuels standard, which requires a minimum quantity of renewable fuel in the nation's fuel supply. Despite this attention from policymakers, relatively little is known about household preferences for biofuels or how expanded markets for these fuels might operate. Such information is critical for designing, implementing, and evaluating policies to promote ethanol and other biofuels.

In chapter II of this dissertation, I address this important research need by estimating demand for ethanol as a gasoline substitute. I find that demand for ethanol is sensitive to relative fuel prices, with elasticities that average from 2.5–3.0. These are the first available estimates in the literature for the price elasticity of household ethanol demand, which is a critical parameter for studies that analyze a retail ethanol subsidy or mandate. Price responses are substantially smaller and less variable than they would be if household preferences for ethanol were identical or nearly identical, and fuel-switching behavior extends over a wide range of prices where ethanol is discounted 0%-25% below gasoline. This implies that household preferences for ethanol are diffuse.

These results have important implications for policy. Previous economic analyses of government interventions that promote ethanol assume that household preferences are identical and depend exclusively on ethanol's fuel economy performance relative to gasoline (Holland et al. 2007). This assumption can yield misleading results if households also value ethanol for its perceived environmental and social benefits. I find that preferences for ethanol as a gasoline substitute vary dramatically, with some households exhibiting a marked preference for ethanol. Households with particularly strong preferences represent "low-hanging fruit" that can be induced to purchase ethanol with less severe distortion of market prices. Accounting for this heterogeneity reduces the economic efficiency costs of an ethanol content standard (i.e., a minimum market share requirement) by as much as 50% relative to previous studies, which incorrectly assume identical preferences. Similar intuition likely applies for policies that promote other "green" substitutes, such as renewable electricity, energy-efficient light bulbs and appliances, hybrid vehicles, and organic foods.

I begin my analysis by developing a model of household utility with inputs of ethanol and gasoline. These inputs are perfect substitutes in producing household transportation services. The key parameter in this model is the relative price at which the household is indifferent between relying entirely on either fuel. When this parameter varies continuously among households, aggregate demand for ethanol is a smooth function of relative fuel prices. The model formalizes the link between the distribution of household preferences for ethanol and the response of aggregate demand to changes in relative fuel prices. This allows me to recover the distribution of household preferences for ethanol as a gasoline substitute from observed price responses. The model also distinguishes price responses associated with fuel switching from those associated with changes in overall fuel demand.

I estimate the model using a unique dataset that contains more than 5000 monthly observations for ethanol prices and sales volumes at nearly 240 individual retail fueling stations during 1997–2006. These data provide a rare opportunity to document household preferences for biofuels, whose market shares have generally been too small to warrant inclusion in household surveys or to be reported separately from gasoline in aggregate measures. I use these data to estimate demand for ethanol as a function of ethanol prices and gasoline prices. Consistent with my theoretical model, which implies that price elasticities might vary dramatically, I allow aggregate ethanol demand to be a flexible function of relative fuel prices in my empirical specification. Previous empirical studies of demand for alternative fuels and gasoline varieties with close substitutes do not allow for this important flexibility.

I use the distribution of household preferences implied by my estimates to simulate the effects of a national ethanol content standard for gasoline. I find that a 25% ethanol content standard would decrease gasoline consumption by 21% and would cut carbon dioxide emissions from gasoline by 12% at an annual economic efficiency cost of about $50 billion. Efficiency losses derive primarily from ethanol's higher marginal cost. The ethanol content standard is costly relative to benefits, even after accounting for heterogeneity. Costs average about $370 per metric ton of carbon dioxide emissions avoided, which is several times higher than even pessimistic estimates of marginal external damages, or about $1.90 per gallon of gasoline saved, which is well beyond reasonable estimates for the external cost associated with petroleum dependence. These results may understate the costs of the ethanol standard, as I do not account for impacts on food prices and as recent studies indicate that ethanol may actually increase carbon dioxide emissions.

The empirical economic literature on demand for biofuels is miniscule.[1] While an immense literature estimates demand for gasoline, the vast majority of studies focus on the response of overall fuel demand to changes in fuel price levels. Because households have relatively few transportation alternatives, fuel demand in the short run is price inelastic.[2] This analysis in contrast focuses on fuel-switching behavior and how demand for ethanol as a gasoline substitute responds to changes in relative fuel prices. Because households that purchase ethanol are able to substitute easily between ethanol and gasoline, demand for ethanol is price elastic.

Within the transportation fuel demand literature, this analysis is most similar to studies that estimate demand for particular fuels with close substitutes, including full-service and self-serve gasoline (Phillips and Schutte 1988) and regular and premium gasoline in both leaded and unleaded varieties (Greene 1989). These studies find own-price and cross-price elasticities that exceed 10 in absolute value. Elasticities also tend to be large for other goods with close substitutes, including breakfast cereals (Nevo 2001), brand-name and generic pharmaceutical products (Ellison et al. 2006), and individual components of money supply (Barnett et al. 1992). I improve on this vein of the gasoline demand literature by formalizing fuel-switching behavior in terms of the distribution of household preferences for alternative fuels. I use a semi-parametric approach and other methods to estimate flexible econometric models that allow elasticities to vary with relative fuel prices.[3]

In chapter III of this dissertation, my coauthor (James Sallee) and I analyze the market for "flexible-fuel" vehicles that are able to burn ethanol and gasoline. While interesting in

[1]Rask (1998) estimates intermediate demand for ethanol as a 10% blending component in gasoline. He does not estimate household demand. Alves and Bueno (2003) estimate aggregate demand for gasoline in Brazil, which requires 25% ethanol blending in all gasoline, and where ethanol comprises roughly 40% of the non-diesel fuels market (Perkins and Barros 2006). They do not estimate price responses for ethanol.

[2]For recent surveys see Graham and Glaister (2002), Espey (1996, 1998), and Dahl and Sterner (1991). Recent studies indicate that the price response may have declined even further in recent decades (Hughes et al. 2008; Kilian 2008)

[3]Hausman and Newey (1995) and Yatchew and No (2001) estimate gasoline demand using a semi-parametric approach and other flexible methods. They do not model fuel switching.

its own right, this market also provides a unique opportunity to estimate the cost of tightening fuel-economy standards in the auto industry. An alternative-fuel incentive program built into federal fuel-economy regulations allows automakers to relax their fuel-economy constraints by adding flexible-fuel capacity to their vehicles. Because the incremental cost of adding flexible-fuel capacity is known, automakers that use this incentive program inadvertently reveal information about how costly it is for them to improve efficiency. I provide an overview of the theory, methods, and findings of this analysis in the background section of chapter III.

CHAPTER II

The Demand for Ethanol as a Gasoline Substitute

The format of the chapter is as follows. Section 2.1 provides background information on the role of ethanol in the fuels market, ethanol's environmental effects, and ethanol production and distribution. Section 2.2 presents a model of household demand for ethanol as a gasoline substitute, aggregates households to give an expression of aggregate demand, and relates the distribution of household preferences for ethanol to aggregate price responses. Section 2.3 describes the data I use to estimate the model, providing descriptive statistics that summarize supply and demand behavior. Section 2.4 outlines the econometric model I use to estimate demand, discusses identification, and presents my econometric results. Section 2.5 uses the distribution of preferences implied by these estimates to simulate the effects of a national ethanol content standard. Several appendixes follow.

2.1 Background

2.1.1 Ethanol's role in the fuels market

Ethanol is an alcohol fuel that in the United States derives primarily from corn. Gasoline blenders mix ethanol with gasoline in ratios of up to 10% to boost oxygen content for compliance with federal air quality regulations. Oxygen helps the fuel burn more completely, reducing carbon monoxide emissions in older engines with carburetors. Blenders also add ethanol to produce mid-grade and premium fuels and to satisfy the federal Renewable Fuel Standard and explicit ethanol content requirements in some states. Virtually all gasoline engines are certified to burn fuel blends that contain 10% ethanol or less. While gasoline blenders sometimes mix ethanol with gasoline on a discretionary basis when ethanol prices are low to extend fuel supplies, ethanol's primary role is as a complement to gasoline in the production of retail fuels. Total U.S. consumption of ethanol for gasoline blending was about 5 billion gallons in 2006, or about 3% of gasoline consumption by volume. Ethanol is heavily subsidized, with direct federal and state payments to ethanol producers, a federal tax subsidy of $0.51 per gallon for blenders that mix ethanol with gasoline, and a tariff of $0.54 per gallon that applies to all but a nominal quantity of imports.

The market for ethanol as a direct substitute for gasoline is small but growing rapidly. Stimulated by rising gasoline prices and supported by federal, state, and local subsidies for installing alternative fueling infrastructure, the number of fueling stations that offer E85—an alternative fuel blend of 85% ethanol and 15% gasoline—nearly doubled during 2006–2007 to over 1200 stations nationwide.[1] On the consumer side of the E85 ethanol market, the federal Alternative Motor Fuels Act of 1988 created strong incentives under the Corporate Average Fuel Economy (CAFE) standards program for carmakers with binding CAFE constraints—notably the big-three American auto companies—to produce vehicles capable of burning this fuel blend. Carmakers produced about 5 million of these so-called flexible-fuel vehicles between 2000 and 2006, and production continues apace. Most flexible-fuel vehicles are large cars, pickups, and SUVs.

The federal Renewable Fuels Standard, which Congress established in 2005, originally required that the fuel industry supply a minimum quantity of renewable fuel each year from 2006–2012. Congress extended and expanded the standard in late 2007. The current

[1] I often refer to E85 simply as "ethanol" or, when necessary to avoid confusion, as "E85 ethanol" or "retail ethanol." I distinguish retail E85 ethanol from "denatured ethanol," which is blended with gasoline in the production of retail fuels. Denatured ethanol is nearly pure alcohol but with a small quantity of gasoline or other chemical added, making it unfit for human consumption.

standard now sets a minimum quantity of renewable fuel each year from 2008–2022, increasing gradually from 9 to 36 billion gallons annually. The quantity standard for 2022 is about 25% of current annual gasoline consumption. The standard is not binding as of 2007, given other sources of ethanol demand. Once binding, however, the standard is likely to be met primarily with ethanol. Although the standard mandates a minimum quantity of renewable fuel, the U.S. Environmental Protection Agency (EPA) implements the standard as a percentage of projected fuel consumption. Below I simulate the effects of a 25% ethanol content requirement for gasoline, which corresponds directly to this standard.

Only owners of flexible-fuel vehicles are able to fill their tanks with gasoline blends that contain more than 10% ethanol without voiding their vehicle warranties. Flexible-fuel vehicles have specialized fuel lines and gas tanks that are resistant to the corrosive properties of ethanol. They also contain sensors that can detect the fraction of ethanol in the fuel and make adjustments to account for ethanol's higher oxygen content. These components, which increase vehicle production costs by no more than \$100–\$200, allow flexible-fuel vehicles to burn E85 ethanol, regular gasoline, or any combination of the two. Ethanol has lower energy content than gasoline, implying fewer miles per gallon, but yields similar power and performance. Published government fuel economy estimates indicate that the ratio of gasoline to E85 ethanol fuel economy for flexible-fuel vehicles is about 1.35, or that E85 ethanol delivers about $1-1/1.35 = 25\%$ lower fuel economy.[2] Households that only care about minimizing fuel costs will therefore require that ethanol's price be discounted at least 25% below gasoline before purchasing ethanol.

[2] Using Environmental Protection Agency (EPA) estimates for combined city and highway driving, I calculate the ratio of regular gasoline to E85 ethanol fuel economy for each flexible-fuel vehicle model offered in 2000 through 2006 (U.S. Environmental Protection Agency 2006). EPA did not test vehicles using both fuels until 2000, but relatively few flexible-fuel vehicle models were offered prior to 2000. I calculate the sales-weighted mean ratio using data for nationwide sales of individual flexible-fuel vehicle models from the U.S. Department of Transportation.

2.1.2 Ethanol's environmental and social effects

Replacing one gallon of gasoline with pure corn-based ethanol reduces net petroleum consumption by 0.95 gallons, after accounting for minor upstream petroleum inputs into ethanol production and for ethanol's lower energy content (Farrell et al. 2006). Some policymakers and industry participants also view ethanol as an increasingly important source of refined motor fuel, given the significant local opposition to expanding or siting new petroleum refineries.

Ethanol's climate benefits are less impressive. Corn collects energy from the sun and absorbs carbon dioxide from the atmosphere as it grows, but ethanol production from corn is energy-intensive. The process requires substantial inputs of fertilizer for corn production and heat for ethanol refining. These inputs derive largely from natural gas given current production techniques. After taking into account these upstream inputs of fossil energy, as well as ethanol's lower fuel economy, replacing gasoline with ethanol only reduces net carbon dioxide emissions by 15% (Farrell et al. 2006). Ethanol may even increase emissions after further accounting for direct and indirect land-use changes associated with growing feedstocks (Searchinger et al. 2008; Fargione et al. 2008).

The local air and water quality benefits of ethanol are mixed. Ethanol is generally perceived to be a cleaner burning fuel than gasoline, leading to improved local air quality, although modern pollution control equipment largely obviates these advantages. Substituting ethanol for gasoline reduces tailpipe emissions of benzene, a known human carcinogen, but increases emissions of acetaldehyde, which is also a possible carcinogen, as well as nitrogen oxide, which is a precursor to ozone and smog. Ethanol has the potential to displace environmentally harmful petroleum refining, but expanding corn production for ethanol leads to greater use of fertilizers and pesticides on potentially marginal and environmentally sensitive land.

Finally, some policymakers worry about the role of ethanol production in driving up food prices and the impact on poor households.

Household preferences for ethanol as a gasoline substitute vary considerably. Some households appear to internalize ethanol's perceived benefits. More than half of drivers in a recent nationwide poll expressed interest in owning a flexible-fuel vehicle (Harris Interactive 2006). Of these, nearly 90% were motivated by reducing oil dependence, while nearly two-thirds wanted to reduce greenhouse emissions. Over 90% of drivers in another poll said they would prefer a flexible-fuel vehicle. When asked about ethanol's benefits, they cited "renewable fuel," "clean fuel," "made in America," and "more economical" with roughly equal frequency (Phoenix Automotive 2006). In addition to these external factors, which vary across households, ethanol's relative mileage varies from vehicle to vehicle and across driving scenarios, even in highly controlled government tests. On the road, some households drive primarily in stop-and-go city traffic, while others log a large fraction of highway miles. These and other differences may affect relative mileage. Households with particularly strong preferences for ethanol might be induced to purchase the fuel with less severe distortion of market prices than the average household.

2.1.3 Ethanol production and distribution

As of 2006 there were about 100 ethanol refineries nationwide. Most refineries are located in the corn belt, although a handful of refineries are located outside the midwest.

Fuel suppliers blend ethanol with gasoline in small ratios to produce fuels with particular environmental and performance characteristics, and E85 ethanol accounts for a small fraction of overall ethanol demand. Most blending occurs at fuel blending and distribution terminals, which are located strategically near population centers throughout the United States. Terminal operators blend gasoline, ethanol, and other components into finished fuel products and then distribute the fuels by tanker truck to individual retail fueling stations. A relatively small share of ethanol blending occurs at ethanol refineries that have infrastructure for fuel blending.

Fuel terminals receive most gasoline by pipeline from oil refineries. Existing pipelines are not suitable for transporting ethanol, however, because ethanol can corrode gasoline pipelines, and because water can accumulate at low points. Gasoline repels water, but ethanol does not. Moreover, existing pipelines are configured to move fuel from large oil refineries to dispersed demand centers, whereas ethanol refineries are considerably smaller and usually located in rural areas. In the corn belt, tanker trucks deliver ethanol from ethanol refineries to fuel terminals. Ethanol traveling from the midwest to the coasts usually goes by rail.

Ethanol is readily available for blending in high ratios wherever large quantities of ethanol are blended with gasoline. In Minnesota, for instance, ethanol is available at virtually every fuel terminal any time of year, because Minnesota requires 10% ethanol blending in gasoline year-round. Terminal operators maintain stocks of gasoline and ethanol and sometimes lease storage facilities to retail chains who manage their own fuel stocks. Ethanol is also readily available at a handful of ethanol refineries that have infrastructure for fuel blending. Ethanol retailers in states like Minnesota have no difficulty resupplying on short notice, given ethanol's wide availability for gasoline blending.

2.2 Theoretical model

I develop a model of demand for ethanol as a gasoline substitute. The model formalizes the link between the distribution of household preferences and aggregate prices responses.

2.2.1 The household's problem

For the moment I assume that each household owns a single flexible-fuel vehicle. The household's utility is quasilinear in transportation services $v(\cdot)$ and other goods:

$$v(e+rg)+x, \tag{2.1}$$

where $v(\cdot)$ is strictly increasing and strictly concave, e is consumption of ethanol, g is consumption of regular gasoline, x is consumption of all other goods, and r is the rate at which the household converts gallons of regular gasoline into ethanol-equivalent gallons. Ethanol and gasoline are perfect substitutes. That is, utility is defined over a linear combination of ethanol and gasoline, which I call ethanol-equivalent fuel. When a household cares only about miles traveled the conversion rate r exactly equals the ratio of the household's fuel economy when burning gasoline to its fuel economy when burning ethanol. This ratio will vary across households due to differences in relative mileage. Additionally, some households will value ethanol for its perceived environmental or social benefits, while the relative convenience of filling up with ethanol will vary considerably, given its limited availability. By embodying fuel economy differences and these other factors, r fully summarizes household preferences for ethanol as a gasoline substitute.

The household's budget constraint is given by

$$y-p_e e-p_g g-x=0, \tag{2.2}$$

where p_e and p_g are the prices of ethanol and gasoline, y is income, and I have normalized the price of the composite good to \$1.

Maximization of household utility in (2.1) subject to the budget constraint in (2.2) is characterized by the following first-order conditions:

$$v'(e+rg)=p_e \text{ if } e>0 \tag{2.3}$$

and

$$v'(e+rg) = \frac{p_g}{r} \text{ if } g > 0, \tag{2.4}$$

where I have implicitly assumed that the household spends some but not all of its income on the composite good.[3] The household equates the marginal utility of ethanol-equivalent fuel consumption to the ethanol-equivalent price of whichever fuel it chooses.

Suppose the household purchases both ethanol and gasoline so that $e > 0$ and $g > 0$. Then both first-order conditions (2.3) and (2.4) must hold simultaneously, which implies that $p_e = p_g/r$, or equivalently that $p_g/p_e = r$. That is, if the household purchases both ethanol and gasoline, then the ratio of gasoline to ethanol prices must equal the conversion rate. Otherwise, the household will be at a corner solution and will purchase ethanol exclusively when $p_e < p_g/r$ and gasoline exclusively when $p_g/r < p_e$. That is, the household will choose the fuel with the lower ethanol-equivalent price. For a household that cares only about mileage, this amounts to choosing the fuel that is least costly per mile. Because the conversion rate r equals the relative price where fuel switching occurs, I also refer to it as the fuel-switching price ratio.

While relative prices determine the type of fuel that a household chooses, quantity demanded depends on absolute price levels, with households equating the marginal utility of ethanol-equivalent fuel consumption to the ethanol-equivalent price of whichever fuel they choose. For households that choose ethanol, the optimal quantity of ethanol demanded is given by

$$e^* = d(p_e), \tag{2.5}$$

where I have defined ethanol-equivalent fuel demand as $d(p) \equiv v'^{-1}(p)$ given ethanol-equivalent fuel price p. The quantity of gasoline demanded for households that choose

[3] Assuming that $v'(0) > 1$ guarantees that the household always purchases at least some fuel and does not spend all its income on the composite good. Assuming that y is sufficiently large, so that $v'(y/p_e) < 1$ and $v'(ry/p_g) < 1$, guarantees that the household never expends its full income on fuel and purchases at least some of the composite good.

gasoline is given by

$$g^* = \frac{d(p_g/r)}{r}, \tag{2.6}$$

where I divide by r to convert ethanol-equivalent gallons into nominal gallons of gasoline.

I assume that households that do not own a flexible-fuel vehicle have the same utility as in equation (2.1) but are not able to purchase ethanol. Thus, equation (2.6) is also gasoline demand for households that do not own a flexible-fuel vehicle.

2.2.2 Aggregate demand

Because ethanol and gasoline are perfect substitutes, households that own flexible-fuel vehicles sort into ethanol buyers and gasoline buyers according to their fuel-switching price ratios. While each individual household rests at a corner solution, aggregate demand will be a smooth function of relative prices when fuel-switching price ratios are distributed continuously.

To move formally from individual to aggregate demand, I first assume that there are N households in the market. More precisely, I assume an infinite number of households of measure N. Each household owns a single vehicle. A fraction ϕ of households own flexible-fuel vehicles. I assume that this parameter is set exogenously by policy and that flexible-fuel vehicles are allocated at random. I discuss the validity of this assumption below. I next assume that the fuel-switching price ratio r is distributed among households according to the differentiable cumulative distribution function $H(r)$ defined on $[0,\infty)$. Recall from above that households will choose ethanol if $r < p_g/p_e$, so the fraction of households with flexible-fuel vehicles that choose ethanol is $H(p_g/p_e)$. Finally, I assume that $v(\cdot)$ is identical for all households, which implies that all households that choose ethanol will consume the same quantity.

Given these assumptions, aggregate demand for ethanol as a function of fuel prices is

$$E(p_e, p_g) = N\phi \int_0^{p_g/p_e} d(p_e)dH(r)$$
$$= N\phi H\left(\frac{p_g}{p_e}\right) d(p_e). \tag{2.7}$$

Demand is the total number of households, multiplied by the fraction that own flexible-fuel vehicles, multiplied by the fraction of these that choose ethanol (which depends on relative prices), multiplied by the level of ethanol consumption among households that choose ethanol (which depends on the absolute price of ethanol). Appendix 2.6 provides similar expressions for aggregate gasoline demand and aggregate welfare, which are important for the policy simulation below.

The only source of heterogeneity in the model is r, the price ratio at which a household is indifferent between relying entirely on ethanol or gasoline. I could further allow for heterogeneity in the propensity to drive, say by multiplying ethanol-equivalent fuel demand by a scale parameter that varied across households. Assuming that r and this scale parameter are independent would give the same expression for aggregate demand as in equation (2.7), multiplied by the expected value of the scale parameter. I address the potential correlation of fuel-switching price ratios and propensity to drive in section 2.4 below.

Taking logs on both sides yields logged aggregate ethanol demand:

$$\ln E(p_e, p_g) = \ln N\phi + \ln H\left(\frac{p_g}{p_e}\right) + \ln d(p_e). \tag{2.8}$$

This equation is critical because it relates fuel prices and ethanol quantities to the distribution of household preferences for ethanol as a gasoline substitute.

Differentiating (2.8) with respect to p_g and then multiplying by p_g yields the gasoline-price elasticity of aggregate ethanol demand:

$$\xi_g = \frac{h\left(\frac{p_g}{p_e}\right)}{H\left(\frac{p_g}{p_e}\right)} \frac{p_g}{p_e}, \tag{2.9}$$

where $h(r) \equiv H'(r)$. This cross-price elasticity quantifies the rate at which consumers switch from regular gasoline to ethanol given a percent increase in the price of gasoline. A 1% increase in gasoline prices leads to a ξ_g% increase in the quantity of ethanol demanded. Observe that this elasticity is also the elasticity of the share of households that choose ethanol with respect to the price ratio.

Differentiating (2.8) with respect to p_e and then multiplying by p_e yields the own-price elasticity:

$$\begin{aligned} \xi_e &= \frac{p_e d'(p_e)}{d(p_e)} - \frac{h\left(\frac{p_g}{p_e}\right)}{H\left(\frac{p_g}{p_e}\right)} \frac{p_g}{p_e} \\ &= \xi_f - \frac{h\left(\frac{p_g}{p_e}\right)}{H\left(\frac{p_g}{p_e}\right)} \frac{p_g}{p_e}, \end{aligned} \tag{2.10}$$

where I have defined $\xi_f \equiv pd'(p)/d(p)$. The first term in (2.10), which I refer to as the price elasticity of individual ethanol-equivalent fuel demand, quantifies the rate at which individual households respond to the price increase by curtailing demand. The second term in (2.10), which is identical to the gasoline-price elasticity above in (2.9) multiplied by negative one, quantifies the rate at which households switch from ethanol to gasoline as the price of ethanol increases. Again, this is the elasticity of the share of households that choose ethanol with respect to the price ratio, this time multiplied by negative one. Together these terms imply that a 1% increase in ethanol prices leads to a $-\xi_e$% decrease in the quantity of ethanol demanded.

As an aside, observe that

$$\frac{h\left(\frac{p_g}{p_e}\right)}{H\left(\frac{p_g}{p_e}\right)} \tag{2.11}$$

is the hazard rate for exiting the ethanol market as the price ratio decreases. That is, expression (2.11) gives the instantaneous rate at which households switch to gasoline given a marginal decrease in the price ratio, conditional on choosing ethanol.

The main benefit of the model is that it formalizes the link between the distribution of fuel-switching price ratios and aggregate price responses. Given any distribution of fuel-switching price ratios, equation (2.9) specifies precisely how the cross-price elasticity of demand varies with relative fuel prices. Equation (2.9) therefore provides a means of recovering the distribution of fuel-switching price ratios from observed market behavior.

Figure 2.1 illustrates the relationship between the distribution of fuel-switching price ratios and the gasoline-price elasticity, as expressed in equation (2.9), for four different hypothetical density functions. When household preferences for ethanol as a gasoline substitute are nearly identical, as in figure 2.1(a), fuel-switching behavior is concentrated around a particular price ratio, which leads to a large and highly variable price response in that neighborhood. When household preferences are perfectly homogenous, as previous studies assume, aggregate demand will mirror individual demand: the market will be at a corner solution, with all households choosing the fuel with the lowest ethanol-equivalent price. In figure 2.1(a), this would mean a mass point of individuals at the same fuel-switching price ratio, an infinite price response at that single point, and a zero elasticity everywhere else. This extreme assumption has important implications for policy analysis. If ethanol has relatively high costs, so that no ethanol is consumed in the unregulated equilibrium, large distortions in market prices may be required to induce households to choose ethanol.

Figures 2.1(b)-2.1(c) show that increasingly diffuse preferences for ethanol lead to price elasticities that are smaller in magnitude and less variable. Fuel switching extends over a wide range of relative prices, and demand is not especially sensitive to prices at any

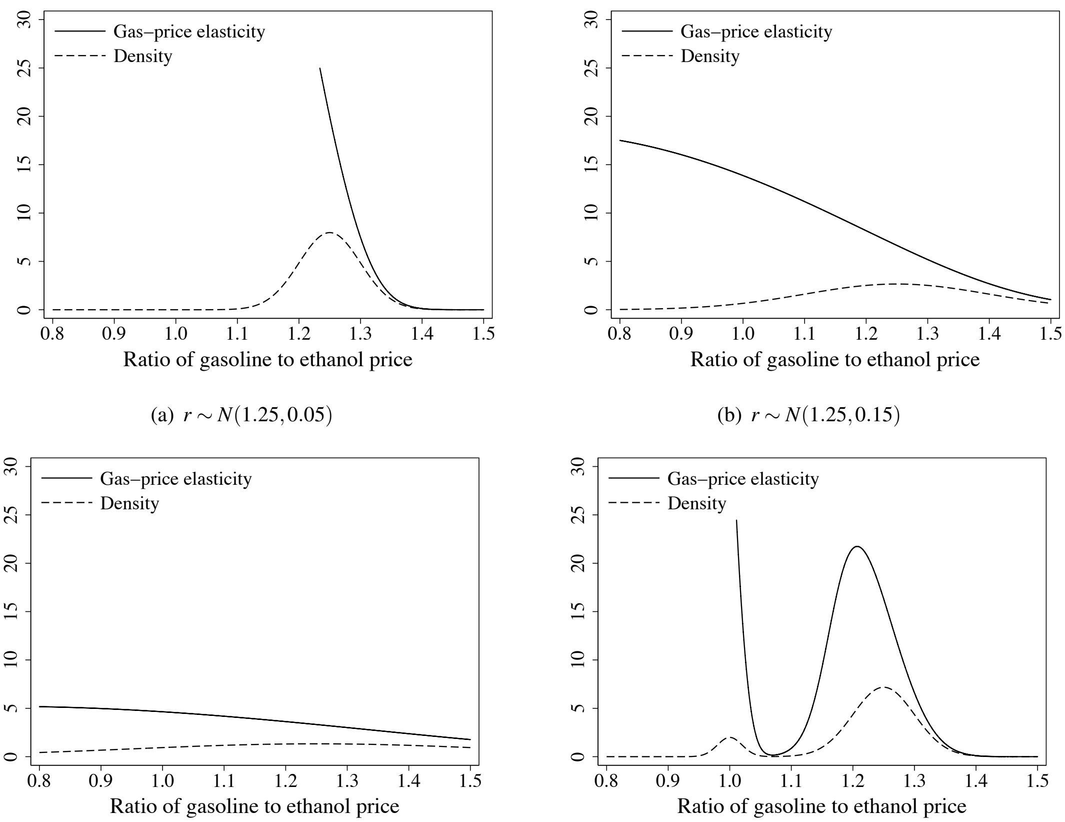

(a) $r \sim N(1.25, 0.05)$

(b) $r \sim N(1.25, 0.15)$

(c) $r \sim N(1.25, 0.30)$

(d) $r \sim 0.1N(1.0, 0.05) + 0.9N(1.25, 0.05)$

Figure 2.1: Hypothetical preferences and elasticity functions

Note: Figure illustrates the relationship between the density function for fuel-switching price ratios and the gasoline-price elasticity of aggregate ethanol demand, as given by equation (2.9), for four hypothetical density functions.

particular point. Households with particularly strong preferences for ethanol represent "low-hanging fruit" that can be induced to purchase the fuel with less severe distortion of market prices. Finally, figure 2.1(d) shows that a distribution with multiple modes can lead to pronounced peaks in the elasticity function.[4] This figure would be consistent with

[4]Some general results are available. Differentiating the gasoline-price elasticity in equation (2.9) gives:

$$\frac{\partial \xi_g(x)}{\partial x} = \frac{(h'x + h)H - xh^2}{H^2},$$

where I have suppressed dependence on x for convenience. At a "peak" or "trough" in the density function, say x^*, which is also an inflection point in the distribution function, the slope of the density function is zero: $h'(x^*) = H''(x^*) = 0$. At such a point, the slope of the elasticity function simplifies to:

$$\frac{\partial \xi_g(x^*)}{\partial x} = \frac{h(H - x^*h)}{H^2},$$

which has the same sign as

$$\frac{H(x^*)}{x^*} - h(x^*)$$

for $x^* > 0$. So the slope of the elasticity function is positive at x^* if the average value of $H(x)$ on $(0, x^*)$ (i.e., the slope of the line some households valuing ethanol primarily for its fuel economy performance and others strongly preferring ethanol for environmental and social reasons.

The model also provides a method for disentangling price responses associated with fuel-switching behavior from price responses associated with overall fuel demand. Adding equations (2.9) and (2.10) demonstrates that the price elasticity of individual ethanol-equivalent fuel demand is simply the sum of the two aggregate elasticities:

$$\xi_f = \xi_e + \xi_g. \tag{2.12}$$

For a precise quantitative interpretation of this elasticity, consider a simultaneous 1% increase in both fuel prices. No fuel switching occurs, because relative prices do not change, but households that choose ethanol reduce their demand by $\xi_f\%$. Equation (2.12) therefore shows how to recover intensive-margin price responses from ethanol demand elasticities, which primarily reflect fuel-switching behavior.

The shape of the elasticity function has important implications for retail pricing behavior. If the elasticity function is highly variable at the profit-maximizing price ratio, the relative price charged by a monopolist ethanol retailer will be unresponsive to marginal costs. This is because variable elasticities coincide with masses of price-sensitive households. The retailer will be reluctant to raise prices when costs increase, so as to avoid driving all consumers to gasoline. At the same time, the retailer will have minimal incentive to reduce prices when costs fall, because lowering prices past the point where households are indifferent will stimulate little additional demand. If elasticities are roughly constant, however, retail prices will be highly sensitive to marginal costs. I make these arguments formally in appendix 2.7 below.

through the origin and $H(x^*)$) exceeds its slope at x^*. One can determine whether the density function has peaks or troughs from sign changes in the slope of the elasticity function.

Table 2.1: Summary statistics

Variable	Mean	Std. Dev.	Min.	Max.
sales volume (gallons)	3250.36	3930.28	6.90	37770.50
retail ethanol price	1.74	0.35	0.74	3.38
retail gasoline price	1.98	0.43	1.10	3.00
retail gasoline / ethanol price	1.14	0.11	0.68	1.86
wholesale ethanol price	1.27	0.56	0.45	3.03
wholesale gasoline price	1.38	0.45	0.44	2.33
wholesale gasoline / ethanol price	1.17	0.33	0.69	2.45
ethanol pump age (months)	32.72	25.11	1.00	110.00
number flexible-fuel vehicles in county	3497.84	5054.92	0.00	24453.00
number ethanol pumps in county	4.45	3.48	1.00	16.00
number gas stations in county	102.10	113.06	4.00	357.00

Note: Table is based on estimation sample of 5027 monthly reports from 237 fueling stations in Minnesota between October 1997 and November 2006. Prices are in 2006 dollars.

2.3 Data and summary statistics

I estimate the model of logged aggregate ethanol demand in equation (2.8) above using monthly data for ethanol prices and sales volumes at a large number of retail fueling stations, gasoline prices near those stations, and a number of ancillary variables. Table 2.1 presents summary statistics for my estimation sample.

2.3.1 Data sources

These data come from several sources. My data for retail ethanol prices and sales volumes come from a Minnesota Department of Commerce (MNDOC) and American Lung Association of Minnesota (ALAMN) monthly survey of retail ethanol stations in Minnesota. Stations that received funding to help defray ethanol infrastructure costs are required to respond, while other stations may participate on a voluntary basis. This requirement is not strongly enforced, however, and stations that are required to report do not always do so. The earliest stations began reporting in October 1997, and the data include records through November 2006. Stations report volume-weighted prices derived from monthly sales volumes and revenues. Retail prices include federal, state, and local fuel

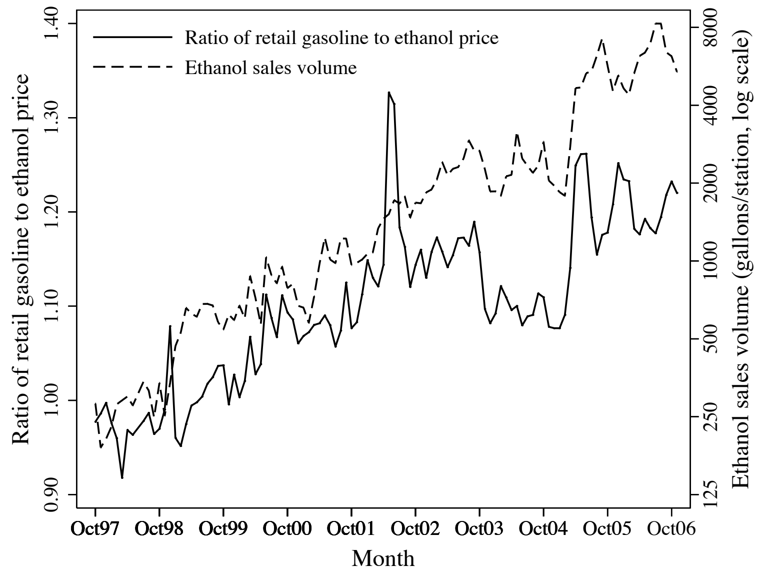

Figure 2.2: Relative retail prices and ethanol sales

Note: Ethanol sales volume is the monthly average volume of ethanol sales among reporting ethanol stations in Minnesota. The ratio of gasoline to ethanol prices is the volume-weighted sample mean price of gasoline divided by the volume-weighted sample mean price of ethanol.

taxes. State and federal fuel taxes did not change during my sample period. The data also record open and close dates for all retail ethanol pumps in Minnesota. I use this information to calculate the total number of stations operating retail ethanol pumps in each county in each month and the length of time that each pump has been operating. While the data do not reveal precise station locations or other identifying information, they do indicate the counties where pumps are located. I match these retail ethanol data to county average retail prices for regular gasoline from Oil Price Information Service (OPIS). I convert all prices to real 2006 prices using the monthly consumer price index from the U.S. Department of Labor.

Figure 2.2 plots relative retail prices from October 1997 through November 2006. Relative prices vary considerably over the sample period, with the relative price of gasoline trending upward. Average ethanol sales also increase steadily over time. The relationship is not necessarily causal, however, as the increase in sales volume is also consistent with a growing stock of flexible-fuel vehicles. I am careful in my estimation to control explicitly for flexible-fuel vehicles and secular trends in fuel demand. Short-run increases in the relative price of gasoline correlate with contemporaneous increases in ethanol sales volumes, which is perhaps more suggestive of a price response. OLS estimates of this relationship are potentially biased, however, if unmodeled shifts in demand cause fuel prices to change. I discuss identification of demand parameters in greater detail below.

As a measure of underlying fuel costs, I obtain wholesale ethanol price data from a trade publication called *Ethanol and Biodiesel News* (previously known as *Renewable Fuels News* and *Oxy-Fuel News* before that). These data measure weekly spot prices at fuel terminals for denatured ethanol in Minneapolis and Fargo. I assign wholesale prices to stations based on proximity to these cities. About four-fifths of stations are located in counties nearest to Minneapolis. I calculate the monthly average of these weekly prices and then subtract the federal ethanol blending subsidy, which fell from \$0.54 per gallon to \$0.51 per gallon during my study period. I obtain wholesale gasoline price data come from the U.S. Energy Information Administration (EIA). These data measure the volume-weighted monthly average spot price in Minnesota. Although wholesale spot price data are available for additional Minnesota cities at a substantial cost from proprietary sources, in practice these prices track each other very closely (Minnesota Department of Agriculture 2003). I use these wholesale price variables to interpret retail pricing behavior.

In addition to these price variables, I obtain data on flexible-fuel vehicle registrations from the Minnesota Department of Public Safety Division of Driver and Vehicle Services. These data record vehicle identification numbers (VINs), original sales dates, and owner zip codes for all vehicles registered in Minnesota as of the summer of 2007. I identify 154,000 flexible-fuel vehicles in the database by cross-referencing VINs with lists of flexible-fuel vehicle models and VIN identifiers from the National Ethanol Vehicle Coali-

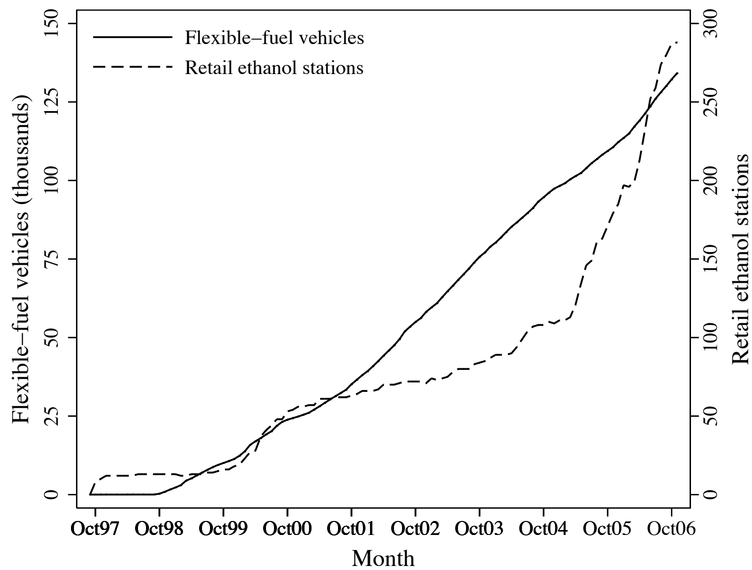

Figure 2.3: Flexible-fuel vehicles and retail ethanol stations

Note: Figure shows growth in number of retail ethanol stations and stock of flexible-fuel vehicles in Minnesota.

tion and from a private firm that collects data on the auto industry. These vehicles represent about 3.3% of the 4.6 million light-duty vehicles registered in Minnesota in 2007. I then use original sales dates to reconstruct a monthly time series for the stock of flexible-fuel vehicles in each county.[5]

Figure 2.3 charts the growth in retail ethanol stations and Minnesota's stock of flexible-fuel vehicles during the study period. The flexible-fuel vehicle stock grows at a roughly constant rate during the sample period, which is consistent with CAFE standards that generated strong incentives for some manufacturers to produce a limited number of flexible-fuel vehicles each year. Growth in the number of retail ethanol stations accelerated in 2000, when ALAMN negotiated an agreement with a particular retail chain to subsidize ethanol pumps at a large number of its stations. Growth accelerated again in 2004-2005.

[5]I am unable to determine whether some vehicles are flexible-fuel vehicles due to missing or invalid VINs, and a relatively small number of flexible-fuel vehicles are excluded due to missing sales dates or zip codes outside Minnesota. I also unable to account for vehicle attrition or historical movements of vehicles in and out of Minnesota and across county lines prior to 2007. Owner addresses also might differ from counties where flexible-fuel vehicles are actually driven. For these various reasons I measure flexible-fuel vehicle stocks with some error.

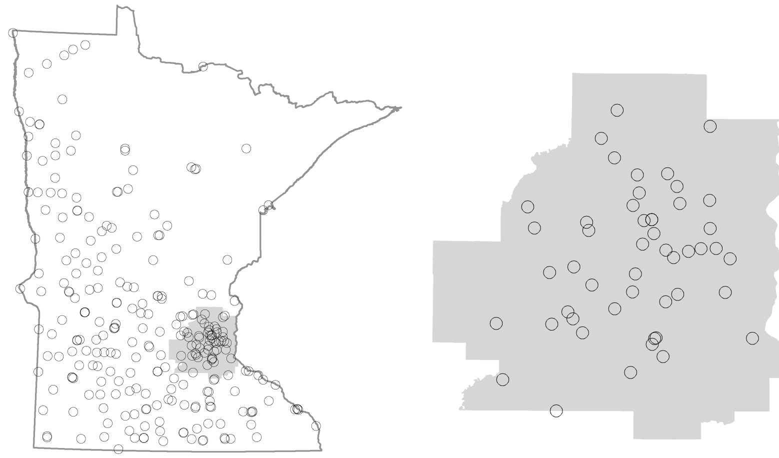

Figure 2.4: Minnesota's retail ethanol stations

Note: Figure shows locations of Minnesota's 264 retail ethanol fueling stations as of August 2006. Minnesota measures 400 miles from north to south and about 250 miles along its southern border. The shaded region is the seven-county metropolitan area of Minneapolis and St. Paul, which measures 65 miles north to south and 60 miles east to west.

High gasoline prices and low wholesale ethanol costs may have contributed to this accelerated growth.

As I note above, I calculate the total number of retail ethanol stations in each county in each month in order to quantify variation in competition. Figure 2.4 maps the locations of all 264 retail ethanol stations in Minnesota as of August 2006 based on a separate list of station addresses from MNDOC. For comparison, I also calculate the total number of retail gasoline stations operating in each Minnesota county in 2006 based on station address information from the Minnesota Department of Commerce Weights and Measures Division. Table 2.1, which assumes the same number of gas stations operating in each county for 1997-2006, shows that there are more than 20 gasoline stations for every ethanol station on average in my sample.[6] While competition in fuel markets is fierce, most ethanol retailers operate as local monopolists in the narrower retail ethanol market.

My analysis covers the time period from October 1997 through November 2006. Dur-

[6]The actual ratio is probably even higher. Although most of the nearly 2900 individual stations operating in 2006 were also operating during 1997-2005 (Buccelli 2007), the total number of retail stations statewide declined about 7% from 1997–2006 (National Petroleum News 2006).

ing this time the number of retail ethanol stations in Minnesota grew from less than 10 to nearly 300. Based on reported open and close dates, there were a total of nearly 9000 potential monthly observations at these stations. Approximately 56% of these potential observations are covered by the MNDOC/ALAMN survey. The remaining 44% are missing, reflecting both stations that never participate in the survey, as well as stations that fail to report in some months. This results in an estimation sample of 5027 observations at 237 stations, implying an average panel size of about 21 months. Some stations operate nearly the entire study period, while others operate for just a few months, as is clear from figure 2.3.

My retail ethanol survey data are subject to several potential layers of sample selection. First, not all stations participate in the survey, and not all participating stations report every month. Most stations do participate, however, and the probability that participating stations report is not correlated with fuel prices. Second, fueling station owners might locate ethanol pumps in areas where preferences for ethanol are strongest. Stations are spread throughout Minnesota, however, covering every major region of the state except the northeast, which is sparsely populated. Finally, sales data reflect the behavior of households that self-selected to own flexible-fuel vehicles and participate in the nascent ethanol market. Selection on price is not problematic, as my analysis will control explicitly for fuel prices. Selection on familiarity with ethanol is also not cause for concern, at least in recent years. Most drivers know whether they own a flexible-fuel vehicle, and most flexible-fuel drivers are well-informed about ethanol and its availability in their area (Phoenix Automotive 2006).

Potentially more problematic is the possibility that flexible-fuel drivers have systematically different preferences than other drivers. Flexible-fuel vehicle owners tend to buy American. Survey evidence indicates that flexible-fuel owners are also more likely to con-

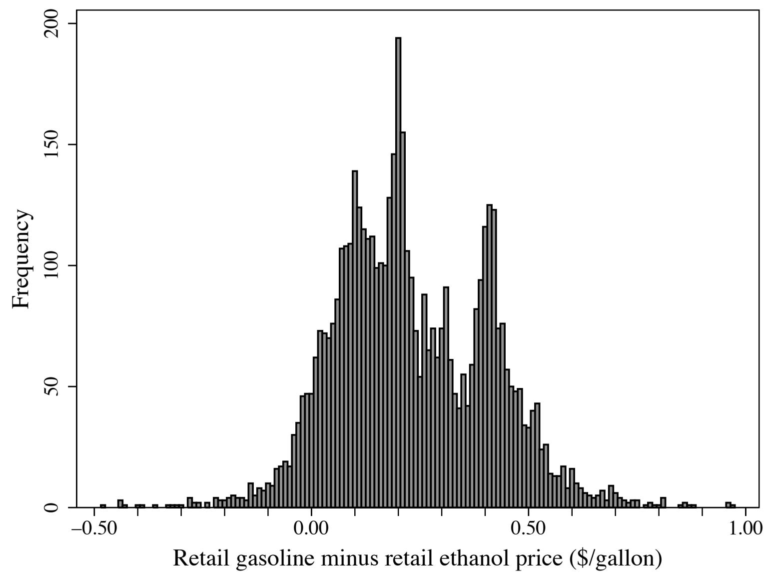

Figure 2.5: Distribution of retail price discount

Note: Figure shows the distribution of ethanol's nominal retail price discount relative to gasoline in the estimation sample.

sider minivans and pickups for their next vehicle purchase, while other drivers are more likely to consider smaller cars (Phoenix Automotive 2006). Furthermore, most flexible-fuel vehicles have identical gasoline-only counterparts, which could lead to sorting directly on flexible-fuel capability. On the other hand, production of flexible-fuel vehicles derives almost entirely from federal fuel-economy regulations. Carmakers sell flexible-fuel vehicles all over the country, even in areas without ethanol, and anecdotally offer flexible-fuel vehicles for the same prices as their gasoline-only counterparts. So sorting on flexible-fuel capability is not necessarily a major problem.

2.3.2 Retail pricing behavior

I spoke with industry representatives and inspected retail pricing behavior closely to identify price variation that is arguably exogenous to demand.[7]

[7] I spoke with representatives from the largest chains in Minnesota that offer retail ethanol, as well as several independently owned and operated stations, representatives from two ethanol refiners that directly supply about one-third of retail ethanol stations in Minnesota, several ethanol industry analysts, and the administrators of the MNDOC/ALAMN survey.

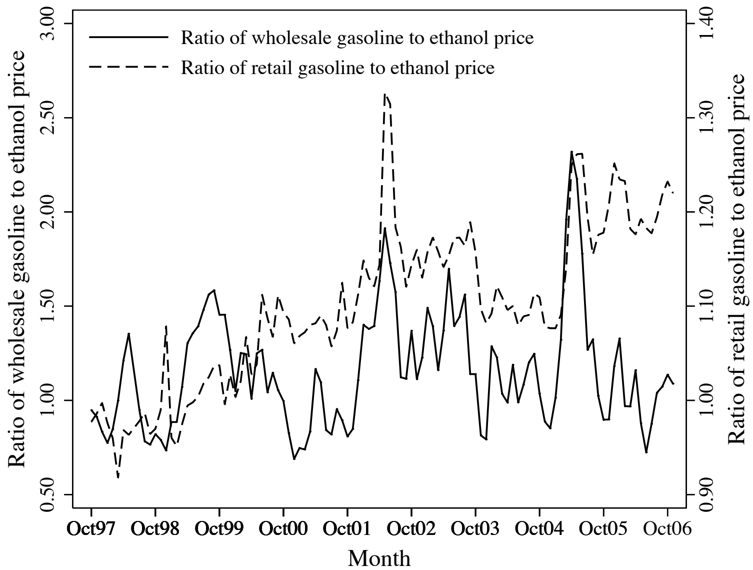

Figure 2.6: Relative wholesale prices and relative retail fuel prices

Note: The ratios of gasoline to ethanol prices are the volume-weighted sample mean prices of gasoline divided by the volume-weighted sample mean prices of ethanol.

Retailers generally price ethanol at a discount to regular gasoline in nominal increments of $0.10 per gallon. This behavior is manifest in figure 2.5, which plots the distribution of nominal price discounts in my sample, and is consistent with how retailers tell me they determine prices. Clustering near salient discounts would be even more pronounced if I had station-level gasoline prices. The discounts in the figure are based on county average gasoline prices.

The sizes of these discounts are governed by broad market conditions. Average discounts generally increase when wholesale ethanol prices fall relative to gasoline, and discounts decrease when ethanol prices rise, as is evident in figure 2.6.

While average discounts appear to respond immediately to changes in underlying costs, industry representatives I spoke with indicated that discounts at individual stations often persist at the same level for months or years at a time. Retailers update discounts infrequently to adjust for broad shifts in their relative fuel costs. Price discounts are not

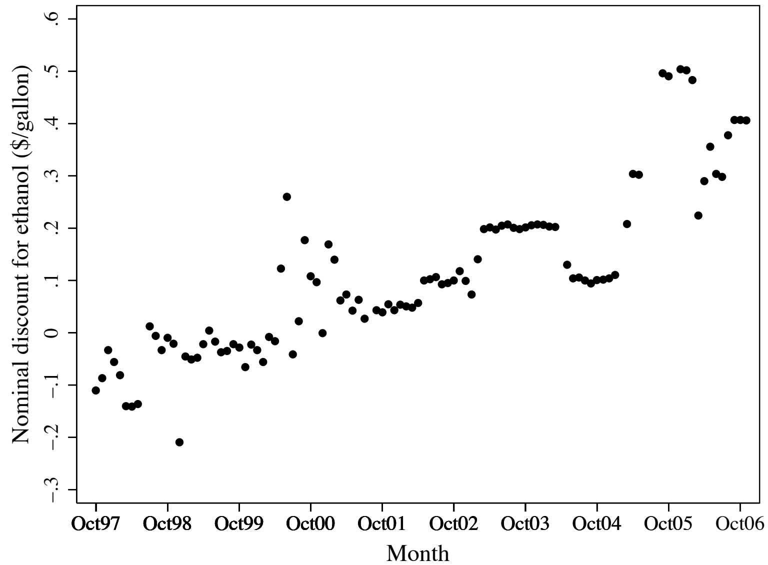

Figure 2.7: Example of one station's nominal price discount

Note: Figure shows ethanol's nominal price discount relative to gasoline for a particular ethanol retailer over time.

adjusted on a daily basis, nor are they adjusted in response to shifts in ethanol-specific demand. This behavior is evident in figure 2.7, which plots the nominal price discount over time for one particular station in the sample. This station has been operating longer than most but its behavior is not atypical. This pricing behavior facilitates identification of demand parameters. Because retailers maintain the same discounts for extended periods of time, unmodeled shifts in ethanol demand will tend not to translate to changes in relative fuel prices.

Changes in market spot prices affect retailers differently, depending on the relationships they have with suppliers. First, some retailers buy ethanol directly from ethanol refineries, while others purchase ethanol at fuel terminals. Second, some retailers purchase ethanol on the spot market and bear the full brunt of variation in spot prices, but many retailers have long-term contracts for wholesale ethanol. Contract prices are less variable than observed spot prices, which explains why large fluctuations in wholesale prices in figure 2.6

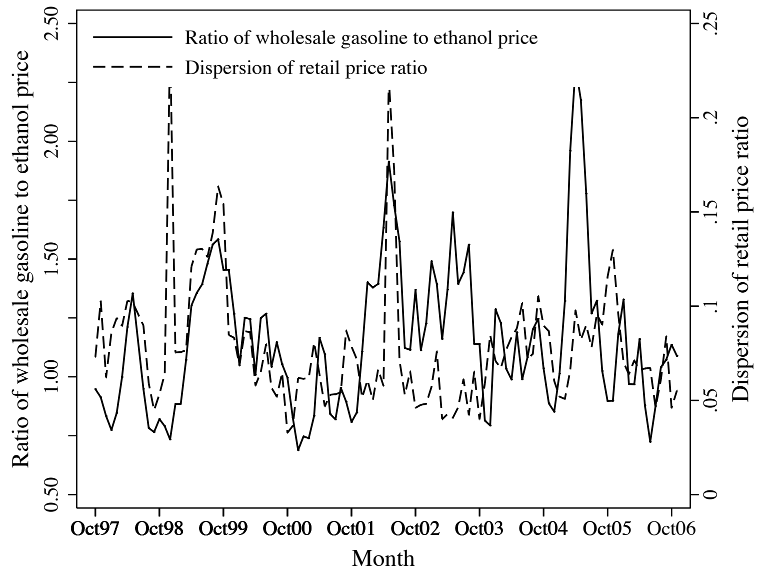

Figure 2.8: Dispersion of relative retail prices

Note: Dispersion of retail price ratio is the monthly standard deviation of the OLS residuals from the retail price ratio regressed on a vector of month dummies and station fixed effects. This variable quantifies differential changes in relative prices across stations.

correlate with comparatively small changes in retail prices. See appendix 2.8 for more detail on this issue and an extended discussion about determinants of wholesale prices. Finally, contracts employ different pricing formulae. Some contracts peg ethanol costs directly to the spot price of gasoline, while others tie costs to some average of gasoline and ethanol. These formulae vary across stations and over time. This variation in supply relationships and contracts leads to cross-sectional variation in retail pricing behavior. Retailers also may face varying degrees of competition, leading to differential pass through of wholesale ethanol costs. See appendix 2.9 for further details.

Variation in pricing behavior is important because it allows me to control for month effects that are common to all stations and still identify demand parameters off of differential changes in fuel prices across stations. To quantify this variation, I regressed relative retail prices on a vector of month dummy variables and station fixed effects. I then computed the standard deviation of the residuals from this regression within each month. I refer to this standard deviation as the dispersion of relative prices. Figure 2.8 shows that price dispersion increases when gasoline spot prices are high relative to ethanol. This behavior is consistent with the different supply relationships that I document. Some stations have ethanol contracts that tie wholesale ethanol costs to gasoline, while other stations purchase ethanol on the spot market. Variation in costs therefore increases whenever spot prices diverge. This behavior is also consistent with differences in competition. When ethanol costs are low relative to gasoline, competitive retailers are forced to reduce prices, while less competitive retailers are able to price closer to gasoline. Price dispersion decreases when the gap between ethanol costs and gasoline prices narrows.

2.4 Econometric estimation and results

2.4.1 Econometric model

I estimate logged aggregate ethanol demand of the following form:

$$\ln E_{it} = \alpha \ln p_{e_{it}} + F\left(\ln \frac{p_{g_{it}}}{p_{e_{it}}}\right) + \beta' X_{it} + \gamma_t + \delta_i + \omega_i(t) + \varepsilon_{it}, \tag{2.13}$$

where: E_{it} is the volume of ethanol sales at fueling station i in month t; $p_{e_{it}}$ is the retail price of ethanol and $p_{g_{it}}$ is the retail price of regular gasoline; X_{it} is a vector of time-varying county and station characteristics; γ_t is a month effect that is constant across all fueling stations; δ_i is a fueling station effect that is constant across all time periods; $\omega_i(t)$ is a station-specific quadratic time trend; ε_{it} is an error term; and the remaining elements are coefficients, vectors of coefficients, and functions to be estimated. Note that equation (2.13) is the empirical analog of logged aggregate demand in theoretical equation (2.8) above.

My theoretical model implies that elasticities may vary dramatically with relative prices. I allow for variable elasticities using several approaches. My main estimates use a flexible polynomial approximation for $F(\cdot)$. I also estimate the model using a cubic spline and semi-parametric approximation of $F(\cdot)$.

The own-price elasticity of ethanol demand is simply $\alpha - F'(\ln p_g/p_e)$. The gasoline price elasticity is $F'(\ln p_g/p_e)$, which is equivalent to the elasticity of the share of households that choose ethanol with respect to relative prices. Following equation (2.12), the price elasticity of individual ethanol-equivalent fuel demand is the sum of the gasoline-price and own-price elasticities, which simplifies here to α. The functional form in (2.13) therefore generates constant-elasticity estimates for the response of overall fuel demand to changes in fuel price levels.

The fueling station effect δ_i controls for persistent differences in fueling station characteristics, such as brand name, location, and amenities. The station effect also controls for persistent determinants of local fuel demand, including household income and other demographics, driving habits, and fuel economy performance. The month dummy variables given by γ_t control for secular trends in demand due to growing awareness of flexible-fuel vehicle capabilities or rising state income levels. The station-specific quadratic time trends $\omega_i(t)$ control for similar factors that evolve at different rates locally. Finally, the month dummies control for potential seasonality in demand, including the well-known surge in driving that occurs each summer.[8]

The vector of time-varying station characteristics X_{it} includes the log of the county's flexible-fuel vehicle stock. The vector also includes the log of the total number of stations that offer retail ethanol in the same county. While a negative coefficient would imply that new stations draw customers away from existing stations, a zero coefficient might only suggest that new stations locate where competition is weak. This measure of competition reflects retailer choices about when and where to install ethanol pumps, and these

[8]The minimum denatured ethanol content of retail ethanol in Minnesota varies seasonally due to cold weather starting issues, ranging from 70% in the winter to 79% in the summer (U.S. Department of Energy 2006). Although the month dummies control for seasonality in the level of demand, they do not control for potential seasonality in the price elasticity of demand due to variation in denatured ethanol content. Variation in ethanol content is relatively minor, however, and unlikely to be transparent to consumers, making it neither problematic nor useful for identification.

decisions presumably depend critically on the locations of existing pumps. Table 2.1 indicates that there are less than 5 retail ethanol stations per county, while there are more than twenty times as many gasoline stations. A finding of significant competition in retail ethanol markets would therefore be surprising. Finally, the vector of time-varying station characteristics includes dummy variables that indicate the length of time that a station has been offering ethanol. These dummy variables differ from the month dummy variables because start dates vary from station to station. Sales will likely be low after a station first opens before customers are fully aware of the new opportunity to purchase ethanol.

2.4.2 Identification

I estimate equation (2.13) using OLS. OLS estimates are potentially biased if unmodeled shifts in ethanol demand correlate with fuel prices. This is a standard endogeneity problem in estimating demand functions. Shifts in ethanol-specific demand would tend to bias estimates of the own-price elasticity toward zero, if such shifts led to higher ethanol prices. In contrast, shifts in overall fuel demand would increase prices for all fuels, in which case relative fuel prices might arguably be exogenous. This would facilitate identification because I am primarily interested in fuel-switching behavior, which depends on relative prices. Endogeneous fuel price levels would nevertheless bias the OLS estimate for the price elasticity of individual ethanol-equivalent fuel demand.

I argue that the price variation I quantify in figure 2.8 is largely orthogonal to the error term in equation (2.13). The station owners I spoke with indicated that they do not update retail ethanol prices in response to local short-term demand shifts. Rather, station owners price ethanol at nominal discounts to regular gasoline in increments of \$0.10 per gallon, often maintain these same discounts for extended periods of time, and only adjust discounts in response to changes in underlying fuel costs. This behavior largely rules out ethanol-specific demand shifts at individual stations being correlated with station-level price changes and biasing OLS estimates.[9]

Changes in underlying fuel costs could still be endogeneous to local demand shifts, however, if such shifts were correlated across many stations (Kennan 1989). A classic example is the surge in summer travel demand that drives up fuel prices. I control for these and other correlated demand shifts using month dummy variables. While these controls throw away potentially useful time-series variation in retail fuel prices, I am able to document a variety of relationships between retail ethanol stations and their wholesale suppliers, which lead to cross-sectional variation in pricing behavior. Figure 2.8 demonstrates that there is substantial variation in relative fuel prices, even after controlling for month and station effects.

Finally, I control for localized demand shifts using station-specific quadratic trends. In using this approach I implicitly assume that local demand shifts evolve slowly relative to shifts in supply. This assumption appears to have rigorous empirical justification, at least in the case of the world oil market (see Kilian (2007)).

2.4.3 Estimation results

Polynomial results

Table 2.2 presents my OLS estimation results based on polynomial approximations for fuel-switching behavior. I control for station effects using fixed-effects and first-difference estimators, which may have different efficiency properties. All logged price variables have been normalized to equal zero at sample mean prices. This allows me to interpret each coefficient estimate in the second row directly as the gasoline-price elasticity of ethanol demand evaluated at sample mean prices.[10]

---

[9]There may be theoretical justification for retailers being unresponsive to local demand shifts when setting relative prices. For a monopolist ethanol retailer, relative retail prices will be invariant to demand shifts that enter multiplicatively by scaling aggregate demand. This is because multiplicative demand shifts do not alter the shape of the own-price elasticity function, leaving the monopolist's first-order pricing condition unchanged. See appendix 2.7 below for benchmark models of retail pricing behavior.

[10]All price variables have been normalized according to:

$$p^* = p/\bar{p},$$

Table 2.2: Main estimation results

Variable	Fixed effects linear	Fixed effects quadratic	Fixed effects cubic	First differences linear	First differences quadratic	First differences cubic
ln(price ethanol)	0.48	0.41	0.35	0.77	0.66	0.64
	(0.64)	(0.66)	(0.64)	(0.48)	(0.48)	(0.47)
ln(gas price / ethanol price)^1	2.60	2.63	3.04	2.41	2.45	2.58
	(0.64)	(0.65)	(0.67)	(0.56)	(0.53)	(0.55)
ln(gas price / ethanol price)^2		-2.58	-2.63		-3.19	-3.22
		(1.06)	(0.86)		(0.91)	(0.86)
ln(gas price / ethanol price)^3			-6.38			-1.97
			(3.73)			(3.65)
ln(number flex-fuel vehicles)	0.07	0.08	0.08	0.08	0.08	0.08
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
ln(number ethanol stations)	-0.14	-0.13	-0.13	-0.13	-0.12	-0.12
	(0.11)	(0.11)	(0.11)	(0.08)	(0.08)	(0.08)
month 1 of operation	-1.20	-1.18	-1.18	-0.99	-0.95	-0.95
	(0.19)	(0.20)	(0.20)	(0.21)	(0.21)	(0.21)
month 2 of operation	-0.28	-0.28	-0.27	-0.14	-0.13	-0.12
	(0.15)	(0.15)	(0.15)	(0.16)	(0.16)	(0.16)
month 3 of operation	-0.14	-0.13	-0.13	-0.00	0.00	0.00
	(0.11)	(0.11)	(0.11)	(0.11)	(0.11)	(0.11)
month 4 of operation	-0.04	-0.04	-0.04	0.04	0.04	0.04
	(0.07)	(0.07)	(0.07)	(0.06)	(0.06)	(0.06)
number observations	5027	5027	5027	4332	4332	4332
number of stations	237	237	237	201	201	201
R-squared	0.84	0.84	0.84	0.34	0.36	0.36
mean gasoline-price elasticity	2.60	2.65	2.88	2.41	2.48	2.55
of ethanol demand	(0.64)	(0.65)	(0.64)	(0.56)	(0.53)	(0.53)
1st-order correlation of residuals	0.39	0.40	0.39	-0.27	-0.27	-0.27
	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)

Note: Dependent variable is logged monthly ethanol sales volume in gallons. Logged price variables have been normalized to equal zero at sample mean prices. Standard errors in parentheses are clustered by station. Fixed-effects regressions use deviations from within-station means to control for station effects. First-difference regressions do not include a constant term, as this term is removed by first differencing. All regressions include month dummy variables and station-specific quadratic time trends. Mean gasoline-price elasticity is the sample mean predicted elasticity. First-order correlation of residuals is the coefficient from the least-squares regression of residuals on their lagged values. See text for details.

Ethanol demand is sensitive to price changes. The coefficient on logged relative prices for fixed-effects estimation of the linear model implies that the gasoline-price elasticity of aggregate ethanol demand is 2.60. This elasticity is equivalent to the elasticity of the share of households that choose ethanol with respect to relative prices. The same coefficient is 2.63 for fixed-effects estimation of the quadratic model and 3.04 for the cubic model. These are gasoline-price elasticities evaluated at mean prices. Sample mean elasticities near the bottom of table 2.2 confirm that average price responses are somewhat smaller in

where p^* is the normalized price variable I use for estimation, p is the price variable prior to normalization, and $\bar{p}$ is the sample mean of p. The normalized prices equal unity at mean prices, while their logged values equal zero. Coefficients on the quadratic and higher-order terms therefore drop out of the elasticity calculations at sample mean prices.

the linear model, which imposes constant elasticities. These results suggest that it may be important to allow elasticities to vary with relative prices.

Mean elasticity estimates are 0.2-0.3 smaller when using first-difference estimation to control for station effects. There are at least two plausible explanations. First, if demand does not respond fully to changes in relative fuel prices within the first month, the estimators may give different results. First-difference estimates exploit differential variation in relative fuel prices in adjacent time periods only, while fixed-effects estimates relate average sales volumes to relative fuel prices in all periods. For this reason fixed-effects estimates may be more robust to delayed responses. I return to this issue below. Second, I calculate relative fuel prices based on county average gasoline prices. While it is unclear that a different level of aggregation is more appropriate, measurement error in relative prices will tend to bias elasticity estimates toward zero. This bias is usually more severe in first-difference estimates (Griliches and Hausman 1986).

Standard errors in table 2.2 are robust to heteroskedasticity and serial correlation. While the first-difference estimates for the price responses have narrower confidence intervals than the fixed-effects estimates, neither estimator is fully efficient. First-order serial correlation in the fixed-effects residuals is 0.39–0.40 and statistically different from zero. First-order serial correlation in the first-difference residuals is -0.27. This coefficient is statistically different from zero, which indicates that the first-difference estimates are not efficient. This coefficient is also statistically different from -0.5, which confirms the inference based on the fixed-effects residuals that the model's errors are serially correlated (Wooldridge 2002). That $-0.27 \approx -(1-0.39)/2$ is consistent with the model's errors following an AR(1) process (Solon 1984).

The coefficients on the logged price of ethanol in the first row give an estimate of the constant price elasticity of individual ethanol-equivalent fuel demand. The fixed-effects estimates imply that this elasticity is 0.35–0.48. These estimates have the incorrect sign but are statistically indistinguishable from the range of estimates in the literature. The first-difference estimates imply an ethanol-equivalent elasticity of 0.64–0.77, again with the incorrect sign. I am still unable to rule out equivalence with previous estimates, however, even with narrower confidence intervals.

There are several plausible explanations for why these point estimates have the incorrect signs. First, my theoretical model implicitly assumes that fuel-switching price ratios and the propensity to consume fuel are uncorrelated. In fact, households that require a larger discount before purchasing ethanol might also drive smaller cars or log fewer miles. This negative correlation between fuel-switching price ratios and fuel demand could generate a positive coefficient on the logged price of ethanol. When the price of ethanol increases, conserving households that dislike ethanol would be the first to switch to gasoline, leaving only fuel-guzzling households in the ethanol market.

Second, the positive coefficient on the logged price of ethanol indicates that a proportional increase in all fuel prices generates additional ethanol demand. Higher overall fuel prices might provide an additional boost for ethanol, if households are frustrated with large oil companies, feel a greater sense of social responsibility to reduce oil consumption, or devote more time to learning about alternative fueling opportunities.

Third, some consumers may respond to linear differences in fuel prices as opposed to differences in relative prices. When the price of gasoline exceeds the price of ethanol, a proportional increase in fuel prices leaves relative prices unchanged, but the absolute increase in gasoline prices is larger. I investigated this issue by adding a variable measuring the linear difference between retail gasoline and ethanol prices to regressions that also included logged relative prices. The coefficient estimates had the incorrect signs and were highly insignificant for both fixed-effects and first-difference estimation, however, suggesting that the model is correctly specified in relative fuel prices (see column 8 of tables 2.5 and 2.6 in appendix 2.12 below).

Fourth, fuel-switching behavior in the short term may be more sensitive to gasoline prices than ethanol prices. Gasoline pumps and price postings are ubiquitous, while ethanol is limited to several hundred stations statewide or less. I tested for delayed price effects by adding one-period, two-period, and three-period lagged price variables to the linear polynomial regressions (see table 2.7 in appendix 2.12 below). The cumulative elasticity of ethanol-equivalent fuel demand continued to have the incorrect sign, but it was smaller in magnitude in the one-period and two-period distributed lag models. The cumulative response to relative prices exceeded the initial response, but not dramatically so. The difference in magnitude between the initial and cumulative response was also larger for the first-difference estimates, which is consistent with the intuition above that the fixed-effects estimates are more robust to delayed responses. None of the cumulative responses was close to being statistically different from the corresponding initial response.

Finally, the incorrect signs may be related to measurement error and bias in the fuel-switching price responses, as I discuss above.

Interpreting the quadratic and cubic terms in table 2.2 is less straightforward. The coefficient on the quadratic term is negative in both the quadratic and cubic regressions, implying that the gasoline-price elasticity function is downward sloping at mean prices. The coefficient on the cubic term for fixed-effects estimation of the cubic model is negative, implying that the gasoline-price elasticity function is convex in logged relative prices. Equivalently, the own-price elasticity is concave. Elasticities appear to be linear functions of logged relative prices based on the first-difference estimates, where the cubic term is small in magnitude and statistically insignificant.

To examine the combined effects of the higher-order terms in the cubic models, I use

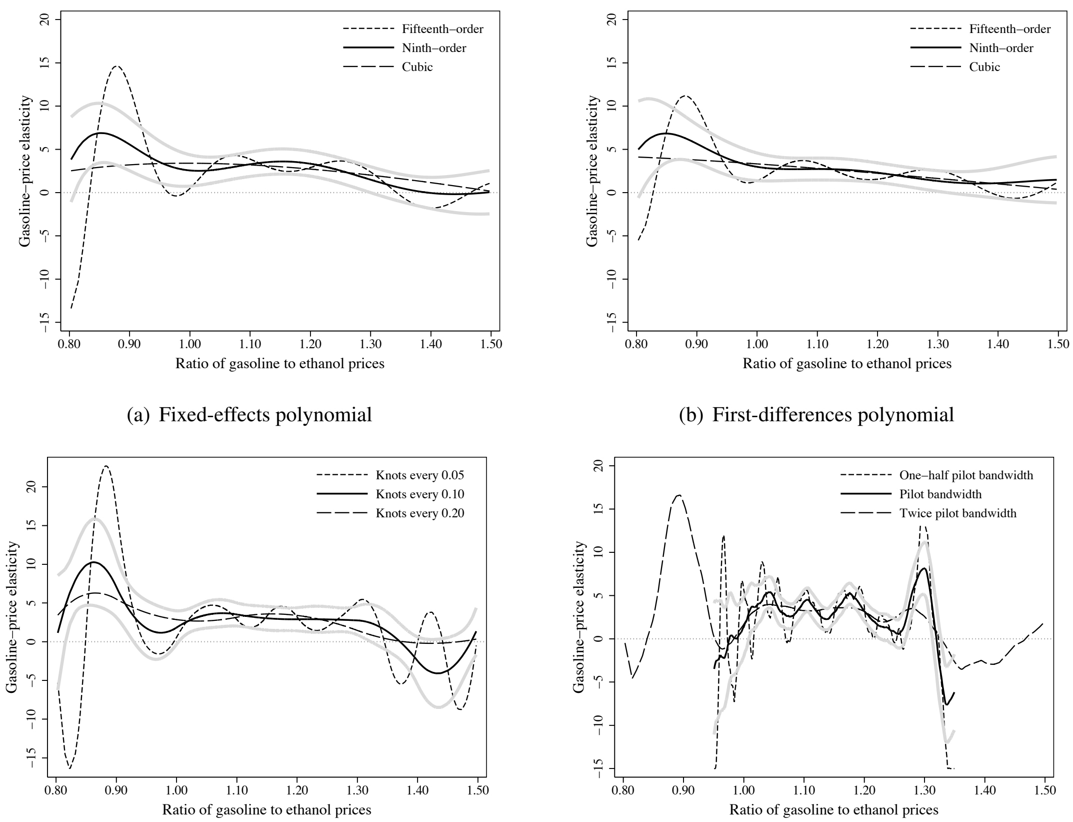

(a) Fixed-effects polynomial

(b) First-differences polynomial

(c) Cubic-spline

(d) Semi-parametric

Figure 2.9: Estimated gasoline-price elasticities

Note: Figure shows gasoline-price elasticity of aggregate ethanol demand based on fixed-effects polynomial, first-differences polynomial, cubic spline, and semi-parametric estimation. Pilot bandwidth is 0.03 for semi-parametric estimates. Solid gray lines are 95% confidence intervals for estimates based on ninth-order polynomials, cubic spline with knots every 0.10, and semi-parametric estimates with pilot bandwidth 0.03. Confidence intervals for semi-parametric estimates are based on standard errors from local polynomial regressions. Semi-parametric estimates using the two narrowest bandwidths are hidden in the extremes of the data, as these estimates were fluctuating wildly outside the range of the figure. See text for details.

the coefficient estimates in table 2.2 to calculate the gasoline-price elasticity of aggregate ethanol demand as a function of relative prices. Figures 2.9(a)–(b) plot this elasticity function for the fixed-effects and first-difference estimates. Elasticity functions based on the cubic model decline slightly in magnitude as the ratio of gasoline to ethanol prices increases.

The cubic model would have difficulty revealing sharp peaks in the elasticity function, so I also estimated the model using more flexible polynomial approximations. Figures

2.9(a)–(b) plot estimates of the elasticity function based on a ninth-order polynomial, as well as a fifteenth-order polynomial, which minimized the Bayesian Information Criterion (BIC) for fixed-effects estimation. A quadratic model minimized the BIC for first-difference estimation. Elasticities based on the higher-order polynomial approximations reveal additional non-linearities but are not statistically different from the more restrictive cubic estimates. Sample mean elasticities for the more flexible fixed-effects estimates increase moderately to about 3.0. Mean elasticities for the more flexible first-difference estimates remain at about 2.5.

Coefficient estimates for price effects were robust to replacing station-specific quadratic time trends with county-specific trends, excluding time trends altogether, and dropping ancillary station and county control variables. Excluding month dummy variables increased the ethanol-equivalent price elasticity in the wrong direction, however, and led to a large increase in the fuel-switching price response (see tables 2.5 and 2.6 in appendix 2.12 below). The month dummies are important.

Returning to the estimates in table 2.2, the coefficients on flexible-fuel vehicle stocks indicate that doubling the number of vehicles leads to a 7%-8% increase in ethanol sales volumes. I had expected to find a coefficient estimate closer to 1, indicating that ethanol sales increase proportionally with the density of potential buyers. I suspect that this estimate is biased toward zero as a result of measurement error, which is exacerbated in panel data models (Hausman 2001). Using my monthly panel of flexible-fuel vehicle stocks, I regressed the logged number of flexible-fuel vehicles on a vector of county and month dummy variables. These controls explained over 99% of the variation in flexible-fuel vehicle stocks. Any residual variation that remains is likely dominated by measurement error, given that I construct my panel using a snapshot of vehicles on the road in 2007.

The next row of coefficients indicate that doubling the number of pumps per county leads to a 12%-14% reduction in sales volumes at individual stations. Because this measure of competition reflects choices about when and where to install new ethanol pumps, the coefficients do not quantify the effect on sales volumes of new pumps located at random, nor do they quantify the effect of new pumps locating in direct proximity to existing stations. Interpreted properly, the coefficient estimates suggest that new pumps locate in areas where competition is weak, drawing only a small number of customers away from existing stations. This result is not surprising, given the small number of ethanol pumps statewide.

Finally, the last set of coefficients indicate that ethanol sales volumes are particularly low in the first months after a pump first begins operating. Sales volumes are close to zero in the first month but quickly increase to normal levels by about the third or fourth month. This rapid increase in sales volume indicates that market participants are well-informed about ethanol's availability.

Coefficients on variables measuring the impact of flexible-fuel vehicle stocks, competition, and the length of time that pumps have been operating were virtually unchanged in models that added higher-order polynomial terms. Coefficients measuring the price elasticity of ethanol-equivalent fuel demand were also unchanged.

Cubic spline and semi-parametric results

Polynomial approximations in general are sensitive to the number of terms, to outliers, and to the local fit of the approximation. To test the performance of the polynomials, I also estimated equation (2.13) using a cubic spline approximation with knots at relative price intervals of 0.05, 0.10, and 0.20 prior to taking logarithms.[11] Cubic spline approximations are more flexible and less sensitive to outliers than polynomials, but are sensitive to the number and placement of knots. Finally, I estimated the model semi-parametrically using

[11] I controlled for station effects using fixed-effects estimation.

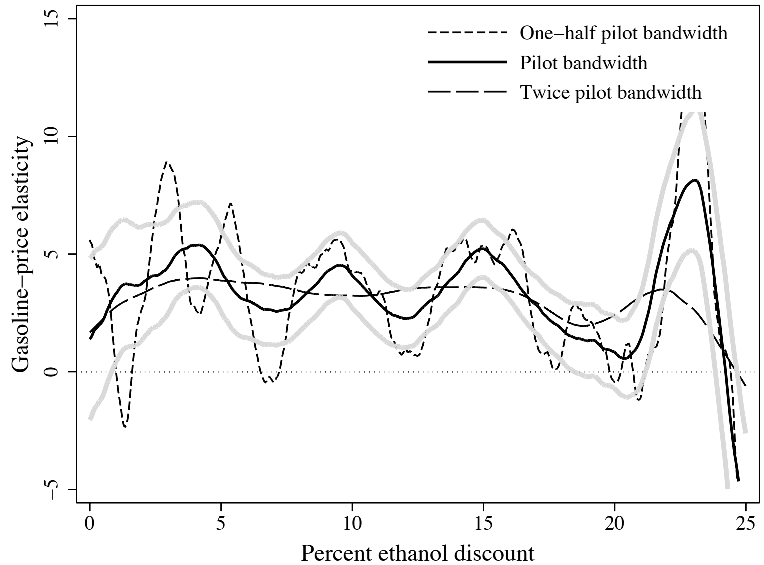

Figure 2.10: Elasticity based on semi-parametric estimates

Note: Figure shows gasoline-price elasticity of aggregate ethanol demand based on semi-parametric estimation using Yatchew's (1997) estimator and local polynomial estimation of non-parametric component. Figure shows estimates using bandwidths of approximately 0.015, 0.03, and 0.06. Solid gray lines are 95% confidence intervals are for estimates that use a bandwidth of 0.03 and are based on standard errors from local polynomial regressions. See text for details.

Yatchew's (1997) estimator for the partial linear model with a bandwidth of 0.03 and bandwidths half and twice as large.[12] Semi-parametric estimators give more detailed local approximations, but estimates are sensitive to the choice of bandwidth. See Hausman and Newey (1995) for a discussion of these tradeoffs in an application on gasoline demand.

Figures 2.9(c)–(d) plot the elasticity estimates from the cubic spline and semi-parametric approaches. For neither the cubic spline nor semi-parametric approaches am I able to reject the least flexible of the functional forms, although the more flexible semi-parametric

[12] For a general partial-linear model given by:

$$y_t = f(x_t) + Z_t\beta + \varepsilon_t,$$

Yatchew's procedure entails: (1) sorting the data by x_t, (2) differencing the data to remove the non-linear component $f(x_t)$ under the assumption that $f(x_t) \approx f(x_s)$ for $x_t \approx x_s$, (3) estimating β parametrically on the differenced data, (4) subtracting the predicted value from this parametric regression from the original dependent variable to yield $y_t - Z_t'\hat{\beta}$, and finally (5) regressing $y_t - Z_t'\hat{\beta}$ on x_t non-parametically using any number of non-parametric regression techniques. I employ tenth-order differencing using Yatchew's (1998) optimal differencing weights, which improves efficiency to within 5% of Robinson's (1988) fully efficient procedure. I control for station effects using station dummy variables. I estimate the non-parametric portion of the model using local polynomial regression, which has attractive properties in the extremes of the data. Polynomials also yield intuitive and convenient estimates for first derivatives. I use a quadratic local polynomial, which is appropriate for estimating first derivatives (Fan and Gijbels 1996), and an Epanechnikov kernel weighting function. I calculate the optimal "pilot" bandwidth using a rule-of-thumb approximation (Fan and Gijbels 1996, p.111).

Table 2.3: Semi-parametric estimation results

Variable	Coefficient
ln(price ethanol)	0.17
	(0.48)
ln(number flex-fuel vehicles)	0.02
	(0.01)
ln(number ethanol stations)	-0.10
	(0.04)
month 1 of operation	-0.98
	(0.05)
month 2 of operation	-0.23
	(0.05)
month 3 of operation	-0.12
	(0.04)
month 4 of operation	-0.04
	(0.04)
Observations	5017
R-squared	0.85

Note: Table is based on first-stage regression from Yatchew's (1997) semi-parametric estimator using tenth-order differencing to remove the nonparametric component. Dependent variable is logged monthly ethanol sales volume in gallons. Standard errors do not adjust for heteroskedasticity or serial correlation; robust standard errors require undoing the effects of differencing in the first stage. Regression includes station dummy variables, month dummy variables, and station-specific quadratic time trends. R-squared reflects all of these control variables. See text for details.

approach reveals additional nonlinearities that the polynomial and cubic spline approaches obscure. Figure 2.10 based on the semi-parametric estimates shows moderately sized peaks in the elasticity function where ethanol is discounted 5%, 10%, and 15% relative to gasoline. These peaks may reflect clustering of fuel-switching behavior around salient switch points. Sample mean elasticities based on the cubic spline and semi-parametric estimates are about 3.0 after excluding roughly 100 outliers in the extremes of the data.

For cubic spline approximation, coefficients measuring the impact of flexible-fuel vehicle stocks, competition, and the length of time that pumps have been operating were virtually identical to the fixed-effects polynomial coefficients, as was the elasticity of ethanol-equivalent fuel demand. The semi-parametric approach also yielded broadly similar coefficient estimates, as table 2.3 shows, although the impact of flexible-fuel vehicle stocks fell nearly to zero. The elasticity of ethanol-equivalent fuel demand still had the incorrect sign, but the coefficient fell by half.

Preference heterogeneity

Figure 2.9 demonstrates that the polynomial, cubic spline, and semi-parametric approaches all yield broadly similar elasticity functions, except in the extremes of the data where I have few observations. For each approach I am unable to reject estimates based on the least flexible model, be it the cubic polynomial, spline with the fewest knots, or semi-parametric approximation with the widest bandwidth. Elasticities are large and variable, though not nearly as large and variable as they would be if household preferences for ethanol were more homogeneous, as in figure 2.1(a). At a minimum, fuel-switching behavior extends over a wide range of relative prices where ethanol is discounted 0%-25% below gasoline. This suggests that preferences for ethanol are diffuse.

Unfortunately, I am not able to reveal the full distribution of preferences for ethanol as a gasoline substitute. This would clearly require an estimate of the elasticity function over all possible fuel-switching price ratios. This information is not available, given the limited range of prices that have been observed historically and the high variance of my estimates in the extremes of the data. That is, for example, I do not observe ethanol discounted 50% below gasoline, so I am unable to estimate the elasticity function or say anything about preferences in that neighborhood. I make this argument formally in appendix 2.10 below.

If I assume that fuel-switching price ratios are distributed normally, however, with a mean of 1.35, which is consistent with relative fuel economy in government tests, then I can calibrate the standard deviation to match what I estimate econometrically. Choosing a standard deviation of 0.43 yields an elasticity function whose sample mean is 3.0, which is consistent with my estimates. The resulting elasticity function is close to linear and has roughly the same overall height and slope as the polynomial and cubic spline estimates in figure 2.9.

These inferences about preferences are based on data for 1997–2006. I tested whether

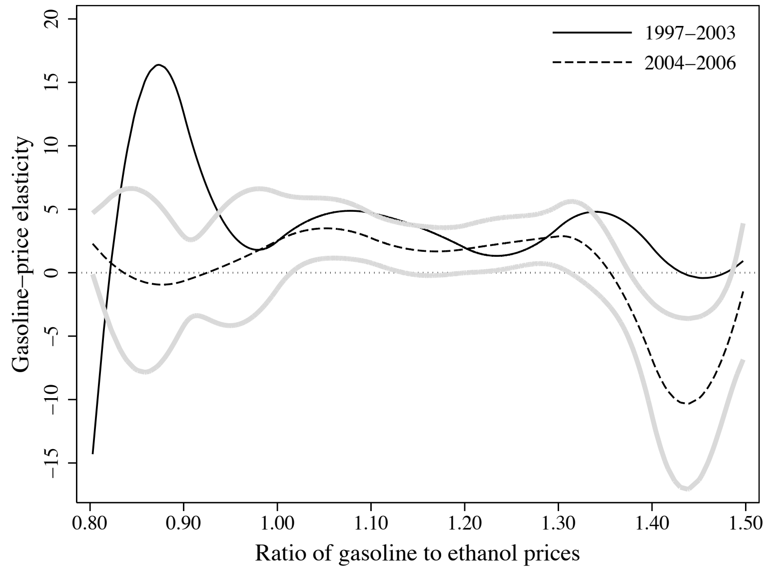

Figure 2.11: Elasticity function over time

Note: Figure shows the gasoline-price elasticity of aggregate ethanol demand estimated separately for 1997–2003 and 2004–2006. Estimates are based on cubic spline approximations with knots every 0.10. Solid gray lines are 95% confidence intervals for estimates using data from 2004-2006. See text for details.

preferences for ethanol have shifted over time by splitting my sample in half and estimating the model separately on data for 1997–2003 and 2004–2006. Preferences may have shifted as the stock of flexible-fuel vehicles has increased or as information about ethanol has improved. Similarly, the distribution may have narrowed as the number of stations has grown and heterogeneity associated with distance has diminished. Figure 2.11 presents elasticity estimates from these time periods based on a cubic spline approximation. Ignoring the extremes of the data where I have few observations, the elasticity function appears to have flattened and declined in magnitude. This is consistent with heterogeneity increasing over time, although the differences are not statistically significant. Estimates based on a polynomial approximation tell a similar story. I did not estimate the semi-parametric model separately for the two time periods, given the small sample sizes.

2.5 Policy simulation

I use my model and estimates to simulate the effects of an ethanol content standard, which mandates that denatured ethanol comprise a minimum fraction of the overall fuel supply. I simulate 10% and 25% standards. The 25% standard is consistent with the federal Renewable Fuel Standard of 36 billion gallons annually for 2022, which represents 25% of current gasoline consumption.[13] I assume in my simulations that the standard is met entirely through increasing the market share of E85 ethanol, although blending denatured ethanol with regular gasoline in other ratios would also be a potential compliance strategy. My model and estimates also could be used to evaluate other government policies that promote retail ethanol.

I simulate these standards assuming that fuel-switching price ratios are normally distributed with mean 1.35 and standard deviation 0.43. This results in a gasoline-price elasticity function with a similar shape as what I estimate econometrically. For comparison to previous analyses that assume identical preferences, I simulate the same standards assuming that fuel-switching price ratios are nearly identical with mean 1.35. I close the model by adding a supply side, drawing on previous work by Holland et al. (2007). I numerically search for retail fuel prices and a shadow value on the ethanol content constraint such that households maximize utility, suppliers maximize profits, the ethanol content standard is met, and markets clear. See the bottom of table 2.4 and appendix 2.11 for further details on the simulation.

Table 2.4 presents the simulation results. Consider first the results for scenario 2, which assumes that households have nearly identical preferences based on ethanol's fuel economy performance relative to gasoline. This constrains the equilibrium price ratio under

[13] Recall that while this standard mandates a minimum *quantity* of renewable fuel, the EPA rulemaking that implements the standard sets a minimum *percentage* of renewable fuel in each compliance period. EPA chooses the standard in advance of each compliance period based on projected fuel demand in an attempt to achieve the quantity specified by legislation. Implementation of the policy is therefore identical to a minimum market share requirement.

Table 2.4: Simulation results

Scenario 1: heterogeneous households		**Ethanol standard**	
	0%	**10%**	**25%**
equilibrium ethanol price ($/gallon)	3.52	3.49	2.88
equilibrium gasoline price ($/gallon)	2.47	2.47	2.90
equilibrium gasoline / ethanol price	0.70	0.71	1.00
quantity denatured ethanol (billion gallons)	12.62	12.92	30.49
quantity pure gasoline (billion gallons)	115.74	115.36	91.23
emissions (million mtCO_2)	1081.41	1079.58	954.90
change consumer surplus (billion $)	0	0.1	-33.34
change producer surplus (billion $)	0	-0.11	-0.43
change tax revenue (billion $)	0	-0.21	-12.41
total cost (billion $)	0	0.22	46.18
cost per gasoline saved ($/gallon)	0	0.59	1.88
cost per emissions reduced ($/mt$CO_2$)	0	121.34	365.04
Scenario 2: homogeneous households		**Ethanol standard**	
	0%	**10%**	**25%**
equilibrium ethanol price ($/gallon)	2.67	2.04	2.45
equilibrium gasoline price ($/gallon)	2.62	2.72	3.31
equilibrium gasoline / ethanol price	0.98	1.33	1.35
quantity denatured ethanol (billion gallons)	5.37	13.80	34.88
quantity pure gasoline (billion gallons)	132.12	125.11	104.44
emissions (million mtCO_2)	1189.40	1169.78	1093.05
change consumer surplus (billion $)	0	-13.01	-90.94
change producer surplus (billion $)	0	0.13	9.56
change tax revenue (billion $)	0	-3.64	-9.38
total cost (billion $)	0	16.51	90.76
cost per gasoline saved ($/gallon)	0	2.36	3.28
cost per emissions reduced ($/mt$CO_2$)	0	841.43	941.93

Note: Table shows simulation results for 10% and 25% ethanol content standards. Scenario 1 assumes fuel-switching price ratios distributed normally with mean 1.35 and standard deviation 0.43, while scenario 2 assumes mean 1.35 and standard deviation 0.01. All simulations assume: that every household owns a flexible-fuel vehicle, or equivalently, that flexible-fuel conversions are costless; a constant price elasticity for individual ethanol-equivalent fuel demand of -0.25; a competitive fuel supply industry producing retail ethanol and retail gasoline to maximize profits subject to the ethanol content standard; constant price elasticities of 1.25 and 2.5 for pure gasoline and denatured ethanol supply; 8.8 kilograms of CO_2 emissions per gallon of gasoline; and that replacing gasoline with ethanol reduces CO_2 emissions by 15% on an energy-adjusted basis. I calibrate the aggregate ethanol-equivalent fuel demand function to 2006 gasoline quantities and retail prices. I calibrate supply functions to 2006 denatured ethanol and gasoline quantities and national average wholesale spot prices net of the ethanol blending subsidy. I add a constant marginal cost for fuel distribution, marketing, taxes, and subsidies, which I calculate as the mean difference between retail and wholesale fuel prices. See appendix 2.11 for further details.

the standard to be near the fuel economy ratio of 1.35. A 25% ethanol content standard reduces gasoline consumption by about 22% and reduces carbon dioxide emissions by about 8%. The policy is costly, however, at $90 billion annually. I calculate total costs based on changes in consumer surplus, producer surplus, and tax revenue net of the federal ethanol subsidy.

Now consider the simulation results for scenario 1, where I have calibrated the standard deviation of preferences to match my econometric estimates. After calibrating the model to my econometric estimates, the surplus cost of a 25% ethanol content standard falls by half. Costs are lower in scenario 1 because households with particularly strong preferences for ethanol represent "low-hanging fruit" that are induced to purchase ethanol with less severe distortion of market prices.[14] About three-quarters of the $46 billion cost falls on consumers, while the rest falls on taxpayers. The fuel supply industry neither benefits nor suffers under the policy, although producer surplus in the table does not distinguish between ethanol and gasoline producers. The ethanol content standard reduces gasoline consumption by about 21% and carbon dioxide emissions by 12%.

The ethanol content standard remains a costly policy, however, even after accounting for household heterogeneity. Surplus costs for the 25% standard average about $1.90 per gallon of gasoline saved. For comparison, a recent study by Harrington et al. (2007) assumes $0.12 per gallon for the external costs of petroleum dependence, though the studies they review estimate a range of $0.08–$0.50 per gallon.[15] Surplus costs average about $370 per ton of carbon dioxide emissions avoided. Again, these costs far outweigh climate damages. A recent meta-analysis suggests that marginal damages are unlikely to

[14]In fact, baseline ethanol consumption is actually higher in scenario 1, and the 10% standard is just barely binding. The expansion of baseline ethanol consumption above current levels occurs because I assume for the simulation that all households own flexible-fuel vehicles, whereas in reality this fraction is quite small. The cost of flexible-fuel vehicle conversions is low but not zero, and so production of these vehicles derives primarily from CAFE incentives. Endogenizing flexible-fuel conversions by adding conversion costs to the analysis would reduce the impact of heterogeneity and increase the cost of complying with an ethanol content standard.

[15]They include petroleum dependence costs in a comprehensive measure of gasoline-related externalities, which they estimate at $2.20 per gallon. The majority of these costs depend on miles driven, however, and therefore hit ethanol even more strongly due to its poor mileage relative to gasoline.

exceed \$15 per ton of carbon dioxide emissions (Tol 2005), while even pessimistic recent estimates put marginal damages at only \$85 per ton (Stern 2006). These estimates are sensitive to assumptions about ethanol's life-cycle emissions. If land-use changes negate ethanol's moderate climate benefits, as recent studies suggest is likely, the ethanol content standard will actually increase greenhouse emissions.

There are several limitations to these results. First, my estimates reflect preferences of households that have self-selected to own flexible-fuel vehicles and participate in the nascent ethanol market. These households likely have the strongest preferences for ethanol and may have a different propensity to drive. Second, I assume that flexible-fuel conversions are costless. Third, much of the observed variation in household preferences for ethanol likely derives from differences in ethanol's convenience. Although I was not able to detect any significant changes between 1997-2003 and 2004-2006, this source of heterogeneity will likely diminish as the ethanol market further expands and ethanol becomes available in more locations. Addressing the first two issues would tend to increase the estimated cost of the standard, while addressing the third would reduce the impact of heterogeneity.

On the supply side, previous research does not estimate ethanol and gasoline supply elasticities particularly convincingly. Second, I do not consider the interaction of the ethanol content standard with preexisting distortions, such as agricultural price supports, nor do I consider other general equilibrium effects. Commodity prices are currently high, however, and price floors are not binding. Finally, I do not consider the potential for a breakthrough technology that facilitates cheap ethanol production from agricultural waste or other feedstocks, which may improve ethanol's cost-effectiveness in mitigating carbon dioxide emissions. While the Renewable Fuel Standard actually mandates that a substantial fraction of the standard be met with such fuels, forcing these technologies prematurely could increase the cost of the standard. Addressing these issues would have an ambiguous effect on the estimated cost of the standard.

2.6 Appendix: Aggregate demand and household welfare

I show above that the aggregate ethanol demand is

$$E(p_e, p_g) = \phi N H\left(\frac{p_g}{p_e}\right) d(p_e). \tag{2.14}$$

Aggregate demand for gasoline reflects households that own flexible-fuel vehicles but choose gasoline, as well as households that do not own flexible-fuel vehicles. Gasoline demand for households that own flexible-fuel vehicles is given by

$$\phi N \int_{p_g/p_e}^{\infty} \frac{d(p_g/r)}{r} dH(r). \tag{2.15}$$

By similar arguments aggregate demand for households that do not own flexible-fuel vehicles is

$$(1-\phi) N \int_{-\infty}^{\infty} \frac{d(p_g/r)}{r} dH(r), \tag{2.16}$$

which is just the total number of households that do not own flexible-fuel vehicles multiplied by their average gasoline consumption. Here I rely on the assumption that flexible-fuel vehicles are allocated at random, which implies that fuel-switching price ratios are distributed identically for households that do and do not own flexible-fuel vehicles. Adding these two expressions gives aggregate gasoline demand:

$$G(p_e, p_g) = \phi N \int_{p_g/p_e}^{\infty} \frac{d(p_g/r)}{r} dH(r) + (1-\phi) N \int_{-\infty}^{\infty} \frac{d(p_g/r)}{r} dH(r). \tag{2.17}$$

Maximized utility for an individual household that chooses ethanol is

$$v(d(p_e)) + y - p_e d(p_e), \tag{2.18}$$

which holds whenever $r \leq p_g/p_e$, while a household that chooses gasoline has utility given by

$$v(d(p_g/r)) + y - p_g \frac{d(p_e/r)}{r}, \tag{2.19}$$

which holds whenever $r > p_g/p_e$. Because I assume that household utility is linearly separable in the composite good, and because I assume an interior solution with respect to consumption of this good, each household's utility function has dollar units. This allows me to compute average welfare:

$$\begin{aligned} &\phi \left\{ \int_{-\infty}^{p_g/p_e} \left[v(d(p_e)) + y - p_e d(p_e) \right] dH(r) + \int_{p_g/p_e}^{\infty} \left[v(d(p_g/r)) + y - p_g \frac{d(p_e/r)}{r} \right] dH(r) \right\} \\ &+ (1-\phi) \int_{-\infty}^{\infty} \left[v(d(p_g/r)) + y - p_g \frac{d(p_e/r)}{r} \right] dH(r) \end{aligned} \tag{2.20}$$

where the top term is average welfare for households that own flexible-fuel vehicles weighted by the fraction of these households, and the bottom term is average welfare for households that do not own flexible-fuel vehicles weighted by the fraction of such households. Average welfare for households that own flexible-fuel vehicles reflects both households that choose ethanol as well as households that choose gasoline. Multiplying by the total number of households N gives aggregate welfare.

2.7 Appendix: Retail supply behavior

How will a retailer facing the demand functions developed above respond to shifting costs? For an ethanol retailer located close to other retailers, competition will drive the retail price of ethanol down to marginal costs:

$$p_e = c_e, \tag{2.21}$$

where c_e is the marginal cost of ethanol. The equilibrium ratio of retail gasoline to ethanol prices is given by:

$$\rho^* = \frac{p_g}{c_e}, \tag{2.22}$$

where $\rho = p_g/p_e$ is the price ratio the retailer chooses, and where I assume for simplicity that the retail price of gasoline p_g is fixed exogenously by conditions in the retail gasoline market.[16] Changes in ethanol's cost relative to gasoline therefore transmit directly to relative retail prices:

$$\frac{\partial \rho^*}{\partial (p_g/c_e)} = 1. \tag{2.23}$$

When ethanol's cost relative to the price of gasoline increases, relative retail prices increase accordingly.

Marginal-cost pricing is not a particularly good model for understanding retail ethanol pricing behavior. Current retail ethanol markets reflect a peculiar mix of monopoly power and competition. Because relatively few stations offer retail ethanol, customer bases overlap only marginally, if at all, allowing ethanol retailers to operate largely as local monopolists. At the same time, these retailers compete directly with nearby gasoline stations in the broader fuels market, because flexible-fuel vehicle owners are able to switch seamlessly between ethanol and gasoline. The monopolist ethanol retailer that only offers ethanol chooses the price of ethanol to maximize profits:

$$\Pi(p_e; p_g) = E(p_e; p_g)p_e - c_e E(p_e; p_g), \tag{2.24}$$

where Π is retailer profit, which is a function of the retail prices of ethanol p_e and regular gasoline p_g, E is the quantity of ethanol demanded as a function of retail prices, and c_e is

[16]This assumption is consistent with the current fuel market, where relatively few stations offer ethanol and ethanol sales volumes are low relative to gasoline. This assumption would not be valid for a significantly expanded retail ethanol market.

the constant marginal cost of offering ethanol. I assume for simplicity that the retail price of gasoline p_g is fixed exogenously by conditions in the retail gasoline market.

The first-order condition of this maximization problem is given by:

$$E + E' p_e - c_e E' \equiv 0, \tag{2.25}$$

where all derivatives are with respect to the retail price of ethanol and I have suppressed the arguments of functions for clarity. Collecting terms that contain E', moving E to the right-hand side, and then dividing by p_e and E' on both sides yields:

$$\frac{p_e - c_e}{p_e} \equiv -\frac{E}{p_e} \cdot \frac{1}{E'}. \tag{2.26}$$

This is equivalent to

$$\mu_e \equiv -\frac{1}{\xi_e}, \tag{2.27}$$

where $\mu_e \equiv (p_e - c_e)/p_e$ is the percent retail markup of ethanol and ξ_e is the own-price elasticity of aggregate ethanol demand. This is the standard monopoly result where the retailer equates the percent retail markup to the negative reciprocal of the price elasticity of demand.[17]

Restating the first-order condition in terms of the price ratio ρ by making the substitutions $p_e = p_g/\rho$ and $\xi_e = -\xi_g + \xi_f$ yields:

$$1 - \frac{\rho}{p_g/c_e} \equiv -\frac{1}{-\xi_g + \xi_f}, \tag{2.28}$$

[17] Pricing behavior is more complicated when the ethanol retailer also offers gasoline. Adding profits from gasoline sales to the maximization problem results in a modified first-order condition:

$$\mu_e + \left(\frac{G'}{E'} \cdot \frac{p_g}{p_e}\right) \mu_g \equiv -\frac{1}{\xi_e},$$

where G' is the change in gasoline sales volume given a marginal increase in the price of ethanol, $\mu_g \equiv (p_g - c_g)/p_g$ is the percent retail markup of gasoline, and all other terms are as above. I again assume that retail gasoline prices are fixed by market competition. When a station's ethanol price has no effect on its gasoline sales, so that $G' = 0$, the first-order condition reduces to the simple case above. When $G' > 0$, however, the optimal price of ethanol is higher, because increasing the price of ethanol drives some consumers to gasoline at the same station. This incentive increases with G'. The incentive to increase ethanol prices and drive consumers to gasoline also increases with gasoline markups μ_g. In practice, retailers are unlikely to retain many customers that switch to gasoline when the price of ethanol increases, given the large number of competing stations that also offer gasoline. So G' will likely be quite small in the current market, and pricing behavior will tend toward the simple case above.

where ρ is the price ratio the retailer chooses. Assuming that the price elasticity of individual ethanol-equivalent fuel demand ξ_f is constant, the implicit function theorem gives the following comparative static for the impact of a change in ethanol's relative cost on the profit-maximizing price ratio:

$$\frac{\partial \rho^*}{\partial (p_g/c_e)} = \frac{\rho^*}{\frac{p_g}{c_e} - \left(\frac{p_g/c_e}{-\xi_g+\xi_f}\right)^2 \xi_g'} > 0, \tag{2.29}$$

where ρ^* is the profit-maximizing price ratio and the inequality assumes that $\xi_g' < 0$ at the optimum. Recall that ξ_f is constant by assumption and that ξ_g and ξ_g' only depend on relative prices.

Expression (2.29) implies that changes in relative costs will have their largest impact on relative retail prices when the gasoline-price elasticity is roughly constant near the optimum, so that ξ_g' is close to zero. In contrast, when the elasticity is highly variable near the optimum, which indicates a large concentration of households near that same fuel-switching price ratio, ξ_g' will be large in magnitude and relative prices will be less responsive to changes in ethanol's costs. In the extreme case where households have identical preferences for ethanol, ξ_g' will be infinitely large in magnitude and relative prices will be invariant to underlying costs. Retailers will be reluctant to raise ethanol prices when costs increase, lest they drive all consumers to gasoline. At the same time, retailers will have no incentive to reduce prices when costs fall, because lowering prices will not stimulate any additional demand.

Figure 2.12 illustrates this first-order condition and comparative static for two hypothetical gasoline-price elasticity functions, where I have set the elasticity of individual ethanol-equivalent fuel demand to a constant -0.25. The figures illustrate that when household preferences are nearly homogeneous, so that price elasticities are highly variable, as in figure 2.12(a), the profit-maximizing price ratio is insensitive to changes in relative

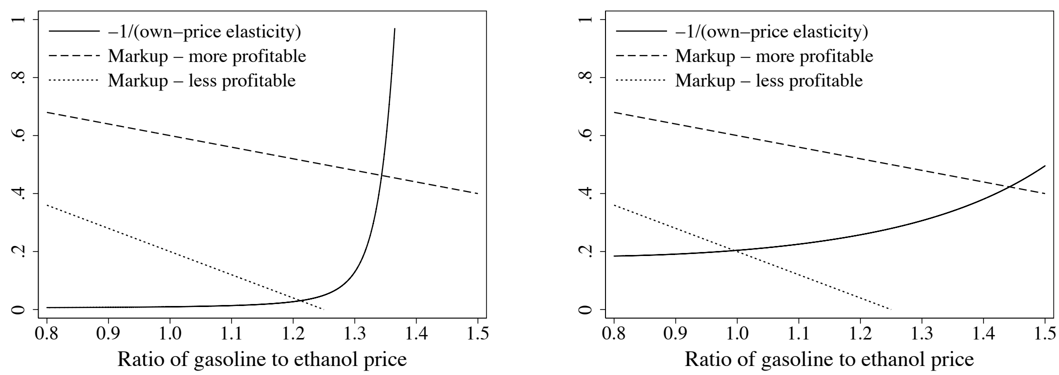

(a) $r \sim N(1.25, 0.05)$

(b) $r \sim N(1.25, 0.30)$

Figure 2.12: Profit-maximizing retail price ratio

Note: Figure illustrates profit-maximizing price ratios for a monopolist ethanol retailer. Profit-maximizing price ratios are given by interesection of markups and the negative reciprocal of the own-price elasticity, as in equation (2.27). More profitable and less profitable cases assume that marginal ethanol costs are 60% and 80% the retail price of gasoline. Elasticity functions assume a constant ethanol-equivalent fuel price elasticity of -0.25.

costs. When household preferences are more diffuse, however, so that price elasticities are less variable, as in figure 2.12(b), shifts in relative costs lead to large changes in the profit-maximizing price ratio.

2.8 Appendix: Aggregate price trends

Figure 2.13 plots average retail ethanol and regular gasoline prices from October 1997 through November 2006. Average ethanol prices track regular gasoline prices closely, albeit at a noticeable discount for most of the period.

Figure 2.14 plots average wholesale prices for the same time period. Wholesale spot prices for denatured ethanol do not always track wholesale gasoline prices closely. This is perhaps not surprising, given that demand for denatured ethanol derives largely from its role as a complement to gasoline production and less from its role as a gasoline substitute. Opportunities for direct substitution do exist, however, and large price differences can create strong incentives for substitution, which is one reason that wholesale ethanol prices track wholesale gasoline prices broadly over time. This is particularly evident in the fall

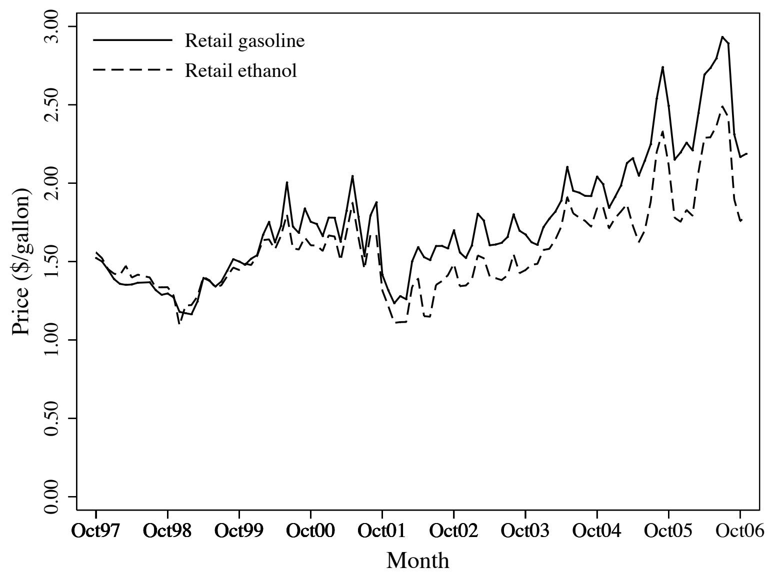

Figure 2.13: Retail fuel prices

Note: Retail ethanol price is the monthly volume-weighted average retail price of ethanol at reporting stations in Minnesota. Retail gasoline price is the monthly county average retail price of regular gasoline, weighted by retail ethanol sales volumes at these same stations. Prices are in 2006 dollars.

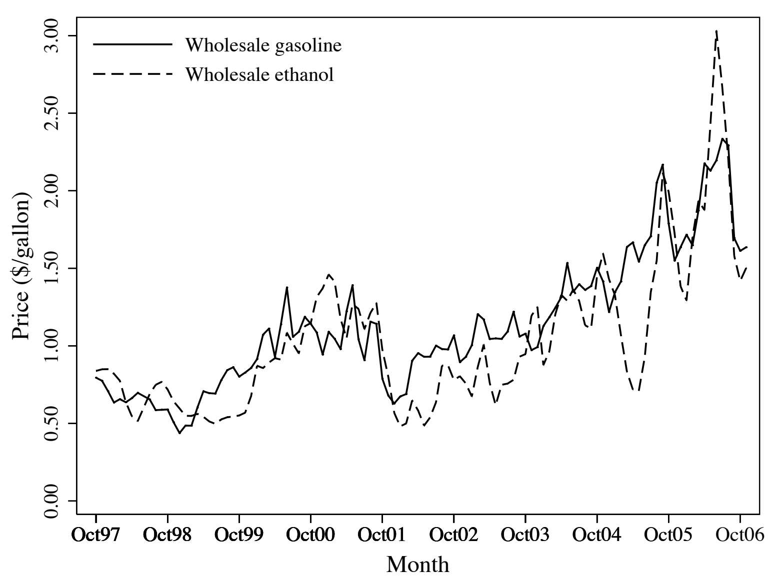

Figure 2.14: Wholesale fuel prices

Note: Wholesale ethanol price is a weighted average of the spot price for denatured ethanol in Minneapolis and Fargo, less the federal ethanol blending tax credit. Wholesale gasoline price is the Minnesota volume-weighted average rack price. Prices are in 2006 dollars.

of 2005, when ethanol helped ease gasoline shortfalls after Hurricanes Katrina and Rita knocked out Gulf Coast petroleum refineries and distribution pipelines. Ethanol prices were low relative to gasoline in the first half of 2005 due to a glut of ethanol. Ethanol prices then spiked to equal gasoline prices as ethanol substituted for gasoline after the hurricanes. Ethanol's margin relative to gasoline eventually returned to pre-hurricane levels as refineries and pipelines came back on line and as imports of refined gasoline arrived from abroad.

A second reason that wholesale prices track broadly is that ethanol and a petroleum-based chemical fuel additive called methyl tertiary-butyl ether (MTBE) are close substitutes in many U.S. regions during much of this time period, creating an avenue for petroleum prices to correlate indirectly with ethanol prices. The importance of this substitution is most evident in the first half of 2006, when fuel suppliers quit using MTBE due to concerns about potential groundwater contamination. Prices surged as ethanol filled the gap left by this key substitute. Ethanol prices fell in the summer months as ethanol refiners scaled up production and as fuel distributors resolved logistical difficulties in transporting ethanol from refineries in the midwest, where ethanol is produced, to markets on the coasts, where MTBE had previously held a large market share.

Figure 2.6 above demonstrates that large fluctuations in relative wholesale prices correlate with comparatively small changes in retail prices. Note that the scale for the wholesale price ratio in figure 2.6 above is five times as large as the scale for the retail price ratio. What explains this behavior? The natural assumption is that ethanol retailers are pricing ethanol based primarily on what flexible-fuel vehicle owners are willing to pay, relative to gasoline, as opposed to what the fuel costs. As I show above in appendix 2.7, when the elasticity is highly variable and retailers are monopolists, the relative price of ethanol will be insensitive to changes in ethanol's costs relative to gasoline. The pricing behavior in figure 2.6 is therefore consistent with a highly variable elasticity function.

There are alternative explanations. First, denatured ethanol costs may differ substantially from publicly reported spot prices. The industry representatives I spoke with indicated that most retailers purchase ethanol in long-term contracts ranging from six months to one year. Contracts often tie costs directly to the price of gasoline, which is consistent with the long-term relationship between ethanol and gasoline prices. Few retailers buy denatured ethanol on the spot market, because spot prices are typically higher than long-term contract prices. Although retailers sometimes purchase ethanol on the spot market to cover minor shortfalls, the quantities involved are generally small. In short, publicly reported spot prices overstate the variability of wholesale ethanol costs relative to gasoline.

Second, retailers may have an incentive to absorb temporary fluctuations in relative costs. Given the relatively small number of ethanol stations, ethanol consumers may drive longer distances or coordinate their daily and weekly activities around refueling with ethanol. If the relative price of ethanol is highly variable, so that households are unsure whether ethanol's discount will be sufficiently generous, they may be less willing to incur these search costs. Station owners therefore have an incentive to allay this uncertainty by maintaining retail ethanol prices that are more consistent with the long-run relationship between gasoline prices and ethanol costs. Short-term profits may suffer, but this strategy helps maintain a consistent customer base. Indeed, several industry representatives I spoke with indicated that some retailers were actually pricing ethanol below costs in late 2005 and early 2006. Ethanol costs were high relative to gasoline, due to the discontinuation of MTBE, but some retailers were willing to incur temporary losses to maintain favorable relationships with their customers. I am also told that the earliest retailers had particularly low sales volumes until they learned to price ethanol at a consistent discount to gasoline. Sales volumes then increased markedly.

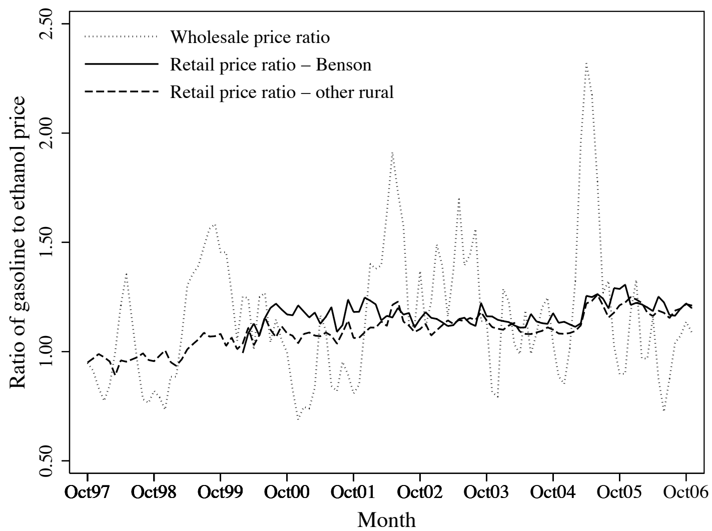

(a) Benson area vs. other rural

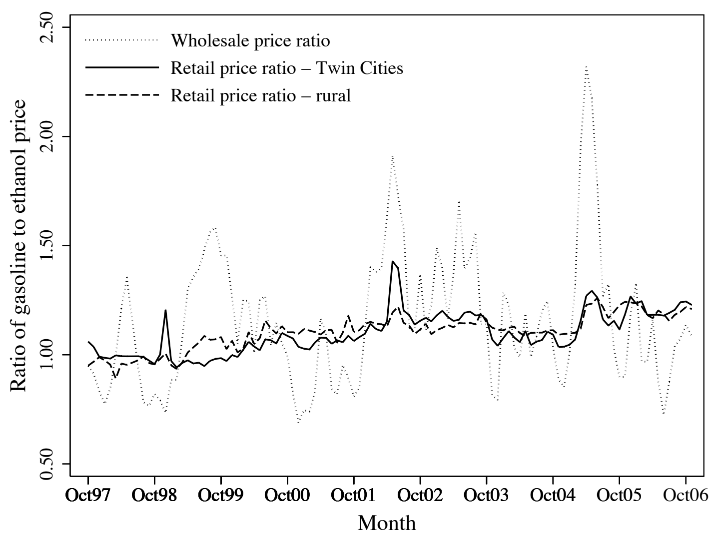

(b) Twin Cities vs. rural

Figure 2.15: Relative wholesale prices and relative retail fuel prices

Note: Top figure shows relative retail prices for stations in counties within 50 miles of Benson and for other rural counties. Bottom figure shows relative retail prices for stations in Twin Cities counties and for stations in rural counties.

2.9 Appendix: Evidence of cross-sectional variation in pricing behavior

About one-third of ethanol retailers in Minnesota purchase ethanol directly from an ethanol refinery in Benson, which is located in the southwestern part of the state. Throughout the entire sample period, this refinery supplied retail ethanol at a fixed nominal discount to the spot price of regular gasoline. The ethanol retailers, in turn, agreed to price retail ethanol at the same discount below regular gasoline at their stations.[18] When retail ethanol prices are tied directly to the price of gasoline, relative prices will be less responsive to changes in ethanol's relative cost. This is apparent in figure 2.15(a), which plots relative retail prices for stations located in counties within 50 miles of the Benson refinery, which are most likely to have contracts with this refinery, and for stations located in other counties outside the Twin Cities. In 2000-2001, when wholesale ethanol costs were high relative to gasoline, stations near Benson priced ethanol at a larger percent discount. This happened again in late 2003-2004 and at times in late 2005-2006.

Figure 2.15(b) plots relative retail prices for stations located inside and outside the Twin Cities, where the density of retail ethanol stations is higher. Stations in the Twin Cities appear to be more sensitive to changes in relative costs. When wholesale ethanol costs are low relative to gasoline, stations in the Twin Cities discount ethanol more heavily than in rural areas. When wholesale ethanol costs are high relative to gasoline, retailers in the Twin Cities do not discount ethanol as generously. This pricing behavior is consistent with retailers in the Twin Cities facing greater competition and therefore being more sensitive to changes in relative costs.

2.10 Appendix: Using elasticity estimates to reveal preferences

It is possible in theory to retrieve the distribution of household preferences from aggregate price repsonses. Recall that equation (2.9) above links the distribution of household

[18] The discount was fixed at $0.20 per gallon for several years, then increased to $0.40 per gallon for several years, and finally fluctuated between $0.35 and $0.70 per gallon for the last several years. This unique pricing agreement ended in the fall of 2007. The ethanol refinery now supplies retail ethanol at market prices, and retailers are free to price ethanol at whatever price the market will bear. I am not aware of any similar agreements between ethanol retailers and their suppliers.

preferences to aggregate price responses:

$$\xi_g(x) = \frac{h(x)}{H(x)} x. \tag{2.30}$$

Dividing both sides by x and using the first-derivative rule for logarithms gives:

$$\frac{\xi_g(x)}{x} = \frac{\partial \ln H(x)}{\partial x}. \tag{2.31}$$

Assume that fuel-switching price ratios are known to range from r_L to r_H. Then integrating both sides with respect to x through $r > r_L$ gives

$$\begin{aligned} \int_{r_L}^{r} \frac{\xi_g(x)}{x} dx &= \int_{r_L}^{r} \frac{\partial \ln H(x)}{\partial x} dx \\ &= \ln H(r) + C, \end{aligned} \tag{2.32}$$

where C is a constant of integration. Finally, taking the exponential of both sides yields

$$\exp\left(\int_{r_L}^{r} \frac{\xi(x)}{x} dx \right) = \exp(C) \cdot H(r). \tag{2.33}$$

Given C and an econometric estimate of $\xi_g(x)$ over the interval $[r_L, r]$, equation (2.33) yields an estimate of the distribution of household preferences.

A boundary condition is required to solve for C. The lower boundary will not work. This is clear from equation (2.32), where the right side is undefined at the lower bound where $r = r_L$ because $H(r_L) = 0$. The lower boundary condition does not work because the elasticity function in equation (2.9) is undefined at r_L. The upper boundary will work, however, provided that an estimate for the gasoline-price elasticity function covering the entire interval $[r_L, r_H]$ is available. At the upper boundary $H(r_H) = 1$ so C is simply the area under the function $\xi_g(x)/x$ on the interval $[r_L, r_H]$. Any other pair of $r^* \in (r_L, r_H)$ and $H(r^*)$ will also work as a boundary condition if $H(r^*)$ is somehow known.

Unfortunately, I am not able to reveal the full distribution of household preferences for ethanol based on my estimates, because I do not have an estimate of the elasticity function over the entire range of possible fuel-switching price ratios.

2.11 Appendix: Simulation details

2.11.1 Minimum ethanol content standard

An ethanol content standard mandates that denatured ethanol comprise a minimum fraction of the overall fuel supply:

$$\frac{\pi_e E + \pi_g G}{E + G} \geq \sigma, \tag{2.34}$$

where E and G are the aggregate retail quantities of ethanol and gasoline, π_e is the percent denatured ethanol content of retail ethanol, π_g is the percent denatured ethanol content of retail gasoline, and σ is the minimum fraction of denatured ethanol in the fuel supply as mandated by the ethanol content standard. I assume that $\pi_e \geq \sigma \geq \pi_g$, where the leftmost inequality guarantees that the ethanol content standard is technically achievable, and the rightmost inequality implies that the standard is not met trivially for any combination of fuels. Rearranging the inequality demonstrates that the standard is equivalent to a minimum market share for retail ethanol:

$$\frac{E}{G} \geq -\frac{\pi_g - \sigma}{\pi_e - \sigma} \tag{2.35}$$

An ethanol content standard is therefore identical to any fuel performance standard that implicitly mandates a minimum market share requirement for ethanol, including a low-carbon fuel standard met through increased ethanol production.

2.11.2 Model of the fuels market

Following Holland et al. (2007) I assume that a competitive fuel supply industry maximizes profits given by:

$$p_e E + p_g G - C(E,G) + \lambda[\pi_e E + \pi_g G - \sigma(E+G)], \tag{2.36}$$

where p_e and p_g are the retail prices of ethanol and regular gasoline, E and G are the aggregate retail quantities of ethanol and regular gasoline, $C(\cdot,\cdot)$ is the fuel industry's cost function, which is increasing in both arguments and convex, λ is the shadow value of the ethanol content constraint, and π_e and π_g are as above. Note that the total quantity of denatured ethanol required to produce the given retail quantities is $\pi_e E + \pi_g G$, while the total quantity of pure gasoline is $(1-\pi_e)E + (1-\pi_g)G$. The cost function reflects denatured ethanol and gasoline refining and distribution costs, as well as the costs of blending, distribution, and retail marketing. The cost function also reflects retail fuel taxes, as well as subsidies for denatured ethanol blending.

The first-order conditions from the fuel supplier profit maximization problem and the household utility maximization problem above together characterize market equilibrium:

$$v'(e) = \frac{\partial C(E,G)}{\partial E} + \lambda[\sigma - \pi_e], \tag{2.37}$$

$$v'(rg)r = \frac{\partial C(E,G)}{\partial G} + \lambda[\sigma - \pi_g], \tag{2.38}$$

and

$$\lambda[\pi_e E + \pi_g G - \sigma(E+G)] = 0, \tag{2.39}$$

where $\lambda \geq 0$. The first condition holds for all consumers with $r \leq p_g/p_e$ who choose ethanol and the second condition holds for all consumers with $r > p_g/p_e$ who choose gasoline. These equilibrium conditions state that each household's marginal willingness to pay for fuel equals the fuel supply industry's marginal cost. The third condition is that either the ethanol content constraint is binding or that the shadow value of the constraint is zero.

The ethanol content standard gives an implicit subsidy of $\lambda[\pi_e - \sigma]$ for the production of retail ethanol, because the denatured ethanol content of retail ethanol exceeds the standard. Conversely, the standard imposes an implicit tax of $\lambda[\sigma - \pi_g]$ on the production of retail gasoline, because the denatured ethanol content of gasoline is less than the standard. The ultimate effect of the standard on equilibrium fuel quantities depends on the stringency of the standard, the fuel industry's cost function, the household's ethanol-equivalent fuel demand function, and the distribution of fuel-switching price ratios.

Holland et al. (2007) use a similar model to evaluate a low-carbon fuel standard met through increased ethanol production. They show that such a standard can never deliver efficient reductions in carbon dioxide emissions, because the standard implicitly subsidizes ethanol while taxing gasoline. Any first-best policy must tax all fuels that contain carbon, including ethanol, based on marginal external damages. They also show that a low-carbon fuel standard might actually increase energy consumption and carbon dioxide emissions, because the fuel supply industry could meet the standard simply by increasing ethanol production. This is similar to the well-known result that a pollution performance standard may create incentives to expand output if the rate of pollution increases less than proportionally with production. These results apply equally to my analysis of an ethanol content standard.

2.11.3 Demand calibration

I assume that the fuel consumption component of individual utility is of the form:

$$v(e+rg) = k^{1/\varepsilon}\frac{\varepsilon}{\varepsilon-1}(e+rg)^{\frac{\varepsilon-1}{\varepsilon}}, \tag{2.40}$$

so that the first-order conditions in (2.3) and (2.4) above yield the following expression for individual ethanol-equivalent fuel demand:

$$d(p) = kp^{-\varepsilon}, \tag{2.41}$$

where k is a constant, p is the ethanol-equivalent price, and $-\varepsilon$ is the constant price elasticity of ethanol-equivalent fuel demand. The assumption that individual demand has a constant price elasticity is consistent with my econometric model, which also generates a constant price elasticity of individual ethanol-equivalent fuel demand. Maximized individual utility is given by

$$\frac{k}{\varepsilon - 1} p^{1-\varepsilon} + y. \tag{2.42}$$

From here, it is straightforward to calculate aggregate quantities of retail ethanol and gasoline demand, as well as aggregate household welfare, based on the general expressions available in appendix 2.6. Given the functional form assumption above, these expressions depend on the price elasticity of individual ethanol-equivalent fuel demand $-\varepsilon$, the scale of fuel demand Nk, the fraction of households that own flexible-fuel vehicles ϕ, and the distribution of fuel-switching price ratios $H(r)$.

I calibrate $-\varepsilon = -0.25$ based previous estimates of this parameter in the literature. I then calibrate Nk based on aggregate gasoline demand and average retail gasoline prices in 2006 under the assumption that $\phi = 0$. This is consistent with current market conditions where few households own flexible-fuel vehicles. I then reset $\phi = 1$ for the simulations. Simulations therefore reflect market conditions in a hypothetical world where the scale of ethanol-equivalent fuel demand is equivalent to current levels but where all households own flexible-fuel vehicles. Finally, I calibrate $H(r)$ by assuming that fuel-switching price ratios are normally distributed with mean 1.35 and standard deviation 0.43, which results in a gasoline-price elasticity function that has roughly the same shape as what I estimate econometrically.

2.11.4 Supply calibration

I assume that marginal costs in the fuel supply industry are given by

$$\frac{\partial C(E,G)}{\partial E} = \pi_e K_e {B_e}^{\eta_e} + (1-\pi_e) K_g {B_g}^{\eta_g} + \psi_e + \tau_e - \pi_e \theta \tag{2.43}$$

and

$$\text{(2.44)} \qquad \frac{\partial C(E,G)}{\partial G} = \pi_g K_e B_e^{\eta_e} + (1-\pi_g) K_g B_g^{\eta_g} + \psi_g + \tau_g - \pi_g \theta,$$

where: π_e and π_g are the denatured ethanol content ratios of retail ethanol and gasoline; $B_e \equiv \pi_e E + \pi_g G$ and $B_g \equiv (1-\pi_e)E + (1-\pi_g)G$ are the quantities of pure ethanol and gasoline required to produce the retail quantities E and G; the functions $K_e B_e^{\eta_e}$ and $K_g B_g^{\eta_g}$ are marginal costs of denatured ethanol and gasoline production, reflecting all costs through delivery to fuel terminals, with η_e, η_g, K_e, and K_g parameters to be calibrated; ψ_e and ψ_g are the constant marginal costs of distributing fuels to retail outlets and retail marketing, to be calibrated; τ_e and τ_g are retail fuel taxes remitted by fuel retailers to state and federal governments; and θ is the federal blending subsidy for denatured ethanol.

I assume that $\pi_e = 0.85$, because E85 ethanol contains 85% denatured ethanol, and calibrate $\pi_g = 0.04$, which was the fraction of denatured ethanol in gasoline in 2006. I assume 8.8 kilograms of CO_2 emissions per gallon of gasoline and that replacing gasoline with ethanol reduces CO_2 emissions by 15% on an energy-adjusted basis. I assume that the constant price elasticity of denatured ethanol supply is $1/\eta_e = 2.5$ and that the price elasticity of gasoline supply is $1/\eta_g = 1.25$, which are the midpoints of the ranges considered by Holland et al. (2007) based on their reading of the previous literature. I then calibrate K_e and K_g based on 2006 production quantities and wholesale spot prices for denatured ethanol and gasoline. I calibrate distribution and marketing costs $\psi_e = \psi_g = \$0.16$ as the current differential between average wholesale prices for retail gasoline and average pre-tax retail prices. I calibrate $\tau_e = \tau_g = \$0.50$ as the average differential between pre-tax and tax-inclusive retail prices. Finally, I calibrate $\theta = \$0.51$, which is the current federal subsidy for denatured ethanol blending.

2.11.5 Numerical solution algorithm

The numerical solution algorithm is as follows:

(1) Choose an initial fuel price vector $p^0 = [p_e^0, p_g^0]$.

(2) Set initial shadow value of ethanol content constraint to zero: $\lambda = 0$.

(3) Compute quantities supplied based on initial price vector and first-order conditions from industry profit maximization problem. If fuel supply industry is not in compliance, increase λ, return to step (2), and iterate until industry is exactly in compliance with the ethanol content standard, yielding retail quantities supplied $S^0 = [S_e^0, S_g^0]$

(4) Compute retail quantities demanded based on initial price vector and first-order conditions from household maximization problem, yielding retail quantities demanded $D^0 = [D_e^0, D_g^0]$.

(5) If the markets clear, that is if

$$D^0 - S^0 = [D_e^0, D_g^0] - [S_e^0, S_g^0] = [0, 0],$$

then stop. Otherwise, update the price vector according to $p^1 = p^0 + \kappa[D^0 - S^0]$, where κ is a positive constant. This moves the price vector in a direction that reduces excess demand. In practice I decrease κ as the number of iterations increases in order to hone in on the market-clearing price vector. Return to step (1), and iterate.

2.12 Appendix: Additional estimation results

Table 2.5: Estimation results—robustness of fixed-effects estimator

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ln(price ethanol)	0.48	0.15	0.01	1.33	2.22	0.39	0.35	0.53
	(0.64)	(0.65)	(0.80)	(0.11)	(0.15)	(0.64)	(0.64)	(0.68)
ln(gas price / ethanol price)^1	2.60	2.52	2.69	3.84	5.63	2.55	3.04	3.16
	(0.64)	(0.67)	(0.85)	(0.34)	(0.40)	(0.66)	(0.67)	(1.44)
ln(gas price / ethanol price)^2							-2.63	
							(0.86)	
ln(gas price / ethanol price)^3							-6.38	
							(3.73)	
gas price minus ethanol price)								-0.31
								(0.78)
ln(number flex-fuel vehicles)	0.07	0.04	0.07	0.05	0.18		0.08	0.07
	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)		(0.03)	(0.03)
ln(number ethanol stations)	-0.14	-0.35	-0.14	-0.06	0.12		-0.13	-0.14
	(0.11)	(0.10)	(0.12)	(0.08)	(0.12)		(0.11)	(0.11)
month 1 of operation	-1.20	-1.06	-0.92	-1.14	-1.11		-1.18	-1.20
	(0.19)	(0.14)	(0.13)	(0.17)	(0.13)		(0.20)	(0.19)
month 2 of operation	-0.28	-0.26	-0.18	-0.32	-0.44		-0.27	-0.28
	(0.15)	(0.09)	(0.08)	(0.12)	(0.10)		(0.15)	(0.15)
month 3 of operation	-0.14	-0.16	-0.10	-0.16	-0.32		-0.13	-0.13
	(0.11)	(0.09)	(0.08)	(0.10)	(0.09)		(0.11)	(0.11)
month 4 of operation	-0.04	-0.08	-0.04	-0.07	-0.26		-0.04	-0.04
	(0.07)	(0.06)	(0.06)	(0.06)	(0.07)		(0.07)	(0.07)
station-specific quadratic trend	x			x		x	x	x
county-specific quadratic trend		x						
month dummies	x	x	x			x	x	x
number observations	5027	5027	5027	5027	5027	5027	5027	5027
number stations	237	237	237	237	237	237	237	237
R-squared	0.84	0.76	0.68	0.80	0.57	na	0.84	0.84
first-order correlation of residuals	0.39	0.54	0.62	0.44	0.66	0.38	0.39	0.39
	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)

Note: Dependent variable is logged monthly ethanol sales volume in gallons. Logged price variables have been normalized to equal zero at sample mean prices. Standard errors in parentheses are clustered by station. All regressions except (6) use deviations from within-station means to control for station effects. Regression (1) is identical to the linear fixed-effects regression above in table 2.2 and is included for comparison. Regression (2) replaces the station-specific quadratic trend with a county-specific quadratic trend. Regression (3) drops all quadratic trends. Regression (4) drops month dummy variables. Regression (5) drops all time controls. Regression (6) drops station and county characteristics. I estimated this regression using dummy variables to control for station effects as the within-station estimator had numerical difficulties; I therefore omit the within R-squared for this regression. Regression (7) is identical to the cubic fixed-effects regression above in table 2.2. Regression (8) adds the retail gasoline minus ethanol price. First-order correlation of residuals is the coefficient from the least-squares regression of residuals on their lagged values. See text for details.

Table 2.6: Estimation results—robustness of first-difference estimator

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ln(price ethanol)	0.77	0.77	0.67	1.06	1.08	0.63	0.64	0.86
	(0.48)	(0.46)	(0.45)	(0.11)	(0.09)	(0.46)	(0.47)	(0.49)
ln(gas price / ethanol price)^1	2.41	2.42	2.32	2.77	2.79	2.30	2.58	3.28
	(0.56)	(0.54)	(0.53)	(0.26)	(0.24)	(0.54)	(0.55)	(1.17)
ln(gas price / ethanol price)^2							-3.22	
							(0.86)	
ln(gas price / ethanol price)^3							-1.97	
							(3.65)	
gas price minus ethanol price)								-0.49
								(0.60)
ln(number flex-fuel vehicles)	0.08	0.07	0.07	0.04	0.05		0.08	0.08
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)		(0.03)	(0.03)
ln(number ethanol stations)	-0.13	-0.18	-0.19	-0.08	-0.09		-0.12	-0.14
	(0.08)	(0.07)	(0.07)	(0.06)	(0.06)		(0.08)	(0.08)
month 1 of operation	-0.99	-0.91	-0.90	-0.91	-0.91		-0.95	-0.99
	(0.21)	(0.14)	(0.12)	(0.22)	(0.13)		(0.21)	(0.21)
month 2 of operation	-0.14	-0.13	-0.12	-0.11	-0.17		-0.12	-0.14
	(0.16)	(0.10)	(0.09)	(0.16)	(0.09)		(0.16)	(0.16)
month 3 of operation	-0.00	-0.01	-0.02	0.01	-0.04		0.00	-0.00
	(0.11)	(0.08)	(0.07)	(0.11)	(0.08)		(0.11)	(0.11)
month 4 of operation	0.04	0.04	0.04	0.06	0.04		0.04	0.04
	(0.06)	(0.05)	(0.04)	(0.06)	(0.05)		(0.06)	(0.06)
station-specific quadratic trend	x			x		x	x	x
county-specific quadratic trend		x						
month dummies	x	x	x			x	x	x
number observations	4332	4332	4332	4332	4332	4332	4332	4332
number of stations	201	201	201	201	201	201	201	201
R-squared	0.34	0.30	0.28	0.25	0.18	0.30	0.36	0.35
first-order correlation of residuals	-0.27	-0.24	-0.23	-0.22	-0.18	-0.25	-0.27	-0.27
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)

Note: Dependent variable is logged monthly ethanol sales volume in gallons. Logged price variables have been normalized to equal zero at sample mean prices. Standard errors in parentheses are clustered by station. Regressions control for station effects through first differencing. Regressions do not include a constant term, which is removed by first differencing. Regression (1) is identical to the linear first-difference regression above in table 2.2 and is included for comparison. Regression (2) replaces the station-specific quadratic trend with a county-specific quadratic trend. Regression (3) drops all quadratic trends. Regression (4) drops month dummy variables. Regression (5) drops all time controls. Regression (6) drops station and county characteristics. Regression (7) is identical to the cubic first-difference regression above in table 2.2. Regression (8) adds the retail gasoline minus ethanol price. First-order correlation of residuals is the coefficient from the least-squares regression of residuals on their lagged values. See text for details.

Table 2.7: Estimation results—dynamic responses

	Fixed effects				First differences			
Variable	L0	L1	L2	L3	L0	L1	L2	L3
ln(price ethanol)	0.48	1.06	1.09	1.43	0.77	0.86	0.97	1.16
	(0.64)	(0.60)	(0.64)	(0.69)	(0.48)	(0.47)	(0.53)	(0.54)
ln(price ethanol)$_{t-1}$		-0.62	-0.26	-0.37		-0.22	-0.21	-0.20
		(0.77)	(0.76)	(0.81)		(0.59)	(0.57)	(0.60)
ln(price ethanol)$_{t-2}$			-0.59	-0.03			-0.12	0.08
			(0.66)	(0.77)			(0.63)	(0.56)
ln(price ethanol)$_{t-3}$				-0.10				0.25
				(0.76)				(0.59)
ln(gas price / ethanol price)^1	2.60	2.57	2.59	2.86	2.41	2.46	2.63	2.77
	(0.64)	(0.57)	(0.59)	(0.63)	(0.56)	(0.59)	(0.61)	(0.67)
ln(gas price / ethanol price)$^1_{t-1}$		0.27	0.27	0.12		0.23	0.19	0.20
		(0.78)	(0.78)	(0.81)		(0.62)	(0.63)	(0.65)
ln(gas price / ethanol price)$^1_{t-2}$			-0.15	0.24			0.05	0.26
			(0.63)	(0.72)			(0.62)	(0.56)
ln(gas price / ethanol price)$^1_{t-3}$				0.29				0.46
				(0.72)				(0.56)
number observations	5027	4332	3843	3462	4332	3843	3462	3161
number of stations	237	201	174	156	201	174	156	140
R-squared	0.84	0.85	0.86	0.87	0.34	0.28	0.29	0.28
ln(price ethanol) cumulative	0.48	0.44	0.24	0.92	0.77	0.64	0.65	1.29
	(0.64)	(1.17)	(1.70)	(2.27)	(0.48)	(0.82)	(1.25)	(1.34)
ln(gas price/ethanol price) cumulative	2.60	2.84	2.71	3.51	2.41	2.70	2.87	3.69
	(0.64)	(1.11)	(1.59)	(2.12)	(0.56)	(0.95)	(1.32)	(1.48)
first-order correlation of residuals	0.39	0.40	0.40	0.40	-0.27	-0.31	-0.33	-0.35
	(0.03)	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.03)	(0.03)

Note: Dependent variable is logged monthly ethanol sales volume in gallons. Logged price variables have been normalized to equal zero at sample mean prices. Standard errors in parentheses are clustered by station. Fixed-effects regressions use deviations from within-station means to control for station effects. First-difference regressions do not include a constant term, which is removed by first differencing. All regressions include logged number of flexible-fuel vehicles in county, logged number of ethanol stations in county, dummy variables for months 1-4 of operation, month dummy variables, and station-specific quadratic trends. Accumulated price responses near the bottom of the table are long-run elasticities given by the sum of the coefficients on the current and lagged price variables. First-order correlation of residuals is the coefficient from the least-squares regression of residuals on their lagged values. See text for details.

CHAPTER III

The Market for Flexible-Fuel Vehicles That Burn Ethanol

3.1 Background

Estimating the cost of regulation is difficult. Few regulations allow trading that could reveal compliance costs through transaction prices, and regulated firms often have little incentive to report costs truthfully. Some regulations, however, feature "loopholes" that allow firms to relax regulatory constraints. When the cost of using a loophole is known, researchers can infer the marginal cost of regulation indirectly for firms that exploit the loophole. My coauthor and I demonstrate that firms in the auto industry reveal the marginal cost of complying with fuel-economy standards when they exploit a loophole that overstates the efficiency of "flexible-fuel" vehicles. Using this approach, we estimate that tightening fuel-economy standards by one mile per gallon would cost domestic automakers no more than \$10–\$30 in profit per vehicle, which is consistent with other recent attempts to measure these costs directly. Our estimates are significantly less than the fine of \$55 for non-compliance. Researchers have used this fine as a measure of compliance costs in the past, even though domestic automakers do not pay fines.

Corporate Average Fuel Economy (CAFE) standards require automakers to achieve a minimum average fuel economy across their entire vehicle fleet. Firms whose fleet average falls below the minimum are subject to a fine. The Alternative Motor Fuels Act (AMFA) modified CAFE regulations starting in 1993 by crediting vehicles capable of burning gasoline and ethanol with about two-thirds better mileage than they actually achieve. These vehicles are known as "flexible-fuel" vehicles.[1] Automakers can make any conventional vehicle a flexible-fuel vehicle through a minor modification, which adds only \$100–\$200 in production cost, as we discuss in detail below. If consumers fill their tanks with gasoline instead of ethanol, a flexible-fuel vehicle is virtually identical to its gasoline counterpart, with negligible differences in performance. The flexible-fuel provision therefore offer automakers a low-cost option for improving fuel economy under CAFE regulations, reducing the need to make vehicles more efficient by adding costly fuel-saving technology or to improve average mileage by selling a larger fraction of small vehicles. The flexible-fuel provision limits the gain in fuel economy that an automaker can achieve using flexible-fuel vehicles to 1.2 miles per gallon.

The original rationale for the flexible-fuel provision was to solve a classic coordination problem. Without vehicles that ran on ethanol, it was thought, retailers would have no incentive to install new fuel pumps for distributing ethanol. Without pumps, consumers would never demand vehicles that burn the fuel. Policymakers hoped that the flexible-fuel provision would solve this coordination problem. By inducing automakers to make flexible-fuel vehicles through a CAFE credit, the provision would put alternative-fuel vehicles on the road, and infrastructure would follow. In reality, ethanol infrastructure has not kept pace with flexible-fuel production, and few flexible-fuel vehicles ever run on ethanol. For this reason, the National Academy of Sciences has advocated eliminating the flexible-fuel provision (National Academy of Sciences 2002), and critics have characterized the provision as a harmful loophole. Because the incremental cost of flexible-fuel capacity is known, however, automakers that exploit this loophole indirectly reveal infor-

[1] Flexible-fuel vehicles actually run on a fuel blend known as "E85," which contains 85% pure ethanol and 15% gasoline. We refer to this fuel as "ethanol" throughout.

mation about the cost of CAFE standards. The purpose of this chapter is to use this insight to estimate the cost of marginally tightening CAFE standards for automakers that produce flexible-fuel vehicles.

We begin by modeling the profit-maximization decision of an oligopolistic automaker. The automaker faces a fuel-economy constraint but can relax the constraint by producing flexible-fuel vehicles. The model provides sufficient conditions under which we can infer the marginal cost of tightening CAFE standards by examining the cost of exploiting the flexible-fuel loophole. If an automaker is constrained by CAFE standards, has installed flexible-fuel capacity on some but not all units for some model, and has not exhausted the maximum gain in fuel economy from producing flexible-fuel vehicles, and if marginal consumers do not value flexible-fuel capacity, then the automaker will equate the marginal cost of improving mileage using the flexible-fuel loophole with the marginal cost of improving mileage through other means.

The empirical portion of the chapter demonstrates that these conditions hold for domestic automakers. Using administrative data from the Department of Transportation, we show that domestic automakers were constrained by CAFE standards and used flexible-fuel vehicles to comply with the regulations. Domestic automakers rarely added flexible-fuel capacity to more than one type of vehicle, and unconstrained Japanese automakers did not produce any flexible-fuel vehicles. These findings are consistent with automakers producing flexible-fuel vehicles to exploit the flexible-fuel loophole. Automakers that produced flexible-fuel vehicles installed flexible-fuel capacity on some but not all units, and automakers rarely exceeded the maximum gain in fuel economy permitted under the provision.

Next, we show that marginal consumers do not value flexible-fuel capacity, using transaction data to analyze both prices and quantities for flexible-fuel vehicles. Automakers sell a large portion of their flexible-fuel vehicles to consumers living in states with virtually no ethanol fueling stations. Consumers in these states are almost certainly not willing to pay more for flexible-fuel vehicles, since they are not able to purchase ethanol. Furthermore, our analysis of transaction prices for flexible-fuel vehicles and comparable gasoline-only vehicles indicates that consumers do not pay more for flexible-fuel capacity, which is consistent with survey evidence that most flexible-fuel owners do not know they own flexible-fuel vehicles.

Because marginal consumers do not value flexible-fuel capacity, and because domestic automakers have exploited the flexible-fuel loophole without exhausting it, the flexible-fuel provision reveals the cost of marginally increasing CAFE standards. The cost of CAFE is a function of vehicle fuel economy and the cost of adding flexible-fuel capacity. Incremental production costs for flexible-fuel vehicles reportedly range from \$100–\$200 or lower. For automakers that produce flexible-fuel vehicles to comply with CAFE standards, this range implies that the marginal cost of tightening the standard for light trucks by one mile per gallon is no more than \$11–\$28 in lost profit per vehicle. The cost of tightening the standard for passenger cars is no more than \$8–\$18. Because the automaker equates the marginal costs of alternative compliance strategies, our cost estimates also reflect lower profit margins on smaller, more efficient vehicles, as well as the gap between incremental production costs and willingness to pay for fuel-saving modifications. Costs are substantially lower than the \$55 fine that automakers would pay if they failed to comply with CAFE standards, which serves as an upper bound on marginal compliance costs.

Our estimate of the marginal cost of tightening CAFE standards is almost identical to recent estimates by Jacobson (2007), who uses a wholly different methodology. Jacobson measures compliance costs by directly estimating demand elasticities and implied marginal costs, in contrast to our loophole approach. We find the similarity of his and our results reassuring.

Our estimates should prove useful in any analysis of the costs of CAFE standards or the benefits of allowing trading of CAFE credits across firms. Our estimates would also be useful in research comparing CAFE to policy alternatives, such as a gasoline tax or fee-bates, which impose fees on inefficient vehicles and offer rebates for efficient ones.[2] More broadly, our approach can inform research measuring the cost of regulation in other industries. A prominent example is "incentive zoning." Zoning regulations typically constrain the height and density of new buildings in a jurisdiction. Under incentive zoning, these constraints are relaxed if developers provide open space, affordable housing, or other public goods.[3] Following our methodology, researchers could estimate the marginal cost to developers of changing height or density restrictions by quantifying how much developers spend on plazas or affordable housing.

Our analysis of the market for flexible-fuel vehicles also contributes to the policy debate and growing literature on alternative fuels policies. Congress renewed the flexible-fuel provision, in spite of the provision's critics, as part of the 2005 energy bill. Liu and Helfand (2008) demonstrate that the flexible-fuel incentive likely increases gasoline consumption and greenhouse emissions because it allows automakers to make inefficient vehicles, but that increasing the incentive could actually reduce production of flexible-fuel vehicles. They do not use the flexible-fuel incentive to estimate the cost of complying with CAFE. The 2005 energy bill also established a renewable fuel standard, which Congress greatly expanded in 2007. The second chapter of this dissertation estimates ethanol demand by owners of flexible-fuel vehicles and uses the demand estimates to evaluate the welfare

[2]Liu and Helfand (2008) assume marginal compliance costs are \$55 in their analysis of AMFA, and Goldberg (1998) assumes that marginal compliance costs are zero if an automaker does not pay a fine and \$55 if it does. These analyses could be improved using our estimates.

[3]Incentive zoning began in Chicago and New York City, where developers were allowed to exceed height restrictions and density limitations on buildings if they provided plaza space on the lot (Weiss 1992; Morris 2000). At least half of all cities and towns with zoning laws reportedly have some incentive zoning program (Morris 2000).

implications of the expanded renewable fuel standard.

The format of this chapter is as follows. Section 3.2 models an automaker's decision to use the flexible-fuel loophole to relax fuel-economy constraints, discusses which vehicles automakers will equip with flexible-fuel capacity, and establishes conditions under which we can infer marginal compliance costs. The next several sections demonstrate that these conditions hold empirically.

Section 3.3 shows that domestic automakers use flexible-fuel vehicles to comply with CAFE standards, that they install flexible-fuel capacity on some but not all units for their flexible-fuel models, and that they have not exhausted the flexible-fuel loophole. Section 3.3 also shows that the set of vehicles we observe with flexible-fuel capacity is broadly consistent with our model's predictions. Section 3.4 argues that marginal consumers do not value flexible-fuel capacity. Section 3.5 then uses publicly reported estimates of the incremental cost of producing a flexible-fuel vehicle to calculate the marginal cost of complying with fuel-economy standards.

3.2 Revealing the cost of fuel-economy standards

3.2.1 Profit-maximizing automaker

We assume that an oligopolistic automaker complying with fuel-economy standards maximizes profits with respect to the prices, mileage, and flexible-fuel shares of the models it produces:

$$\pi = \sum_{j \in \mathcal{M}} \Big(p_j - c_j(m_j) - \alpha_j \theta_j \Big) q_j(p,m) - \sum_{j \in \mathcal{M}} I(\theta_j > 0) \cdot F_j \tag{3.1}$$

where: $\mathcal{M}$ is the set of models the automaker produces; p_j is the price the automaker charges for model j; m_j is the model's fuel economy in miles per gallon; q_j is its sales quantity, which depends on the prices p and mileage m for all models of all producers; c_j is the constant marginal cost of the gasoline-only version of the model, which depends on the model's mileage; $\theta_j \in [0,1]$ is the model's flexible-fuel share, or the fraction of units with flexible-fuel capacity; α_j is the incremental production cost of outfitting one such unit with flexible-fuel capacity; and F_j is the fixed cost of engineering the model to have flexible-fuel capacity, which the automaker pays if the model's flexible-fuel share exceeds zero, as denoted by the indicator function $I(\theta_j > 0)$. Profits equal the sum over all models of price minus average variable cost multiplied by quantity, minus engineering fixed costs.[4] Following previous studies in this literature we assume that the set of models is fixed.

It is convenient to interpret α_j as the incremental production cost of adding special components that allow a flexible-fuel vehicle to burn ethanol. In addition to having larger fuel injectors, flexible-fuel vehicles have fuel-system components made from materials that are more resistant to the corrosive nature of ethanol. Earlier models also had special fuel sensors to detect the percent ethanol content of fuel running through the engine. Incremental costs vary from model to model, depending on a model's engine technology and sales volume.[5] Often more important than the hardware changes themselves, however, is the engineering time and effort needed to add flexible-fuel capacity. In addition to making minor design changes, outfitting a new model with flexible-fuel capacity requires modifying on-board software, doing additional engine calibration work, and performing extra emissions testing. These up-front fixed costs can be substantial.

We assume that each model has a unique engine size or technology, leading to a separate fixed cost for each model. In reality, different models often share the same engines, implying substantial overlap in fixed costs. Thus, when we analyze actual flexible-fuel production below, we focus on flexible-fuel shares for specific engine sizes, which proxy

[4]We could easily generalize the model by making marginal costs an increasing function of quantity. This would not alter any of the conclusions, however, and would complicate the notation.

[5]High sales volumes allow automakers to attract multiple parts suppliers for flexible-fuel components, which can bring down incremental costs through competition.

for models with shared fixed costs.

An implicit assumption in the model is that consumers do not care about flexible-fuel capacity one way or the other. Quantities do not depend on flexible-fuel shares, which implies, for example, that no consumer would switch from a Honda Accord to a Chevy Impala if General Motors increased the fraction of Impalas with flexible-fuel capacity. Similarly, we do not include separate prices for flexible-fuel vehicles and their gasoline-only counterparts. Since consumers regard the vehicles as perfect substitutes, no consumer would pay more or less for an Impala with flexible-fuel capacity, and the automaker sets a single price for all Impalas. In reality, some consumers surely prefer flexible-fuel vehicles, while other consumers may even have a distaste for such vehicles. We think the fraction of such consumers is small, however, given the dearth of ethanol pumps nationwide. We therefore ignore these consumers momentarily and proceed as if no consumer cared about flexible-fuel capacity. After deriving our results, we argue that our key result—that flexible-fuel costs reveal the marginal cost tightening CAFE standards—holds as long as marginal consumers do not value flexible-fuel capacity. Later, we present empirical evidence that marginal consumers indeed do not value flexible-fuel capacity.

3.2.2 Fuel-economy standards

The automaker faces a fuel-economy constraint that sets a minimum average mileage for the automaker's fleet. The constraint features a "loophole" that gives extra credit for flexible-fuel vehicles. The constraint takes the following form:

$$1 \Big/ \left(\sum_{j \in \mathcal{M}} \frac{q_j(p,m)}{Q} \cdot \frac{\theta_j \beta + (1-\theta_j)}{m_j} \right) - \sigma \geq 0, \tag{3.2}$$

where σ is the fuel-economy standard in miles per gallon, m_j is the mileage of model j, $\beta \in [0,1]$ is the incentive for flexible-fuel vehicles, $Q = \sum_{j \in \mathcal{M}} q_j(p,m)$ is the automaker's total sales volume, and all other parameters are as above. The constraint requires that an automaker's AMFA fuel economy—that is, the sales-weighted harmonic-average mileage of the automaker's vehicles, calculated using flexible-fuel incentives—exceed the CAFE standard of σ. Equivalently, the constraint prevents sales-weighted average fuel consumption per mile, calculated using flexible-fuel incentives, from exceeding $1/\sigma$. Current legislation fixes the flexible-fuel incentive at $\beta \approx 0.6$, giving automakers with binding constraints a strong implicit subsidy to produce flexible-fuel vehicles.[6] For a sense of how strong this incentive is, note that adding flexible-fuel capacity increases a vehicle's effective mileage by about $1/0.6 - 1 \approx 67\%$, which amounts to treating a flexible-fuel Hummer like a Toyota Camry or a flexible-fuel Camry like a Toyota Prius. Increasing a model's flexible-fuel share increases average mileage because the standard treats flexible-fuel vehicles as though they achieve better mileage than they actually do.

It is convenient to rewrite this first constraint as follows:

$$1/\left(\sum_{j\in\mathcal{M}} \frac{q_j(p,m)}{Q}\frac{1}{m_j} - (1-\beta)\sum_{j\in\mathcal{M}} \frac{q_j(p,m)}{Q}\frac{1}{m_j}\theta_j\right) - \sigma \geq 0, \tag{3.3}$$

which clarifies that flexible-fuel vehicles relax the constraint by reducing sales-weighted average fuel consumption per mile.

The automaker is limited in its ability to improve fuel economy using the flexible-fuel loophole. This limit acts like a "backstop" on actual fuel economy by adding a second constraint:

$$1/\left(\sum_{j\in\mathcal{M}} \frac{q_j(p,m)}{Q}\cdot\frac{1}{m_j}\right) - (\sigma - \phi) \geq 0, \tag{3.4}$$

where $\phi > 0$ is the limit on using the flexible-fuel incentive, and all other parameters are as above. This constraint requires that actual sales-weighted harmonic-average mileage

[6] In practice $\beta = arg + (1-a)$, where $a \in [0,1]$ is the assumed fraction of miles that the vehicle drives using E85 ethanol, $r > 1$ is the ratio of ethanol to gasoline fuel consumption per mile, and $g \in [0,1]$ is the assumed gasoline content of E85. The credit's logic is that it purports to count only gasoline consumption when determining a vehicle's contribution toward average fuel economy. Current legislation fixes $a = 0.50$, which dramatically overstates the fraction of miles that flexible-fuel vehicles actually run on ethanol, and sets $g = 0.15$, which is the fraction gasoline content of E85. In practice r varies slightly among flexible-fuel vehicles, averaging about 1.35, which implies that flexible-fuel vehicles achieve about 35% higher fuel economy on gasoline or 1-1/1.35 = 25% lower fuel economy on ethanol. We assume for simplicity that r is the same for all vehicles so that β is also the same for all vehicles.

exceed the less-stringent standard of $\sigma - \phi < \sigma$, or that sales-weighted average fuel consumption per mile not exceed $1/(\sigma - \phi) > 1/\sigma$. Equivalently, the constraint requires that the automaker's actual fuel economy not fall short of the nominal fuel-economy standard by more than ϕ miles per gallon. Legislation fixes this limit at $\phi = 1.2$ miles per gallon.[7]

It is helpful to think of the automaker as solving a two-stage maximization problem. First, the automaker pays the fixed costs to engineer flexible-fuel capacity on whichever models it chooses. Then the automaker sets flexible-fuel shares for these models. Variable profits in the second stage depend on the combination of models engineered to be flexible-fuel capable in the first stage. Thus, the automaker chooses this combination of models optimally in the first stage to maximize second-stage variable profits minus first-stage fixed costs. We remain agnostic as to the competitive behavior automakers use to arrive at an equilibrium in vehicle prices, quantities, and fuel economy. We simply assume that some equilibrium mapping from prices and mileage to sales quantities exists and that automakers choose flexible-fuel shares optimally given this mapping.

The Lagrangian for the automaker's second-stage maximization problem is given by:

$$
\begin{aligned}
\mathcal{L} = & \sum_{j \in \mathcal{M}} \Big(p_j - c_j - \alpha_j \theta_j\Big) q_j \\
& + \lambda \left[1 \Big/ \left(\sum_{j \in \mathcal{M}} \frac{q_j}{Q} \frac{1}{m_j} - (1-\beta) \sum_{j \in \mathcal{M}} \frac{q_j}{Q} \frac{1}{m_j} \theta_j \right) - \sigma \right] \\
& + \mu \left[1 \Big/ \left(\sum_{j \in \mathcal{M}} \frac{q_j}{Q} \cdot \frac{1}{m_j} \right) - (\sigma - \phi) \right],
\end{aligned} \tag{3.5}
$$

where λ and μ are the shadow prices on the constraints, all other variables are as above, and we have suppressed the arguments of functions for convenience. Flexible-fuel shares are choice variables only for models on which the automaker has paid the fixed engineering costs; flexible-fuel shares are zero for other models. When the constraints are binding, the

[7] It is strange that the standard regulates fuel consumption per mile, yet caps the flexible-fuel incentive in miles per gallon. For automakers that have a binding fuel-economy constraint, however, the two are equivalent. It is similarly strange that penalties for non-compliance scale proportionally with miles per gallon below the standard. This implies that the penalty per gallon of fuel consumption is higher in the passenger car fleet, which has a higher fuel-economy standard.

shadow prices implicitly tax inefficient models and subsidize efficient models. The shadow prices also quantify the marginal cost, in terms of lower profits, resulting from tighter fuel-economy standards. Equivalently, the shadow prices quantify the marginal benefit of looser standards. We estimate the first of these shadow prices by examining the loophole for flexible-fuel vehicles. This shadow price is revealed by the first-order conditions for flexible-fuel shares immediately below. We show empirically that the second shadow price is usually zero.

3.2.3 First-order conditions for flexible-fuel shares

Differentiating the Lagrangian with respect to a model's flexible-fuel share leads to the following first-order condition:

$$-\alpha_k + \lambda \frac{1-\beta}{m_k Q} M^2 = 0. \tag{3.6}$$

where q_k factors out of both terms, and M is the automaker's sales-weighted harmonic-average mileage calculated using flexible-fuel incentives, which is given by the first term in equation (3.2). This first-order condition holds with equality for models whose flexible-fuel shares are strictly greater than zero and strictly less than one. At corner solutions the equality becomes an inequality. The first term is the incremental cost of flexible-fuel capacity. In the second term, β is the share of a flexible-fuel vehicle's fuel consumption per mile that contributes toward the automaker's fleet average, so $(1-\beta)/(m_k Q)$ is the reduction in average fuel consumption per mile that the automaker achieves when it adds flexible-fuel capacity to another unit. Multiplying by M^2 converts this value to a marginal improvement in mileage, while multiplying by the shadow price on the first constraint λ converts this improvement into dollars of marginal benefits. The automaker simply equates the incremental cost of flexible-fuel capacity with the marginal benefit of a flexible-fuel vehicle in terms of relaxing the first constraint.

Note that the assumption that consumers ignore flexible-fuel capacity is critical here. If consumers valued flexible-fuel capacity, then this first-order condition would contain additional terms reflecting marginal revenue changes.

The first-order conditions for flexible-fuel shares reveal the shadow price on the first constraint, which is the key insight of this chapter. Rearranging equation (3.6) gives:

$$\lambda = \frac{\alpha_k}{1/m_k} \frac{Q}{(1-\beta)M^2}, \tag{3.7}$$

which holds with equality for models at an interior flexible-fuel share. The shadow price λ on the first constraint equals the incremental cost of adding flexible-fuel capacity divided by the corresponding improvement in AMFA fuel economy that flexible-fuel capacity affords. The automaker equates the marginal benefit of relaxing the constraint with the marginal cost of relaxing the constraint using the flexible-fuel loophole. This equation holds regardless of whether the second constraint is binding or not. Again, at corner solutions this equality becomes an inequality.

While our mathematical model assumes that consumers do not value flexible-fuel capacity, the key result in equation (3.7) will hold as long as marginal consumers do not value flexible-fuel capacity. That is, suppose that prices for flexible-fuel vehicles and their gasoline counterparts are the same in equilibrium, and that sales quantities do not change when an automaker marginally increases the flexible-fuel share on one of its flexible-fuel models. Then a constrained automaker will still equate the marginal cost of a flexible-fuel vehicle with the marginal benefit in terms of relaxing the first constraint, even if some consumers prefer flexible-fuel vehicles or have a distaste for flexible-fuel capacity. We demonstrate in section 3.4 below that marginal consumers indeed do not value flexible-fuel capacity.

3.2.4 Revealing marginal compliance costs

If an automaker produces flexible-fuel vehicles to comply with CAFE standards, is at an interior flexible-fuel share for some model, and does not exhaust the flexible-fuel loophole, and if in equilibrium marginal consumers do not value flexible-fuel capacity, then we are able to pin down the cost of marginally tightening CAFE standards. Differentiating the automaker's Lagrangian in equation (3.5) at the optimum with respect to the nominal fuel-economy standard gives marginal compliance costs in terms of lost profit:

$$\frac{\partial \mathcal{L}^*}{\partial \sigma} = -\lambda - \mu. \tag{3.8}$$

Assuming the automaker does not exhaust the flexible-fuel loophole, the shadow price on the backstop constraint μ is zero, and we can ignore the second term. Marginal compliance costs then equal the shadow price on the first constraint only. It does not matter on the margin whether the limit on on using the flexible-fuel loophole ϕ changes or stays the same, as the backstop constraint is not binding.[8]

Substituting for the shadow price using equation (3.7) and then dividing by total production yields marginal compliance costs per vehicle as a function of known parameters:

$$\frac{\partial \mathcal{L}^*}{\partial \sigma}\frac{1}{Q} = -\frac{\alpha_k \cdot m_k}{(1-\beta)\sigma^2}, \tag{3.9}$$

where we have replaced average mileage M with the fuel-economy standard σ because the first constraint is binding. Marginal compliance costs are then a simple function of mileage and the incremental cost of adding flexible-fuel capacity for any model with an interior flexible-fuel share. The automaker equates the marginal cost of relaxing the constraint using the flexible-fuel loophole with the marginal cost of relaxing the constraint through other means, such as by directly improving mileage or by selling a larger share of small

[8]Note that we calculate the marginal cost of tightening the CAFE standard σ while holding the limit on using the flexible-fuel loophole ϕ constant. Other policy changes are possible and in general have different costs. For example, the marginal cost of reducing the limit ϕ is μ, which we are not able to estimate using our methodology. The marginal cost of tightening the CAFE standard σ while holding backstop fuel economy $\sigma - \phi$ constant is λ, regardless of whether the backstop is binding or not.

vehicles. Constrained automakers that exploit the flexible-fuel loophole therefore reveal their marginal compliance costs, as long as they do not exhaust the loophole by running up against the backstop constraint. If the backstop constraint were binding, then marginal compliance costs would also depend on the second shadow price, which we are not able to estimate using our methodology.

Observe that λ is the marginal benefit of relaxing the first constraint, which sets a minimum AMFA fuel economy. Note, however, that AMFA fuel economy differs from actual fuel economy, as is clear from the first constraint in equation (3.3). Thus, while our methodology yields the marginal cost of improving AMFA fuel economy, it does not necessarily give the marginal cost of improving actual fuel economy, which may be a parameter of interest. While not exact, however, our estimates closely approximate the marginal cost of improving actual fuel economy when the automaker produces a small number of flexible-fuel vehicles. This is clear from the constraint in equation (3.3), where the difference between AMFA fuel economy and actual fuel economy shrinks to zero with sales quantities for flexible-fuel vehicles.

More formally, suppose the automaker produces only one type of vehicle. Then the first constraint weighted by its shadow price simplifies to

$$\lambda \left[\frac{m}{\theta\beta + (1-\theta)} - \sigma \right], \tag{3.10}$$

where m is the automaker's actual mileage, θ is its flexible-fuel share, and the first term inside the brackets is the automaker's AMFA mileage. Differentiating with respect to actual mileage gives:

$$\frac{\lambda}{\theta\beta + (1-\theta)}, \tag{3.11}$$

or the marginal benefit of relaxing the constraint by improving actual mileage, which the automaker will set equal to marginal costs. Suppose that θ is small, say 0.15, which is the maximum flexible-fuel share for a binding light-truck standard of $\sigma = 20.7$ miles per gallon and maximum fuel economy gain of $\phi = 1.2$ miles per gallon. Then the marginal cost of improving actual fuel economy exceeds the marginal cost of improving AMFA fuel economy by a factor of just $1/[0.15 \cdot 0.6 + (1-0.15)] \approx 1.06$. The maximum flexible-fuel share for cars is even lower than 0.15, and in practice flexible-fuel shares average less than 0.06 during our study period.

There are two cases in which we are able to bound marginal compliance costs, even though we are unable to infer costs precisely. First, if the backstop constraint is binding, then the cost of improving fuel economy using the flexible-fuel loophole gives a lower bound on marginal compliance costs. The shadow price on the backstop constraint is nonzero, and while no first-order condition reveals this shadow price, it must be positive. Because the automaker complies with the fuel-economy standard and does not pay fines, we also know that costs are bounded above by the level of the fine.[9] Second, if a constrained automaker produces no flexible-fuel vehicles, and if fixed engineering costs are zero, then the cost of improving fuel economy using flexible-fuel vehicles gives an upper bound on marginal compliance costs. The shadow price on the backstop is zero, but the first-order conditions for flexible-fuel shares do not hold with equality. This upper bound on marginal compliance costs does not hold for a general model with fixed costs for adding flexible-fuel capacity.

3.2.5 Which models get flexible-fuel capacity?

In general, the combination of models with flexible-fuel capacity depends on fuel-economy standards, fixed and incremental costs for flexible-fuel vehicles, consumer demand, and production costs for different models of varying efficiency. Nevertheless, some

[9]Technically, not complying with CAFE standards is a civil infraction. Some analysts suggest that domestic automakers fear that breaking the law would make them liable for damages to stockholders, meaning that the cost of non-compliance may exceed the $55 fine (Kleit 2002). Jacobson (2007) models this legal risk and potential loss of reputation as a fixed cost of non-compliance.

generic results are possible. We begin by assuming that the automaker has already paid the fixed engineering costs on some subset of models. What can we say about flexible-fuel shares for this subset?

Returning to equation (3.7), note that the first-order conditions at corners imply that the shadow price on the first constraint is less than incremental costs per mileage gain for models whose flexible-fuel shares are zero. The shadow price exceeds incremental costs per mileage gain for models whose flexible-fuel shares are one. These observations imply a particular ordering of flexible-fuel shares by incremental costs. Focusing on the subset of models on which the automaker has paid the fixed engineering costs, suppose that some model has a nonzero flexible-fuel share. Then any other model with a lower incremental cost per fuel consumption in gallons per mile will have a flexible-fuel share of one. Similarly, if any model has an interior flexible-fuel share, any other model with a higher incremental cost per fuel consumption will have a flexible-fuel share of zero. In other words, the automaker begins installing flexible-fuel capacity on a new model only after installing flexible-fuel capacity on all units for models that have lower incremental costs per fuel consumption.[10] The intuition for this result is simple. A flexible-fuel vehicle's impact on average efficiency is proportional to the vehicle's fuel consumption. The automaker simply adds flexible-fuel capacity in order of ascending cost per impact. Note that if incremental costs are the same for all models or increase less than proportionally with fuel consumption, then the automaker will install flexible-fuel capacity on its most inefficient vehicles first.

Clearly, if the flexible-fuel share for any model is zero in the second stage, then the automaker would never have chosen to pay the fixed engineering costs in the first place.

[10]Formally, suppose that $\alpha_k/(1/m_k) < \alpha_l/(1/m_l)$ and that $\theta_l > 0$. Then it must be that $\theta_k = 1$. Why? Suppose instead that $\theta_k < 1$. Then the first-order conditions imply that $\lambda \leq \alpha_k Q/((1/m_k)(1-\beta)M^2)$, since $\theta_k < 1$, and that $\lambda \geq \alpha_l Q/((1/m_l)(1-\beta)M^2)$, since $\theta_l > 0$. This implies that $\alpha_l/(1/m_l) \leq \alpha_k/(1/m_k)$, which is a contradiction. By similar arguments, if $\alpha_j/(1/m_j) > \alpha_i/(1/m_i)$ and $\theta_i \in (0,1)$, then $\theta_j = 0$.

This leads to a brief discussion of fixed costs and the combination of models engineered to receive flexible-fuel capacity. If fixed costs for flexible-fuel vehicles are zero, then the automaker's two-stage maximization problem reduces to the second-stage problem above. The automaker simply adds flexible-fuel capacity in order of ascending incremental cost per fuel consumption, as we show above.

Suppose instead that fixed costs are positive but that the automaker could fully exploit the flexible-fuel incentive for any single model outfitted with flexible-fuel capacity. That is, for any single model engineered to have flexible-fuel capacity, optimal second-stage behavior would have the automaker producing an interior flexible-fuel share, either because extra flexible-fuel vehicles were no longer useful to meet the first CAFE constraint or because the automaker was up against the backstop constraint. Then, if the automaker produces any flexible-fuel vehicles at all, it will install flexible-fuel capacity only on one model, irrespective of fixed costs. Why? Suppose the automaker paid the fixed cost on two or more models. The results above tell us that the automaker would apply flexible-fuel capacity in order of ascending cost per fuel consumption. Given the assumption that any single model is sufficient, the automaker would never apply flexible-fuel capacity to a second model, and it never would have paid the fixed costs in the first place. If we additionally assume that fixed costs are the same for all models, then the automaker would only install flexible-fuel capacity on the model with the lowest incremental cost per fuel consumption. Any substitute vehicle would have the same fixed cost, by assumption, but the same gain in fuel economy would be more costly for other models, which have higher incremental costs per fuel consumption.

The assumptions and qualifications above hint at several additional points. First, the gain in variable profit from installing flexible-fuel capacity on any single model must exceed its corresponding fixed cost, otherwise the automaker could not justify paying the fixed cost. This point holds more generally. For any optimal combination of models with flexible-fuel capacity, the loss in variable profit from excluding any subgroup of models must exceed the subgroup's collective fixed costs, by the definition of an optimum.

Second, automakers will clearly avoid installing flexible-fuel capacity on models with especially high fixed costs, even if incremental production costs are low. Similarly, low fixed costs can compensate for high incremental costs.

Finally, automakers will tend to avoid installing flexible-fuel capacity on models whose sales volumes are low, even if such models have relatively low fixed or incremental production costs. Suppose, for example, that a model's sales share were near zero. Then even if the automaker installed flexible-fuel capacity on every unit, the impact on average fuel economy would be negligible. This is clear from the constraint in equation (3.3), where the impact of flexible-fuel capacity is weighted by a model's sales share. The automaker might rather install flexible-fuel capacity on a model with high fixed costs but higher sales volume than pay fixed costs on a model yielding such a negligible gain in average mileage. Sales volumes also influence the decision to install flexible-fuel capacity in the first place. Small automakers may be unable to justify producing any flexible-fuel vehicles if sales volumes and profits are low relative to fixed engineering costs.

3.2.6 Additional considerations

Actual fuel-economy standards are more complicated than we describe above. First, automakers also receive extra credit for vehicles that burn natural gas, electricity, and other alternative fuels. These vehicles all contribute toward the backstop limit of 1.2 miles per gallon. We could model these vehicles explicitly, but it would not change our key result that the incremental cost of installing flexible-fuel capacity reveals marginal compliance costs for automakers satisfying the conditions we set forth. In practice, automakers produce few alternative-fuel vehicles besides flexible-fuel vehicles. Because ethanol flexible-fuel vehicles can use existing gasoline infrastructure, and because incremental costs are relatively low, flexible-fuel vehicles have proven more attractive to automakers than natural-gas or electric vehicles.

Second, fuel-economy standards regulate light-duty trucks and passenger cars separately. The nominal fuel-economy standard for passenger cars is 27.5 miles per gallon, while the standard for light trucks is 20.7 for most of our study period. Both fleets qualify for the same flexible-fuel incentive, and the limit of 1.2 miles per gallon applies to both fleets separately. Mathematically, this implies a constraint on AMFA fuel economy and corresponding shadow price for each fleet, as well as a backstop constraint and corresponding shadow price for each fleet. Weights are then given by by a model's sales share within its respective fleet. In what follows we distinguish between light-truck and passenger-car fleets. Each of the above results applies separately to the passenger-car and light-truck fleets.[11]

Finally, fuel-economy regulations allow "banking" and "borrowing." An automaker that exceeds the standard in one year earns credits that it can use to comply in an earlier or future year. For example, if an automaker falls short of the standard, it can use banked credits from a previous year to avoid paying fines. Banked credits expire after three years. If the automaker does not have banked credits, it can borrow credits in the short term and earn the credits back in a subsequent year. It must earn the credits back within three years to avoid paying fines.[12] Fines are \$55 for every mile per gallon below the standard and scale with total production. Credits earned for light trucks can not be applied to passenger cars, and vice-versa.

[11] Fuel-economy standards also regulate domestic and import passenger cars separately. This distinction is not thought to be particularly important for present-day automakers (see Jacobson (2007)). The one possible exception is Chrysler, which from 1999–2007 produced import passenger-car vehicles through its European subsidiary Mercedes-Benz and paid CAFE fines on this fleet in 2004–2006. Chrysler separated from Mercedes starting in 2008.

[12] Moreover, a model's "year" is itself a choice parameter that automakers can manipulate to comply with fuel-economy regulations. For example, suppose that Ford's lineup of trucks for the 2008 model year was relatively efficient. Ford could stop selling the 2007 version of its gas-guzzling Excursion early in the 2007 calendar year, and begin selling the 2008 version to include with its relatively efficient 2008 models.

3.3 Automakers exploit the flexible-fuel loophole

Recall from the previous section that we are able to infer the marginal cost of complying with CAFE standards as long as four conditions hold. First, constrained automakers must exploit the flexible-fuel loophole to comply with CAFE standards. Second, automakers must offer a model with an interior flexible-fuel share. Third, automakers must not exhaust the flexible-fuel loophole by hitting the backstop constraint. Fourth, and finally, marginal consumers must not value flexible-fuel capacity. We demonstrate that the first, second, and third of these conditions hold using administrative data from the Department of Transportation's National Highway Safety and Transportation Administration (NHTSA). These data record model names, production quantities, AMFA fuel economy, actual fuel economy, fuel type, and other vehicle attributes by model year. NHTSA collects these data to determine whether firms comply with CAFE standards. We demonstrate that the fourth condition holds using vehicle transaction data below.

3.3.1 Constrained automakers exploit the flexible-fuel loophole to comply with CAFE standards but do not exhaust the loophole

Table 3.1 summarizes fuel-economy performance and flexible-fuel production across automakers during 1993–2006. For both passenger-car and light-truck fleets, the table shows an automaker's actual fleet-average fuel economy ignoring the flexible-fuel incentive, the difference between actual fuel economy and the fuel-economy standard, the fraction of the automaker's vehicles that are flexible-fuel vehicles, and whether the automaker pays fines during 1993–2006. The table also shows each automaker's total production and market share during this time period, as well as the fraction of each automaker's production that is light trucks.

Automakers produce flexible-fuel vehicles only when necessary to comply with fuel-economy standards. Table 3.1 shows that all three domestic automakers produced vehicles

Table 3.1: Fuel-economy performance and flexible-fuel production 1993–2006

	Passenger cars				Light trucks				All vehicles		
Firm	Actual MPG	Over std.	% FFV	Paid fine?	Actual MPG	Over std.	% FFV	Paid fine?	Sales (mil.)	% mkt.	% truck
Domestic											
GM	27.9	0.4	0.0	no	20.5	-0.3	5.2	no	61.0	29.0	44.8
Ford	27.2	-0.3	2.5	no	20.3	-0.4	5.8	no	47.6	22.6	53.0
Chrysler	27.3	-0.2	0.9	yes	20.5	-0.3	6.0	no	33.3	15.8	65.8
European											
VW	28.8	1.3	0.0	no	19.9	-0.9	0.0	yes	3.9	1.8	3.4
BMW	25.8	-1.7	0.0	yes	20.6	-0.3	0.0	yes	2.6	1.2	13.4
Volvo	25.7	-1.8	0.0	yes					0.5	0.2	0.0
Porsche	23.8	-3.7	0.0	yes	18.5	-2.4	0.0	yes	0.3	0.1	18.3
Japanese											
Toyota	31.0	3.5	0.0	no	22.6	1.8	0.0	no	21.8	10.3	38.2
Honda	32.0	4.5	0.0	no	25.0	4.1	0.0	no	14.9	7.1	23.2
Nissan	29.1	1.6	0.0	no	21.5	0.7	1.8	no	11.0	5.2	39.0
Total	28.6	1.1	0.6		20.8	0.1	4.5		210.7	100.0	45.1

Note: Table summarizes fuel-economy performance and flexible-fuel production during the 1993–2006 model years. Actual MPG is sales-weighted harmonic-average mileage ignoring flexible-fuel incentives. Fuel economy in excess of the standard is based on sales-weighted standards because the light-truck standard is increasing over time. Table omits several small European automakers with market shares less than 0.1% (e.g., Ferrari) and eight Japanese automakers automakers with market shares ranging from 0.1%–1.6%. (e.g., Hyundai and Subaru). Table does not distinguish between domestic and import passenger-car fleets; all fines for passenger cars were for imports. Chrysler includes Mercedes-Benz for 1999–2006. See text for details.

whose average fuel economy was below the standard during 1993–2006, and all three domestic automakers produced flexible-fuel vehicles. The only domestic fleet above the standard was the General Motors passenger-car fleet, and General Motors did not produce any flexible-fuel cars. In general, the domestic automakers would have paid fines based on their actual fuel economy but did not, thanks to the incentive for flexible-fuel vehicles.[13] This evidence suggests that automakers only produce flexible-fuel vehicles to comply with fuel-economy standards, which is consistent with statements by automakers that flexible-fuel production would fall dramatically if the incentive were eliminated (U.S. Department of Transportation et al. 2002).

Full-line Japanese automakers, such as Honda and Toyota, exceed fuel-economy stan-

[13] The only domestic automaker to pay fines was Chrysler, which produced flexible-fuel vehicles but still paid fines on its import passenger-car fleet. Chrysler's import passenger-car fleet is dominated by Mercedes-Benz, which consistently paid fines prior to merging with Chrysler in 1999.

dards and never produce flexible-fuel vehicles. In fact, the only Japanese automaker that produces flexible-fuel vehicles is Nissan, which did not produce flexible-fuel vehicles until 2005–2006 when its actual fuel economy fell below the light-truck standard for the first time, as we show below. This evidence suggests that adding flexible-fuel capacity does not increase profitability in the absence of a binding fuel-economy constraint and the flexible-fuel provision. If the value of adding flexible-fuel capacity exceeded its incremental cost in the present market equilibrium, we would expect Toyota and Honda to offer models with flexible-fuel capacity, assuming the gain in variable profits exceeded flexible-fuel fixed costs. Below we present evidence that marginal consumers are not willing to pay more for vehicles with flexible-fuel capacity.

While European automakers consistently fall short of fuel-economy standards and regularly pay fines, they do not produce flexible-fuel vehicles. This is not surprising. European sales volumes are low, especially for light trucks. For example, Volkswagen has less than 2% market share in the United States, and its trucks account for just 3% of its U.S. imports, while BMW has about 1% market share and 13.4% of its vehicles are trucks. As we note above, fixed engineering costs act as a barrier to installing flexible-fuel capacity when an automaker has low sales volumes. Moreover, while the European automakers could reduce fines by producing flexible-fuel vehicles, in many cases they could not avoid paying fines entirely. Porsche, for example, fell short of the passenger-car standard by 3.7 miles per gallon and could only improve its fuel economy by 1.2 miles per gallon using flexible-fuel incentives. Thus, marginal compliance costs for these automakers would still equal fines for non-compliance.

Figures 3.1–3.4 provide more detail by plotting AMFA fuel economy calculated using flexible-fuel incentives, actual fuel economy, and fuel-economy standards over time for automakers that produce flexible-fuel vehicles.[14] The figures make clear that domestic automakers regularly depend on flexible-fuel vehicles to comply with fuel-economy standards. For example, Chrysler would have fallen short of the light-truck standard every year from 1994–2002 were it not for the flexible-fuel loophole, while Ford would have missed the light-truck standard every year from 1997–2006. Because automakers can bank or borrow for up to three years, flexible-fuel vehicles that increase fuel economy in a year in which an automaker is already above the standard may still be valuable. For example, Chrysler's flexible-fuel cars in 2003–2005 made up for deficiencies in 2000–2002 and 2006.

Figures 3.1–3.4 also plot the difference between AMFA fuel economy and actual fuel economy in each year, as well as the limit of $\phi = 1.2$ miles per gallon. NHTSA ignores any gain in fuel economy above this threshold when calculating an automaker's compliance in a given year, and an automaker is not able to bank or borrow anything above this limit. Automakers therefore have no incentive to produce above the limit unless marginal consumers value flexible-fuel capacity. As expected, automakers rarely exceed this limit. Chrysler came close with its light-truck fleet in 2002 but did not exceed the limit. Ford and General Motors briefly exceeded the limit for their light-truck fleets in 2003–2004, but reduced flexible-fuel shares in 2005. Note that the gain in mileage from using the flexible-fuel loophole is roughly proportional to the fraction of vehicles with flexible-fuel capacity, assuming the gain in mileage is relatively small. This implies, for example, that Chrysler, which gained about 0.5 miles per gallon using the flexible-fuel provision in 2004, could have roughly doubled its production of flexible-fuel vehicles in 2004 without exceeding the limit.

In general, the figures show that fuel-economy standards were binding for domestic

[14] Our calculations for AMFA fuel economy include a small number of natural gas vehicles and other alternative-fuel vehicles.

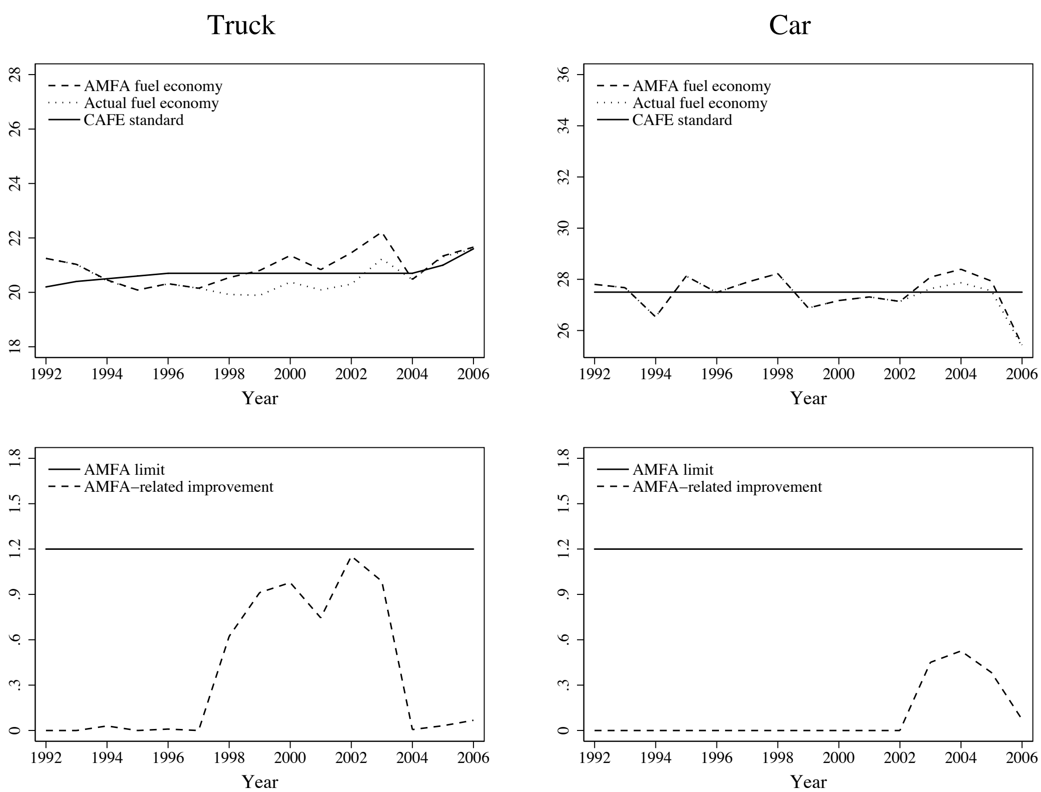

Figure 3.1: Chrysler fuel economy

Note: Top two figures show AMFA fuel economy and actual fuel economy for Chrysler's light-truck and passenger-car fleets for model years 1992–2006. Figures also show fuel-economy standards. AMFA incentives began in 1993. Bottom two figures show annual increase in fuel economy attributable to AMFA incentives and the 1.2 mile-per-gallon limit. Regulations ignore any gain above this limit when calculating an automaker's annual fuel economy. Chrysler merged with Mercedes-Benz in 1998 and began averaging fuel economy with Mercedes in the 1999 model year. Fuel economy for 1992–1998 does not include Mercedes.

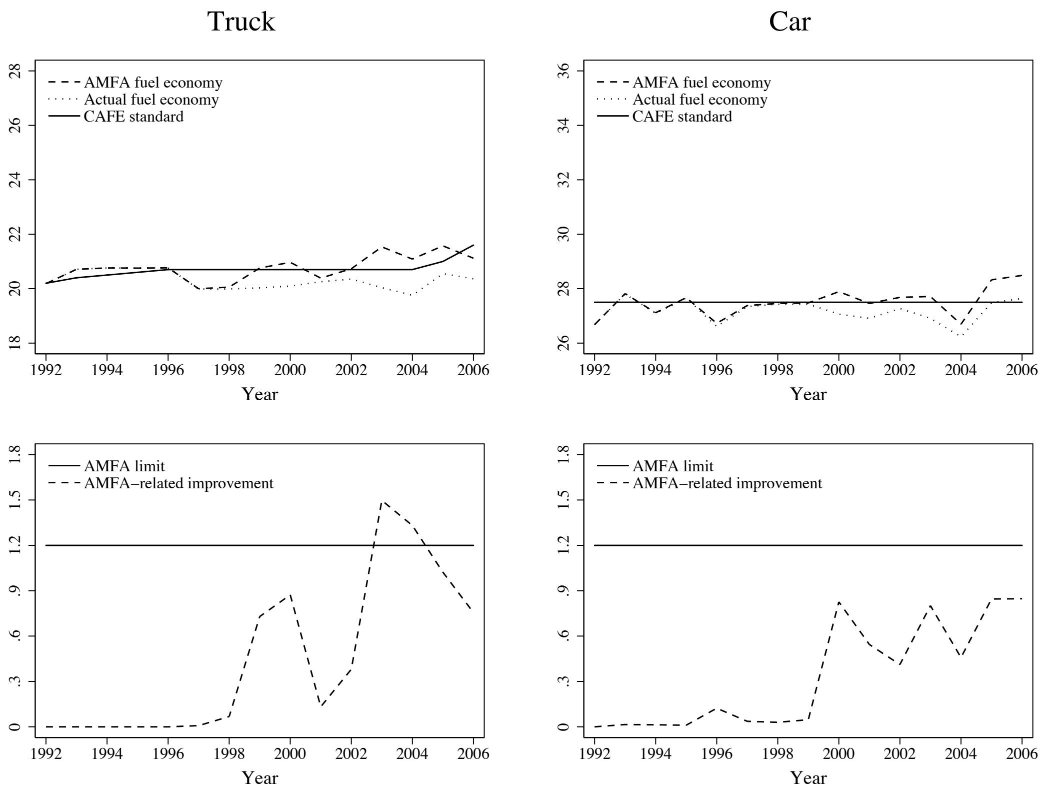

Figure 3.2: Ford fuel economy

Note: Top two figures show AMFA fuel economy and actual fuel economy for Ford's light-truck and passenger-car fleets for model years 1992–2006. Figures also show fuel-economy standards. AMFA incentives began in 1993. Bottom two figures show annual increase in fuel economy attributable to AMFA incentives and the 1.2 mile-per-gallon limit. Regulations ignore any gain above this limit when calculating an automaker's annual fuel economy.

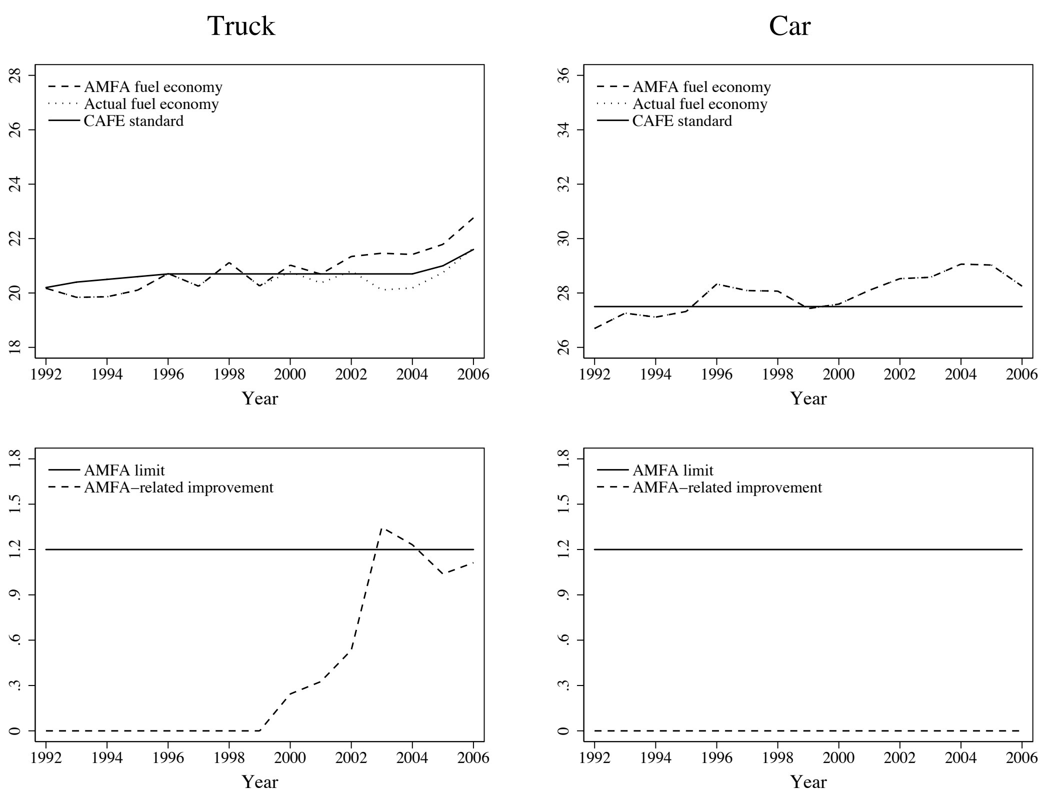

Figure 3.3: General Motors fuel economy

Note: Top two figures show AMFA fuel economy and actual fuel economy for General Motors light-truck and passenger-car fleets for model years 1992–2006. Figures also show fuel-economy standards. AMFA incentives began in 1993. Bottom two figures show annual increase in fuel economy attributable to AMFA incentives and the 1.2 mile-per-gallon limit. Regulations ignore any gain above this limit when calculating an automaker's annual fuel economy.

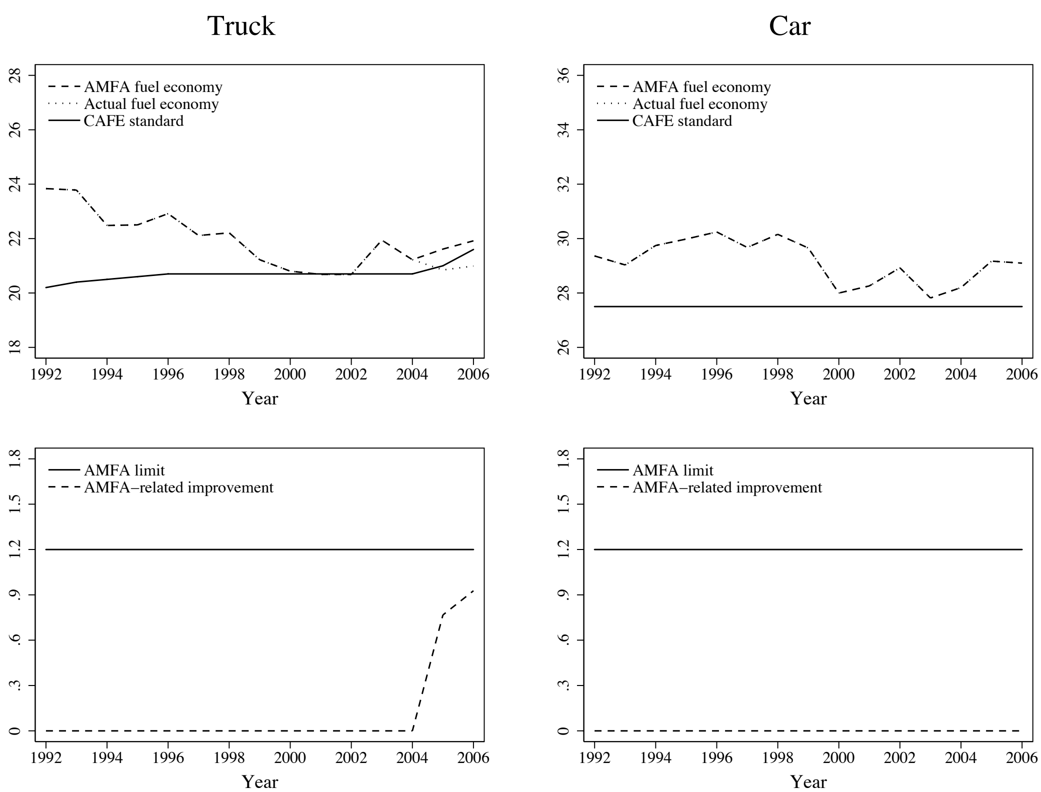

Figure 3.4: Nissan fuel economy

Note: Top two figures show AMFA fuel economy and actual fuel economy for Nissan's light-truck and passenger-car fleets for model years 1992–2006. Figures also show fuel-economy standards. AMFA incentives began in 1993. Bottom two figures show annual increase in fuel economy attributable to AMFA incentives and the 1.2 mile-per-gallon limit. Regulations ignore any gain above this limit when calculating an automaker's annual fuel economy.

automakers during 1993–2006 and that automakers would have paid fines were it not for flexible-fuel vehicles. The figures also show that automakers rarely exhaust the flexible-fuel loophole. These are two of the four conditions we need to infer marginal compliance costs by analyzing the cost of exploiting the flexible-fuel loophole.

3.3.2 Automakers are at interior flexible-fuel shares and rarely install flexible-fuel capacity on more than one model

Our model makes several broad predictions about flexible-fuel shares when automakers exploit the flexible-fuel loophole to comply with CAFE standards. First, fixed engineering costs imply that automakers will not install flexible-fuel capacity on multiple models when any single model is sufficient to take full advantage of the flexible-fuel incentive. Table 3.2 shows that in any given year automakers rarely install flexible-fuel capacity on more than one engine size per fleet. They never install flexible-fuel capacity on more than two engine sizes. Engine size serves as a proxy for engine type, which is the relevant "model" when thinking about fixed engineering costs for flexible-fuel production, as we discuss above. Once an automaker has engineered an engine to be flexible-fuel capable, it can apply flexible-fuel capacity to any model that shares the same engine at roughly the same incremental cost, achieving roughly the same increase in average fuel economy. In some cases there is a one-to-one relationship between engine size and model name, in which case the distinction is irrelevant, but some models with different names share the same engines (e.g., the Ford Explorer and Explorer Sport Trac), and some models are effectively the same vehicle (e.g., the Ford Explorer and Mercury Mountaineer). Note that automakers do not necessarily incur fixed engineering costs in every year, as vehicle characteristics remain largely unchanged between major redesigns.

Our model also predicts that if fixed engineering costs are constant, and if any single model is sufficient to exploit the flexible-fuel loophole, then an automaker will install

Table 3.2: Number of engine sizes with flexible-fuel capacity

	Chrysler Domestic cars	Chrysler Import cars	Chrysler Trucks	Ford Domestic cars	Ford Trucks	General Motors Trucks
1996				1/8		
1997				1/6		
1998			1/9	1/6		
1999			1/12	1/6	1/9	
2000			1/13	1/7	1/10	1/10
2001			1/14	1/6	1/13	1/10
2002			1/15	1/6	2/11	1/12
2003	1/6	1/11	1/15	1/7	2/11	1/10
2004	1/6	2/12	1/11	1/6	1/11	1/12
2005	1/8	2/9	2/13	1/7	1/10	1/11
2006	1/6		2/14	2/7	1/9	1/13

Note: Table shows number of engine sizes that have flexible-fuel capacity for each fleet in each year, as well as the total number of engine sizes. Table omits fleets with no flexible-fuel vehicles.

flexible-fuel capacity on the model with the lowest incremental cost per fuel consumption. If incremental costs increase less than proportionally with fuel consumption, then automakers will tend to install flexible-fuel capacity on their most inefficient models. Finally, automakers will tend to avoid installing flexible-fuel capacity on models with low sales volumes. Figure 3.5 plots flexible-fuel shares and sales-weighted average fuel consumption per mile by engine size for vehicles produced from 1993–2006. Flexible-fuel vehicles are not particularly inefficient relative to other vehicles, although there is no reason to expect a strong relationship, given the potential for wide variation in fixed and incremental production costs across models. Automakers do avoid installing flexible-fuel capacity on models with low sales volumes. Sales volumes in the figure are proportional to circle sizes. Note that the figure does not control for the number of years that various models were offered, however, so sales volumes for some engine sizes may appear artificially low.

Finally, figure 3.5 indicates that flexible-fuel shares were less than one for all engine sizes on which automakers installed flexible-fuel capacity from 1993–2006. This is the third of four conditions we need to infer the marginal cost of tighter CAFE standards. In

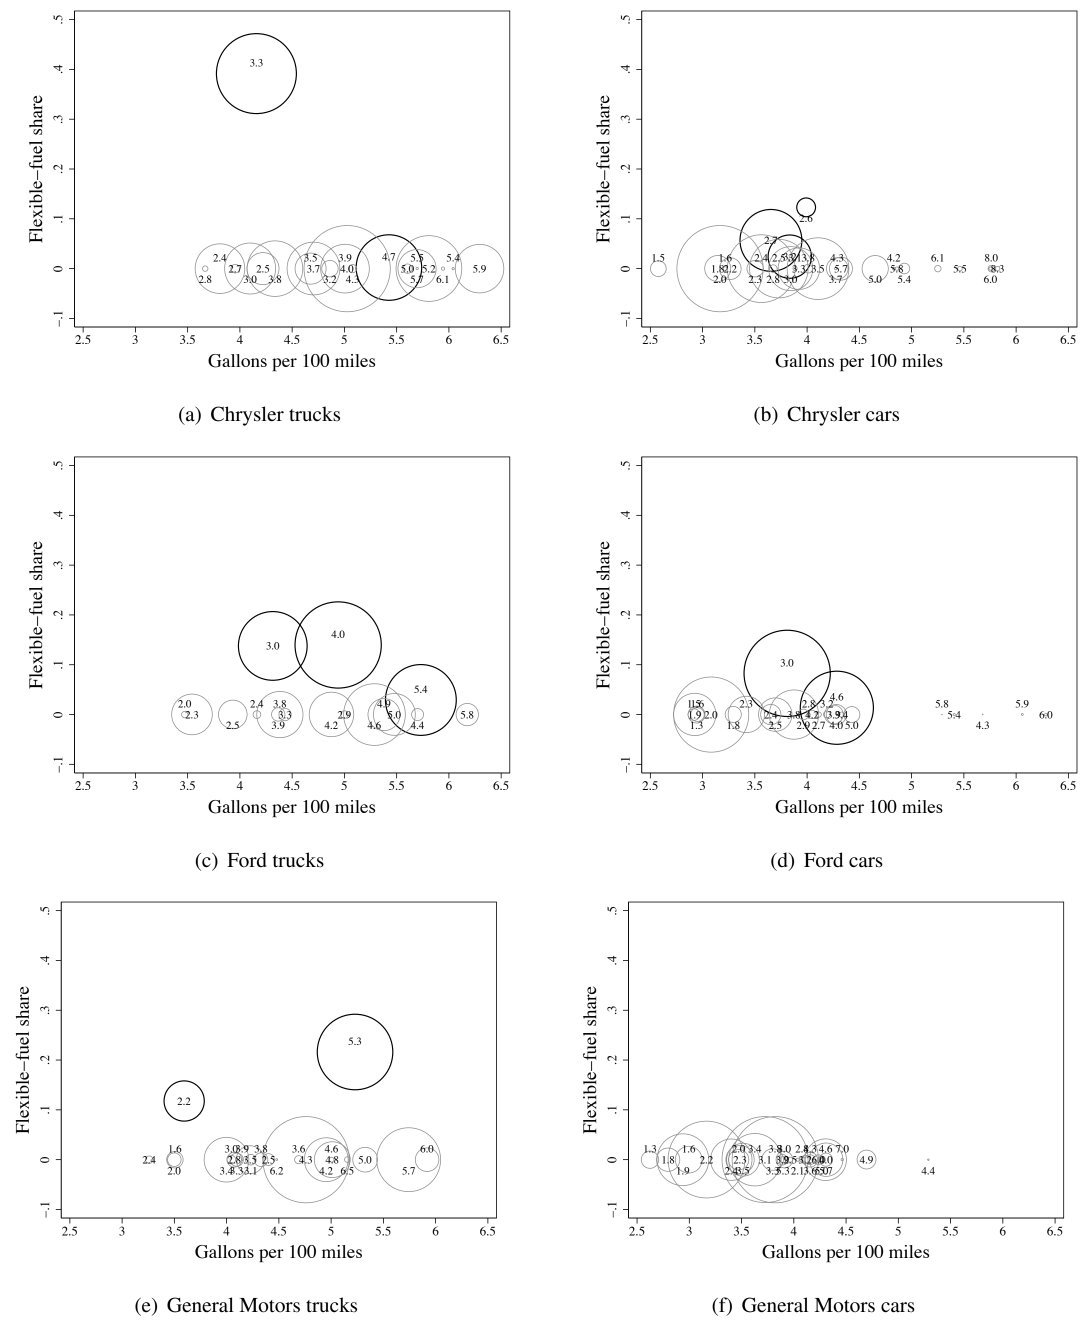

(a) Chrysler trucks

(b) Chrysler cars

(c) Ford trucks

(d) Ford cars

(e) General Motors trucks

(f) General Motors cars

Figure 3.5: Flexible-fuel shares by engine size and fuel consumption

Note: Figure is based on NHTSA fuel-economy compliance data for 1993–2006 model years. Flexible-fuel share is the fraction of units for each engine size that has flexible-fuel capacity. Fuel consumption per 100 miles is the sales-weighted average for each engine size. Circle sizes are proportional to sales. Circle labels are engine sizes in liters. Dark circles indicate engine sizes with nonzero flexible-fuel shares. Specific models are as follows. For Chrsyler: 3.3L truck is the Caravan; 4.7L truck is the Durango and Ram Pickup; 2.7L car is the Stratus and Sebring; 3.2L car is the Mercedes C240; 2.6L car is the Mercedes C230. For Ford: 3.0L truck is the Ranger and Mazda B3000; 4.0L truck is the Explorer and Mountaineer; 5.4L truck is the F150 pickup; 3.0L car is the Taurus and Sable; 4.6L car is the Town Car, Grand Marquis, and Crown Victoria. For General Motors: 2.2L truck is the S10 and Sonoma; 5.3L truck is the Suburban, Tahoe, Yukon, Avalanche, Sierra, and Silverado. Compliance data for the 2007 model year and beyond are not yet available.

summary, automakers respond as predicted to flexible-fuel incentives, and the first three conditions we need to infer marginal compliance costs hold. It only remains to show that marginal consumers do not value flexible-fuel capacity.

3.4 Marginal consumers do not value flexible-fuel capacity

We provide empirical evidence based on nearly one-million new vehicle transactions that marginal consumers do not value flexible-fuel capacity. This is important because our key result—that constrained automakers equate the marginal cost of improving fuel economy using flexible-fuel vehicles with the marginal cost of improving fuel economy through other means—depends on the marginal consumer having zero valuation. When combined with the evidence we present above that automakers use flexible-fuel vehicles to comply with CAFE standards, install flexible-fuel capacity on some but not all units for their flexible-fuel models, and do not exhaust the flexible-fuel loophole, the zero valuation implies that we can pin down marginal compliance costs exactly.

Rosen (1974) shows that what matters for equilibrium prices in a hedonic framework is the valuation of marginal agents. We suspect that some consumers would pay more for a vehicle with flexible-fuel capacity, but that the flexible-fuel loophole leads automakers to supply flexible-fuel vehicles in sufficiently large quantities that marginal consumers are indifferent.[15] In fact, automakers sell many flexible-fuel vehicles to consumers who have no access to ethanol, and we estimate that the price premium for flexible-fuel vehicles is approximately zero. These results imply that marginal consumers indeed do not value flexible-fuel capacity.

These findings are consistent with evidence that many consumers are unaware that they own flexible-fuel vehicles, particularly in earlier years. For example, a report by several

[15] If automakers are able to price discriminate, they may be able to generate revenue from flexible-fuel capacity, even if marginal consumers are indifferent.

federal government agencies in 2002 concluded that "many people who have purchased flexible-fuel vehicles do not know they could use E85" (U.S. Department of Transportation et al. 2002), and a major ethanol-producing firm found that about 70% of flexible-fuel vehicle owners surveyed in 2005 did not know they owned flexible-fuel vehicles (Wald 2005).

3.4.1 New vehicle transaction data

Our vehicle transaction data come from an industry source that collects data directly from a nationally representative sample of dealers. The data contain detailed information on new vehicle prices and characteristics for millions of transactions from 2000–2007. The data record sales prices, manufacturer rebates, trade-in prices, and trade-in market values. This information allows us to adjust prices for manufacturer rebates and any difference between trade-in prices and actual trade-in market values. Automakers frequently offer financing incentives instead of manufacturer rebates. We observe interest rates and other information for dealer-financed transactions, allowing us to control for financing incentives.[16] Finally, the data record the calendar date of each transaction and the state in which the transaction took place, as well as the buyer's age and gender. Some observations also include manufacturer-suggested retail prices. We deflate all prices by the consumer price index for all urban consumers and all items from the U.S. Bureau of Labor Statistics.

To isolate the value of flexible-fuel capacity, we identify flexible-fuel vehicles and comparison vehicles in our transaction data that we observe to be identical along every observable dimension except fuel type. The transaction data include each vehicle's truncated vehicle identification number (VIN), which provides information about a vehicle's make,

[16] We calculate the value of financing incentives in dealer-financed transactions by comparing a car buyer's actual stream of monthly payments to the payment stream she would have faced at a market interest rate. We calculate the actual stream of monthly payments using the loan's size, term, and dealer APR. We calculate an alternative stream of payments using the market-average APR for new car loans through commercial banks from the Federal Reserve Board. The Fed reports average interest rates every three months. We calculate interest rates for intervening months using linear interpolation. Finally, we calculate the present value of each payment stream using a 4% annual rate of pure time preference. The value of the financing incentive is the difference between these two present values. These calculations are identical to Corrado et al. (2006).

model, model year, body style, number of doors, drive type, transmission, engine displacement, number of cylinders, and aspiration (e.g., turbo-charged). The data also record each vehicle's fuel type, distinguishing between gasoline-only vehicles, flexible-fuel vehicles, diesels, gasoline-electric hybrids, and other fuels. We cross-reference these fuel types with information from the National Ethanol Vehicle Coalition, which gives model names, model years, engine sizes, and flexible-fuel VIN identifiers (usually the 8th digit of the VIN) for ethanol-gasoline flexible-fuel vehicles. We omit flexible-fuel models that do not also appear in the Coalition's list, as some models in our data are actually natural gas dual-fuel vehicles.

For the flexible-fuel models that remain we identify comparison gasoline-only vehicles in our transaction data based on model name, model year, engine size, and other observable characteristics, including truncated VINs that exclude flexible-fuel VIN identifiers. We group vehicles into vehicle types based on these observable characteristics. Restricting the sample to these vehicles gives a preliminary sample size of nearly 900,000 transactions. About one-quarter of these transactions are for vehicles where we observe a single VIN or more than two VINs per vehicle type. We omit these observations to minimize the possibility of unobservable characteristics correlating with flexible-fuel capacity. This gives a final estimation sample of nearly 666,000 observations.

Table 3.3 presents summary statistics for our final sample, while Table 3.4 presents model names and quantities for flexible-fuel models and comparison vehicles. At the end of the day, the detailed transaction data allow us to identify and compare, for example, the price of a gasoline-only 2006 Ford F150 extended-cab pickup with a 5.4L V8 engine and manual transmission to the price of a flexible-fuel 2006 Ford F150 extended-cab pickup with a 5.4L V8 engine and manual transmission. The data do not, however, include information about various options that may be installed. We investigate whether the value of

Table 3.3: Summary statistics

Variable	Mean	Std. Dev.	Min.	Max.	Obs.
flexible-fuel vehicle	0.62	0.48	0	1	665887
transaction price	31229.62	8352.92	-10994.38	75002.08	665887
suggested retail price	37139.43	8694.88	0	206606.53	249706
manufacturer rebate	2415.44	2021.8	0	11409.17	665887
inventory days	73.08	85.63	1	805	644291
loan at dealer	0.75	0.43	0	1	576274
interest rate (% APR)	5.51	4.22	0	29.99	427656
down payment	7347.18	7399.25	-10994.01	55054.91	432137
monthly payment	590.1	195.7	13.39	3836.9	427656
loan term (months)	62.44	11.14	12	96	427656
trade-in vehicle	0.51	0.5	0	1	665887
trade-in balance	1347.07	2730.48	-23974.05	30170.07	341293
age of buyer	44.94	13.43	16	107	580964
female buyer	0.27	0.45	0	1	596987
ethanol availability (%)	0.17	0.68	0	8.10	665887

Note: Table shows summary statistics for final estimation sample based on flexible-fuel vehicles and their gasoline-only counterparts. See text for details.

such options is correlated with flexible-fuel capacity below.

In addition to these transaction data, we collect information on ethanol refueling locations from the Department of Energy's Alternative Fuels Data Center. These data record station addresses. The data do not systematically record open dates, but they do record the date when each station was added to the database. We assume that these add dates approximate open dates and calculate the total number of ethanol stations in each state in each month. The Department of Energy began collecting these data in 1995, and new stations are added regularly, so our calculations based on add dates give a fairly accurate picture of how ethanol availability evolved during our sample period from 2000–2007. We calculate percent ethanol availability by dividing by the total number of retail gasoline stations in each state using information from National Petroleum News.[17]

Table 3.4: Flexible-fuel models in the data sample

Model	**Gasoline-only**	**Flexible-fuel**	**Total**
Armada	1,165	1,506	2,671
Aspen	119	565	684
Avalanche	1,473	16,208	17,681
B3000	1,350	773	2,123
Caravan	5,959	13,612	19,571
Cherokee	393	1,311	1,704
Commander	111	667	778
Crown Victoria	85	236	321
Dakota	39	46	85
Durango	2,677	212	2,889
Explorer	41,677	70,855	112,532
Express	65	58	123
F150	26,579	26,322	52,901
Grand Marquis	1,471	4,583	6,054
Impala	796	19,731	20,527
Monte Carlo	17	1,376	1,393
Mountaineer	6,026	4,647	10,673
Ranger	4,355	1,743	6,098
S10	3,245	6,986	10,231
Sable	1,237	37	1,274
Savana	62	12	74
Sebring	4,693	2,180	6,873
Sierra	7,028	4,120	11,148
Silverado	16,801	9,217	26,018
Sonoma	735	1,680	2,415
Stratus	1,872	11	1,883
Suburban	17,112	61,764	78,876
Tahoe	45,605	75,919	121,524
Taurus	8,626	8,574	17,200
Terraza	149	14	163
Titan	19,926	20,342	40,268
Town Car	1,009	2,585	3,594
Town & Country	2,416	10,126	12,542
Uplander	125	46	171
Voyager	1,510	6,565	8,075
Yukon	24,055	40,695	64,750
Total	250,563	415,324	665,887

Note: Table shows flexible-fuel models and quantities in estimation sample. Sample excludes flexible-fuel models with a single VIN or more than two VINs per vehicle type. See text for details.

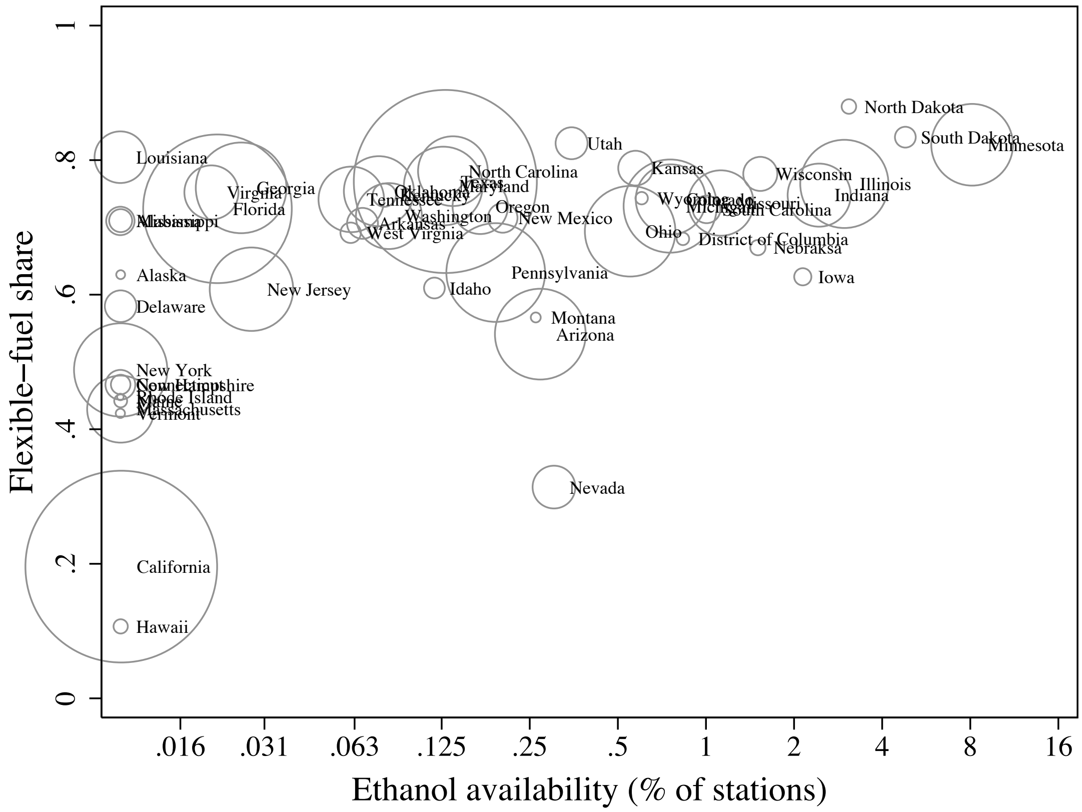

Figure 3.6: Flexible-fuel shares and ethanol availability

Note: Flexible-fuel share is the fraction of vehicles in the estimation sample that have flexible-fuel capacity. Ethanol availability is the maximum fraction of stations that offer ethanol at any time during 2000–2007. Sizes of circles are proportional to the number of observations. Figure sets availability to 0.01% for 13 states with zero ethanol stations to be compatible with log scaling. These states appear along the left-hand side of the figure. California's peak availability is small but not zero.

3.4.2 Many flexible-fuel vehicle buyers do not have access to ethanol

Our first step is to analyze the relationship between the availability of retail ethanol in a consumer's state of residence and the geographic allocation of flexible-fuel vehicles. Our reasoning is that if a large number of vehicles are sold in states that lack ethanol, it is highly unlikely that marginal consumers value flexible-fuel capacity. Our analysis indicates that while there is a positive correlation between ethanol availability and flexible-fuel sales across states, this relationship is weak.

[17] Although National Petroleum News reports data annually, we divide by the mean number of retail gasoline stations in each state from 2000–2006, because the data collection process appears to vary from year to year.

Figure 3.6 plots flexible-fuel shares and peak ethanol availability by state. We calculate flexible-fuel shares based on our estimation sample of flexible-fuel vehicles and comparison gasoline-only vehicles. Flexible-fuel shares for these vehicles range from 0.6–0.8 in most states. Flexible-fuel shares are substantially lower in California, where many flexible-fuel vehicles fail the state's strict emissions laws, and in Hawaii and Nevada. For the remaining states there appears to be a slight positive correlation between flexible-fuel share and ethanol availability, but the correlation is weak. Doubling ethanol's availability ten times over only correlates with a 30% increase in flexible-fuel shares, and flexible-fuel shares are high all over the country.

A full 17% of the flexible-fuel vehicles in our estimation sample sell in states where there are no ethanol pumps or just a single pump at the end of the sample period, while more than 86% sell in states where ethanol is available at less than 1% of stations. It is difficult to imagine that more than a handful of consumers in these states are willing to pay for flexible-fuel capacity. Thus, automakers deciding on how many flexible-fuel vehicles to produce must have expected that the price premium for marginal vehicles would be zero.

We also test the relationship between flexible-fuel quantities and ethanol pumps statistically. We calculate the flexible-fuel share for each vehicle type in each state in each year. We also calculate the peak fraction of fueling stations that have an ethanol pump in each state in each year.[18] We then regress flexible-fuel shares on availability, controlling for vehicle-specific fixed effects. We do not include year controls because vehicle types already differentiate by model year.

Table 3.5 presents the estimation results. The coefficient on ethanol availability in regression (1) implies that increasing ethanol's market penetration by 1% correlates with

[18]Using the mean fraction of fueling stations with an ethanol pump in each year does not alter the results appreciably.

Table 3.5: Where are flexible-fuel vehicles allocated?

Controls	(1) State dummies excluded	(2) State dummies included
percent ethanol availability	0.047	0.048
	(0.014)	(0.011)
observations	26105	26105
groups	517	517
R-squared (within)	0.02	0.16

Note: Dependent variable is flexible-fuel share within each vehicle-state-year group. Both regressions include vehicle-specific fixed effects. Standard errors in parentheses are clustered by state.

an increase in flexible-fuel shares of 0.047. This relationship might be biased by unobserved determinants of flexible-fuel shares across states, such as California's strict emissions laws. Regression (2), however, which includes state dummy variables, finds a nearly identical correlation. Flexible-fuel shares correlate with differential changes in ethanol availability across states over time, as well as with differences in availability across states in any given year. While these coefficient estimates are consistent with automakers allocating vehicles based in part on preferences, flexible-fuel shares are high everywhere—even in states with virtually no ethanol pumps. If automakers are "overproducing" flexible-fuel vehicles to exploit the flexible-fuel loophole, then marginal consumers in these and other states are unlikely to value flexible-fuel capacity. Our estimates for the flexible-fuel price premium are consistent with this hypothesis.

3.4.3 Consumers do not pay extra for flexible-fuel capacity

Given that automakers sell a large fraction of flexible-fuel vehicles to consumers who lack access to ethanol, one would expect the equilibrium price of flexible-fuel capacity to be zero. In the presence of market power or price discrimination, however, consumers in states with ethanol availability, such as Minnesota, might pay a premium, even if marginal consumers in other states do not. We compare the prices of vehicles with and without flexible-fuel capacity and find that their prices are not statistically different.

Anecdotal evidence from media reports and from a report by several federal agencies suggests that automakers sometimes increased the manufacturer's suggested retail price (MSRP) for flexible-fuel vehicles, but then netted-out these price increases with targeted rebates (U.S. Department of Transportation et al. 2002). In other media reports, automakers claim that the cost of flexible-fuel capacity is not passed on to consumers (Kohn 2000; Williams 2008). We checked the MSRP of several flexible-fuel vehicles in May 2008 and found that list prices were the same as comparable gasoline-only vehicles. We examine whether or not transaction prices and list prices vary with fuel type using our vehicle transaction data.

We estimate the price premium for flexible-fuel vehicles using the following econometric specification:

$$p_{ijst} = \gamma FFV_{ijst} + \delta_{jst} + \varepsilon_{ijst}, \tag{3.12}$$

where: p_{ijst} is the sales price that we observe in transaction i for vehicle type j in state s and in month t; FFV_{ijst} is a dummy variable that equals one if the vehicle in the transaction is a flexible-fuel vehicle and zero otherwise; δ_{jst} is a vehicle-state-month fixed effect; and ε_{ijst} is an error term. We estimate the model using least-squares estimation and vehicle-state-month fixed effects.

The coefficient of interest is γ. This coefficient is the average price premium for flexible-fuel vehicles relative to comparable gasoline-only vehicles sold in the same place at the same time. This coefficient measures the marginal willingness to pay for flexible-fuel capacity. We identify this parameter based on the difference in price between flexible-fuel vehicles and gasoline-only vehicles that are identical based on other observable characteristics.[19] Estimating implicit prices for a vehicle's fuel type or mileage can be chal-

[19]These characteristics include make, model, model year, trim level, engine displacement, cylinders, body style, number of doors,

lenging, given the strong collinearity these characteristics usually share with other attributes (see Espey and Nair (2005)). This is not a problem here. Because we observe thousands of transactions for flexible-fuel models and their identical gasoline-only counterparts, we are able to control for vehicle attributes non-parametrically using vehicle-specific fixed effects and still estimate flexible-fuel premiums precisely. Most previous studies use annual cross-sectional data for vehicle list prices and control for vehicle attributes parametrically.

Our transaction-level microdata also allow us to control flexibly for a vast number of potentially confounding variables. The vehicle-state-month fixed effects given by δ_{jst} are equivalent to including vehicle, state, and month dummy variables, as well as all relevant two-way and three-way interactions of these variables. These controls eliminate nearly all sources of confounding variation one could imagine.

The error term ε_{ijst} reflects unobserved vehicle characteristics such as carpet floor mats, tinted windows, or other options that do not come standard in observed trim levels. The error term also reflects transaction-level variation in final sales price, deriving for example from differences in negotiating skill across dealers and buyers. The identification assumption is that this error term is uncorrelated with flexible-fuel capacity, conditional on state, month, and vehicle type: $E[\varepsilon_{ijst} \cdot FFV_j | \delta_{jst}] = 0$. If unobserved vehicle characteristics or other determinants of prices are correlated with flexible-fuel capacity, then least-squares estimates of the flexible-fuel premium γ will be biased.

Table 3.6 presents the estimation results for the model in equation (3.12). The coefficient in regression (1) indicates that the marginal consumer demands a \$22 price discount to purchase a flexible-fuel vehicle during the sample period, although this coefficient is not statistically different from zero.[20] When we restrict the analysis to cash transactions

drive type, transmission type, aspiration (e.g., turbo-charged), and truncated VINs excluding flexible-fuel identifiers.

[20]These results are consistent with earlier work by Liu (2007), who estimates flexible-fuel premiums using annual nationwide data for suggested retail prices from 1996–2001. She estimates a premium of \$0.37.

Table 3.6: Flexible-fuel premium

Controls	(1) Primary regression	(2) Cash sales only	(3) MSRP prices
FFV	-22.07	-38.13	154.21
	(28.29)	(60.19)	(42.85)
observations	665887	143869	249706
groups	99398	52264	55249
R-squared (within)	0.00	0.00	0.00

Note: Dependent variable in regression (1)–(2) is sales price net of manufacturer rebates, financing incentives, and trade-in overallowance. Regression (2) estimates the model using transactions where the purchaser paid cash at the dealer (i.e., did not borrow or lease from the dealer), so financing incentives do not apply. Dependent variable in regression (3) is is the manufacturer's suggested retail price (MSRP). All regressions control for vehicle-state-month fixed effects. Standard errors in parentheses are clustered by vehicle-state-month cells. See text for further details.

in regression (2) the flexible-fuel premium falls slightly to $-\$38$ but is statistically indistinguishable from the estimate in regression (1). These results suggest that neither dealer-financed sales nor our adjustment for financing incentives change the estimates appreciably.

Some sources indicate that list prices include a flexible-fuel premium, while other sources indicate they do not. Table 3.6 shows that MSRPs are about \$150 dollars higher for flexible-fuel vehicles in our sample. Assuming that automakers do not include a flexible-fuel premium in list prices, then this coefficient would imply that flexible-fuel capacity is correlated with unobserved options packages that consumers value. The flexible-fuel premium of $-\$22$ we estimate in regression (1) would be biased upward. That is, consumers would require an even larger discount to purchase flexible-fuel vehicles than what we estimate. Because incremental production costs reportedly range from \$100–\$200, however, the \$154 premium is also consistent with anecdotal evidence that some automakers raised MSRPs for flexible-fuel vehicles but rebated the difference. This would explain why MSRPs are higher for flexible-fuel vehicles (regression 3), while transaction prices

Table 3.7: Are flexible-fuel transactions different?

	(1) Days on lot	(2) Dealer loan?	(3) Interest rate	(4) Total down	(5) Monthly amount	(6) Loan term	(7) Trade auto?	(8) Trade balance	(9) Age of buyer	(10) Female buyer?
FFV	-29.43	-0.011	-0.03	46.81	0.03	-0.17	-0.0002	-19.13	0.15	-0.001
	(1.07)	(0.003)	(0.03)	(50.34)	(1.08)	(0.08)	(0.0030)	(21.07)	(0.08)	(0.003)
obs.	644291	576274	427656	432137	427656	427656	665887	341293	580964	596987
grps.	97097	95135	79231	79731	79231	79231	99398	75134	90386	94463
R-sq.	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Note: Dependent variables are: (1) days that vehicle was in dealer's inventory prior to sale (2) indicator variable that equals one if buyer took out loan from dealer and zero if buyer purchased vehicle with cash; (3) APR interest rate conditional on loan from dealer; (4) down payment conditional on loan from dealer; (5) monthly payment conditional on loan from dealer; (6) loan term in months conditional on loan from dealer (7) indicator that equals one if buyer sold dealer a trade-in vehicle at time of purchase and zero otherwise; (8) trade-in amount minus trade-in market value conditional on trade-in vehicle; (9) age of first buyer listed on purchase agreement; (10) indictor variable that equals one if first buyer listed is female and zero otherwise. All regressions include vehicle-state-month fixed effects. Standard errors in parentheses are clustered by vehicle-state-month cells. See text for details.

are not (regressions 1–2). We are hesitant to read too deeply into this difference, however, as it is likely the product of sample selection. When we limit our analysis to the MSRP sample, transaction prices are $121 higher for flexible-fuel vehicles.

If consumers had specific preferences for flexible-fuel vehicles, we would expect these preferences to correlate with consumer characteristics, such as age or income. This would lead to sorting on flexible-fuel capacity. Similarly, if automakers installed flexible-fuel capacity on models with low-value or high-value options packages, these packages would correlate with consumer characteristics, which would also lead to sorting. Our data allow us to test this hypothesis. Using the same econometric specification as in equation (3.12) above, we estimate the correlation between flexible-fuel capacity and other transaction characteristics. Sallee (2007) uses a similar approach to test whether the typical Prius buyer who purchased her vehicle when tax incentives were available is different from the typical buyer who purchased his vehicle when incentives were not available.

Table 3.7 presents the results of these regressions. The first regression indicates that flexible-fuel vehicles sold 29 days earlier than comparable gasoline-only vehicles sold at the same time in the same state. This is a fairly sizable difference given that vehicles in our sample remain in a dealer's inventory an average of 73 days. There are several plausible explanations. First, flexible-fuel capacity may be correlated with other desirable vehicle options that we do not observe. This seems unlikely, given that we find a zero price premium for flexible-fuel vehicles on average. Second, flexible-fuel vehicles may go disproportionately to dealers with high turnover. Finally, flexible-fuel vehicles may spend fewer days in inventory if some consumers specifically request flexible-fuel vehicles. As long as these consumers are not marginal they could increase turnover without affecting prices.

None of the other transaction characteristics differ between flexible-fuel and gasoline-only vehicles. Flexible-fuel buyers are no more or less likely to finance their vehicles through dealers. Interest rates are no different for flexible-fuel buyers, nor are down payments, monthly payments, or loan durations. Flexible-fuel and gasoline-only buyers trade in used vehicles just as often, and trade-in balances do not differ systematically. Finally, flexible-fuel and gasoline-only buyers are the same age and gender on average. In summary, we detect no observable differences between car buyers that purchase flexible-fuel vehicles and those that buy identical gasoline-only vehicles. Flexible-fuel vehicles spend fewer days in inventory, however, which is consistent with some buyers having specific preferences for flexible-fuel capacity. These buyers do not affect prices.

Overall, our analysis of prices and quantities suggests that automakers do not charge more for flexible-fuel vehicles, and, more specifically, that the marginal consumer does not value flexible-fuel capacity. This justifies the formulation of our model, which implicitly assumes that consumers ignore flexible-fuel capacity. When combined with the evidence we present above that automakers exploit the flexible-fuel loophole to comply with CAFE standards, install flexible-fuel capacity on some but not all units, and do not exhaust the flexible-fuel loophole, we have shown that the four conditions we need to pin down marginal compliance costs actually hold.

3.5 Estimating marginal compliance costs

Using our methodology, we now calculate marginal compliance costs for automakers that produced flexible-fuel vehicles. Equation (3.7) from above, which we repeat here for convenience, shows that the cost per vehicle of marginally increasing the CAFE standard is a function of both flexible-fuel vehicle attributes and regulatory parameters:

$$\frac{\partial \mathcal{L}^*}{\partial \sigma}\frac{1}{Q} = -\frac{\alpha \cdot m}{\sigma^2(1-\beta)}, \tag{3.13}$$

where α is the incremental cost of adding flexible-fuel capacity, m is actual fuel economy, σ is the nominal fuel-economy standard, and β is the AMFA incentive for flexible-fuel vehicles. Again, this result assumes that the backstop constraint on actual mileage is not binding. In theory, this equality holds separately for all models whose flexible-fuel shares are greater than zero and less than one. As we show above, however, an automaker will typically only have one model with an interior flexible-fuel share, both in theory and in practice.

We calculate marginal compliance costs by plugging in parameter values as follows. While we do not observe the incremental cost of adding flexible-fuel capacity to particular models, reports from various sources put these costs at anywhere from \$25–\$300 per vehicle.[21] We use a range of \$100–\$200 per vehicle, which we think gives a conservatively high estimate of incremental costs. We calculate relevant mileage as the sales-weighted harmonic-average mileage of an automaker's flexible-fuel vehicles. We assume,

[21] Reliable sources put costs as high as \$150–\$300 per vehicle before automakers began producing flexible-fuel vehicles in large quantities (U.S. Environmental Protection Agency 1990) to as low as \$25–\$50 currently (Alson 2008), while NHTSA put the range at \$100–\$200 when it ruled to extend the flexible-fuel provision in 2004 (U.S. Department of Transportation 2004). Recent reports in the popular press quoting automakers themselves are consistent with these ranges, with costs ranging from "\$70 to \$100 per vehicle, depending on engine size" (Williams 2008), to "at most a few hundred dollars more per car" (Barrionuevo and Maynard 2006). Some sources report costs at "high sales volumes," implying that some cost estimates include average fixed costs. Rubin and Leiby (2000) cite a consulting report from 1995 that estimated fixed costs of \$4.2 million per model annually and incremental production costs of \$240 per vehicle.

Table 3.8: Marginal compliance costs

Automaker	Vehicle fleet Trucks	Cars
Chrysler	\$14–\$28	\$9–\$18
Ford	\$12–\$24	\$8–\$17
General Motors	\$11–\$22	\$9–\$18
Nissan	\$11–\$21	
CAFE fine	\$55	\$55
Jacobson (2007)	\$11–\$23	\$4–\$36

Note: Table shows estimates of marginal compliance costs based on equation (3.13). Ranges assume an incremental cost of \$100–\$200 for adding flexible-fuel capacity. Calculations use the sales-weighted mean fuel-economy standard for each automaker and the sales-weighted mileage for each automaker's flexible-fuel vehicles during 1993–2006. Estimate for the General Motors passenger-car fleet assumes that GM applied flexible-fuel capacity to a car of average mileage when it began producing flexible-fuel vehicles in 2007. Ford and General Motors exhausted the flexible-fuel loophole for their light-truck fleets in 2003–2004, so marginal compliance costs in at least those years are higher than what we estimate here. Chrylser pays fines on its import passenger-car fleet; costs based on our methodology are for domestic fleet only. Firms that serially pay fines have marginal cost equal to CAFE fine of \$55. Table also includes Jacobson (2007) estimates for comparison.

as above, that the flexible-fuel incentive is $\beta = 0.6$. Finally, while the standard for passenger cars remains 27.5 miles per gallon during the entire study period, the light-truck standard increases gradually from 20.4–21.6 miles per gallon. We therefore calculate the sales-weighted harmonic-average standard for each automaker. We calculate costs separately for light-truck and passenger-car fleets.

Table 3.8 presents our estimates of marginal compliance costs for automakers that produce flexible-fuel vehicles. Tightening the light-truck standard by one mile per gallon would cost these automakers at most \$11–\$28 in lost profit per vehicle, while tightening the standard for passenger cars would cost at most \$8–\$18 per vehicle. The ranges for each automaker derive from the assumed range of \$100–\$200 for incremental production costs. Costs are not identical because the mileage of flexible-fuel vehicles varies from automaker to automaker, as does the average light-truck standard. We do not calculate compliance costs on a year-by-year basis, because banking and borrowing provisions allow automakers to equate marginal compliance costs over time.

Again, because the automaker is optimizing on the margin, these costs equal the marginal cost of improving AMFA fuel economy using flexible-fuel vehicles, as well as the marginal cost of improving AMFA fuel economy through other means. Recall that while AMFA fuel economy is not equivalent to actual fuel economy, the two are quite close in practice. Our estimates therefore reflect lower profit margins on smaller, more efficient vehicles, as well as the difference between production costs and willingness to pay for marginal improvements in vehicle efficiency.

Costs are substantially lower than the $55 fine that automakers pay when they are out of compliance, which previous researchers use as a measure of compliance costs. Thus, were it not for fixed engineering costs, European automakers could reduce compliance costs by producing flexible-fuel vehicles.

Nissan first produced flexible-fuel trucks in 2005–2006, revealing marginal compliance costs for at least those years. General Motors first began producing flexible-fuel cars in the 2007 model year. Assuming that the passenger-car standard is binding for General Motors toward the end of the study period, and that General Motors applied flexible-fuel capacity to cars with average mileage, then our methodology implies that GM's marginal cost of compliance is $9–$18 per vehicle. Marginal costs for these automakers were as low as zero in earlier years when their fuel economy exceeded the standard and they did not produce flexible-fuel vehicles.

For constrained automakers that exhaust the flexible-fuel loophole, we are able to estimate lower bounds for marginal compliance costs. Ford and General Motors both exhausted the flexible-fuel loophole for their light-truck fleets in 2003–2004, so their marginal costs are probably higher than the estimates in table 3.8 for at least those years. Their costs are bounded from above by the $55 penalty for non-compliance, ignoring any implicit fixed cost for non-compliance. Marginal compliance costs are zero for unconstrained automakers, such as Honda and Toyota. Marginal compliance costs are $55 per vehicle for automakers that serially pay fines, including Volkswagen and Porsche.

Jacobson (2007) estimates marginal compliance costs for domestic automakers during 1997–2001 using a wholly different methodology based on estimated demand elasticities and implied markups.[22] He finds that tightening the fuel-economy standard for light trucks by one mile per gallon would cost domestic automakers $11–$23 per vehicle, depending on the automaker, while tightening the standard for passenger cars would cost $4–$36. These estimates are very close to what we estimate based on incremental costs for flexible-fuel vehicles. This is precisely what theory would predict, given that most domestic automakers produced flexible-fuel vehicles during 1997–2001. We find this overlap reassuring.

Our cost estimates nevertheless have several limitations. First, like other estimates in this literature, our estimates only reflect the cost of marginal increases in CAFE standards. They do not reflect engineering investments, capital expenditures, and other fixed costs that may be required for aggressive increases in mileage. These costs could be substantial. Second, our estimates reflect compliance costs during our study period and do not necessarily hold for current or future years. Vehicle characteristics, consumer preferences, and technology evolve over time. Moreover, the structure of CAFE regulation is currently in flux, as regulators move toward "size-based" standards, which mandate higher mileage for firms that produce smaller vehicles. These reforms will undoubtedly impact compliance costs. Finally, our estimates do not reflect changes in consumer surplus resulting from tighter fuel-economy standards.

To put our cost estimates in context, we provide back of the envelope calculations for

[22]He estimates a system of demand elasticities for new vehicles, assumes that oligopolistic automakers engage in Nash-Bertrand pricing behavior, and then solves each automaker's system of first-order conditions to impute marginal costs. He then assumes that the share of a model's markup that goes to dealers is constant across an automaker's models. This is his key identification assumption. He is then able to regress observed dealer markups over invoice on imputed costs and fuel consumption for each model. The estimated parameter on fuel consumption yields the shadow value of the fuel-economy constraint.

the marginal external benefits of tighter fuel-economy standards assuming that automakers are forced to comply by improving actual fuel economy. Tighter fuel-economy standards reduce U.S. gasoline consumption, which lowers world oil prices, mitigates adjustment costs associated with oil price shocks, and reduces carbon dioxide emissions. Tighter standards reduce the cost of traveling a mile, however, which leads to increased travel and offsetting externalities, including noise, congestion, and traffic accidents. Net benefits are therefore highly sensitive to the elasticity of miles with respect to mileage. Generous assumptions would put benefits at roughly \$0.30 per gallon, costs at \$0.10 per mile (Harrington et al. 2007), and the elasticity response at 0.1 (Small and Dender 2007). Assuming that the average truck travels 190,000 miles in its lifetime, the external benefit for light trucks is \$23 per vehicle. The external benefit for cars is $-\$1$, assuming a car travels 160,000 miles.[23] We are unable to perform a formal benefit-cost test, as our cost estimates do not include changes in consumer surplus. Jacobson (2007) finds that consumers bear over 80% of the welfare loss of tighter standards, however, which suggests that fuel-economy standards are unlikely to pass a benefit-cost criterion, even though the cost to producers is small.

[23]We obtain information on average lifetime miles weighted by survival rates from the U.S. Department of Transportation (2008). The marginal external benefit per vehicle is given by:

$$\frac{\partial E}{\partial \sigma} = b\frac{M}{\sigma^2}(1-\xi) - k\frac{M}{\sigma}\xi \tag{3.14}$$

where $\partial E/\partial \sigma$ is the marginal externality, b is the marginal external benefit of reducing gasoline consumption, k is the marginal external cost of increasing miles traveled, σ is the fuel-economy standard, M is miles traveled, and ξ is the elasticity of miles with respect to mileage. Discounting benefits at an annual rate of say 3% would reduce the magnitude of the benefit estimate slightly but would not change its sign.

CHAPTER IV

Conclusion

In chapter II of this dissertation, I develop a model that explicitly links the distribution of household preferences for ethanol as a gasoline substitute to aggregate price responses. The model allows me to extract information about micro-preferences from aggregate data on ethanol quantities and relative fuel prices. I do not need to observe gasoline quantities, as in other methodologies that match predicted and observed market shares. I estimate the model using data for ethanol sales volumes and relative fuel prices at a large number of retail fueling stations. I use a semi-parametric approach and other methods to estimate elasticities flexibly as a function of relative fuel prices, thereby revealing the distribution of household preferences for ethanol. Future research could apply this model and methodology to estimate preferences for other goods with perfect subsitutes.

I find that demand for ethanol as a gasoline substitute is sensitive to relative fuel prices, with elasticities that range from 2.5–3.0. Price responses are considerably smaller and less variable than they would be if household preferences for ethanol as a gasoline substitute were nearly identical. Fuel-switching behavior extends over a wide range of relative prices, and demand is not especially responsive to price changes at any particular point. These results imply that preferences are heterogeneous. Some households require that ethanol be discounted heavily relative to gasoline, while others require smaller discounts or may even be willing to pay a premium for ethanol.

These results have important implications for policy analysis. Accounting for heterogeneity cuts the cost of a national ethanol content standard in half. While the average household may require a large subsidy before choosing ethanol, households with strong preferences switch to ethanol with minimal price distortion. Similar intuition likely applies for policies that promote higher market shares for other "green" substitutes, such as renewable electricity, energy-efficient lighting and appliances, hybrid vehicles, or organic foods. Researchers should focus on marginal households when assessing the impacts of policy; assuming mean preferences for all households can yield misleading results.

The ethanol content standard nevertheless remains a costly policy. Costs per gallon of gasoline saved or ton of carbon emissions avoided exceed most conventional estimates of marginal external damages by a wide margin. Even after revising the analysis in ethanol's favor, the ethanol content standard can not be justified on efficiency grounds. If land-use changes associated with growing feedstocks negate ethanol's climate benefits, then the policy actually increases greenhouse emissions. Policymakers should seek to regulate emissions or tax externalities directly, as the ethanol content standard is likely to do more harm than good.

In chapter III of this dissertation, my coauthor and I analyze the market for flexible-fuel vehicles that burn ethanol. While interesting in its own right, this market indirectly provides information about the cost of tightening fuel-economy standards. Efforts to reduce gasoline consumption in the United States largely focus on mandating vehicle efficiency through Corporate Average Fuel Economy (CAFE) standards. The merits of these standards are not always clear, in part because it is difficult to measure the cost of regulation in the absence of market prices and because automakers have an incentive to overstate the costs of compliance. Domestic automakers claim that aggressive increases in CAFE standards would cost them tens-of-billions of dollars in profit, force them to close plants and cut tens-of-thousands of jobs, increase car prices by thousands of dollars, and "cripple" the domestic auto industry (Byrne 2003; Bloomberg News 2007; Shepardson 2007). Automakers do not publicly state net compliance costs in terms that we estimate here, but they always claim that the costs are high.

We estimate the cost of marginally tightening CAFE standards as revealed by profit-maximizing behavior in the auto industry. We demonstrate that automakers exploit an incentive or "loophole" in CAFE regulation that allows them to relax CAFE standards by producing flexible-fuel vehicles. We show theoretically that constrained automakers will equate the marginal cost of improving fuel economy using flexible-fuel vehicles with the marginal cost of improving fuel economy through other means. Thus, because we can observe the cost of producing a flexible-fuel vehicle, automakers that produce flexible-fuel vehicles indirectly reveal their marginal compliance costs. Based on this approach, we estimate that tightening CAFE standards by one mile per gallon would cost domestic automakers at most \$10–\$30 in profit per vehicle. Our estimate is similar to another recent estimate in the literature, which was obtained using alternative methods.

The CAFE standards program is an important policy in a high-profile industry. Recent legislation has scheduled significant increases in CAFE standards for the coming decades, and policymakers are likely to pursue further increases in vehicle efficiency, given that the personal transportation sector accounts for a large share of petroleum consumption and greenhouse gas emissions. The standards are politically controversial, however, because domestic automakers are perceived as less capable of producing efficient vehicles than their Japanese counterparts. Research that reliably quantifies the cost of CAFE regulation is therefore important. Our estimates for marginal compliance costs directly address this research need. Future research can use our estimates and approach when estimating compliance costs and comparing CAFE to alternative policies, such as an increase in the tax on gasoline.

Our methodology may also prove useful for researchers investigating the costs of regulation in other industries. For example, "incentive zoning" laws in some cities allow developers to relax height and density restrictions on new structures by setting aside open space or providing other public goods. Researchers could infer the benefits of relaxing zoning restrictions by observing how much developers spend on public goods to relax these restrictions. In addition, when government budgets become tight, and funding for direct provision of public goods becomes scarce, policymakers may seek to create incentives for public goods in other areas by modifying existing regulations. Firms that use these incentives may reveal information about compliance costs. Finally, energy and environmental regulation is likely to increase significantly in the coming decades, as policymakers around the world grapple with climate change. We suspect that some of these regulations will feature incentives and loopholes like the one we use here to uncover CAFE compliance costs.

Bibliography

Alson, J. (2008, May). Personal communication with Jeff Alson of the U.S. Environmental Protection Agency.

Alves, D. C. and R. D. L. d. S. B. Bueno (2003). Short-run, long-run and cross elasticities of gasoline demand in Brazil. *Energy Economics 25*(2), 191–199.

Barnett, W. A., D. Fisher, and A. Serletis (1992, December). Consumer theory and the demand for money. *Journal of Economic Literature 30*(4), 2086–2119.

Barrionuevo, A. and M. Maynard (2006, May 1). Dual-fuel vehicles open mileage loophole for carmakers. *New York Times*.

Bloomberg News (2007, June 16). Fuel plan would 'cripple' Chrysler, chief says. *Bloomberg News*.

Buccelli, M. (2007, February). Personal communication with Mark Buccelli of the Minnesota Department of Commerce, Weights and Measures Division.

Byrne, R. (2003). Life in the slow lane: Tracking decades of automaker roadblocks to fuel economy. Union of Concerned Scientists Report.

Corrado, C., W. Dunn, and M. Otoo (2006, January). Incentives and prices for motor vehicles: What has been happening in recent years? Federal Reserve Board, Finance and Economics Dicussion Series, Working Paper 2006-09.

Dahl, C. and T. Sterner (1991, July). Analyzing gasoline demand elasticities: A survey. *Energy Economics 13*(3), 203–210.

Ellison, S. F., I. Cockburn, Z. Griliches, and J. Hausman (2006, Autumn). Characteristics of demand for pharmaceutical products: An examination of four cephalosporins. *The RAND Journal of Economics 28*(3), 426–446.

Espey, M. (1996). Explaining the variation in elasticity estimates of gasoline demand in the United States: A meta-analysis. *Energy Journal 17*(3), 49–60.

Espey, M. (1998). Gasoline demand revisited: An international meta-analysis of elasticities. *Energy Economics 20*(3), 273–295.

Espey, M. and S. Nair (2005). Automobile fuel economy: What is it worth? *Contemporary Economic Policy 23*(3), 317–323.

Fan, J. and I. Gijbels (1996). *Local Polynomial Modelling and Its Applications*, Volume 66 of *Monographs on Statistics and Applied Probability*. Boca Raton, FL: CRC Press.

Fargione, J., J. Hill, D. Tilman, S. Polasky, and P. Hawthorne (2008, February). Land clearing and the biofuel carbon debt. *Science 29*(319), 1235–1238.

Farrell, A. E., R. J. Plevin, B. T. Turner, A. D. Jones, M. O'Hare, and D. M. Kammen (2006). Ethanol can contribute to energy and environmental goals. *Science 311*, 506–508.

Goldberg, P. K. (1998, March). The effects of the corporate average fuel efficiency standards in the U.S. *The Journal of Industrial Economics 46*(1), 1–33.

Graham, D. J. and S. Glaister (2002, January). The demand for automobile fuel: A survey of elasticities. *Journal of Transportation Economics and Policy 36*(1), 1–26.

Greene, D. L. (1989). Motor fuel choice: An econometric analysis. *Transportation Research A 23A*(3), 243–253.

Griliches, Z. and J. A. Hausman (1986). Errors in variables in panel data. *Journal of Econometrics 31*, 93–118.

Harrington, W., I. Parry, and M. Walls (2007). Automobile externalities and policies. *Journal of Economic Literature 45*, 374–400.

Harris Interactive (2006, July 26). Two-thirds of adult vehicle owners are familiar with flexible fuel vehicles and more than half are interested in purchasing one. http://www.harrisinteractive.com/NEWS/allnewsbydate.asp?NewsID=1074.

Hausman, J. (2001, Autumn). Mismeasured variables in econometric analysis: Problems from the right and problems from the left. *Journal of Economic Perspectives 15*(4), 57–67.

Hausman, J. A. and W. K. Newey (1995, November). Nonparametric estimation of exact consumers surplus and deadweight loss. *Econometrica 63*(6), 1445–1476.

Holland, S. P., C. R. Knittel, and J. E. Hughes (2007, May). Greenhouse gas reductions under low carbon fuel standards? NBER Working Paper 13266.

Hughes, J. E., C. R. Knittel, and D. Sperling (2008, January). Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal 29*(1).

Jacobson, M. R. (2007, December). Evaluating U.S. fuel economy standards in a model with producer and household heterogeneity. UC San Diego Working Paper.

Kennan, J. (1989). Simultaneous equations bias in disaggregated econometric models. *The Review of Economic Studies 56*(1), 151–156.

Kilian, L. (2007, November). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. University of Michigan Working Paper.

Kilian, L. (2008, October). The economic effects of energy price shocks. University of Michigan Working Paper.

Kleit, A. N. (2002, Fall). CAFE changes, by the numbers. *Regulation 25*(3), 32–35.

Kohn, J. (2000, June 19). GM: Some SUVs to use ethanol. *Automotive News*.

Liu, Y. (2007). *The Causes and Effects of the Alternative Motor Fuels Act*. Ph. D. thesis, University of Michigan.

Liu, Y. and G. E. Helfand (2008). The effects of the CAFE standard and incentive program on alternative fuel vehicles and greenhouse gas emissions. University of Michigan Working Paper.

Minnesota Department of Agriculture (2003, May). Economic impact of the ethanol industry in Minnesota. Technical report.

Morris, M. (2000). *Incentive Zoning: Meeting Urban Design and Affordable Housing Objectives*. APA Planning Advisory Service.

National Academy of Sciences (2002). *Effectiveness and Impact of Corporate Average Fuel Economy (CAFE) Standards*. Washington, DC: National Academies Press.

National Petroleum News (2006, July). Retail market: Market facts. *National Petroleum*

News 98(8), 96–109.

Nevo, A. (2001, March). Measuring market power in the ready-to-eat cereal industry. *Econometrica 69*(2), 307–342.

Perkins, M. and S. Barros (2006, February). Brazil sugar ethanol update. Global Agriculture Information Network Report BR6001, U.S. Department of Agriculture, Foreign Agricultural Service.

Phillips, O. R. and D. P. Schutte (1988, February). Identifying profitable self-service markets: A test in gasoline retailing. *Applied Economics 20*(2), 263–272.

Phoenix Automotive (2006, May). E85 ethanol flexible-fuel vehicle survey. Prepared for the American Lung Association of Minnesota.

Rask, K. N. (1998, June). Clean air and renewable fuels: The market for fuel ethanol in the U.S. from 1984-1993. *Energy Economics 20*(3), 325–345.

Robinson, P. M. (1988). Root-n-consistent semi-parametric regression. *Econometrica 55*, 931–954.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy 82*(1), 34–55.

Rubin, J. and P. Leiby (2000). An analysis of alterntive fuel credit provisions of U.S. automotive fuel economy standards. *Energy Policy 13*, 589–601.

Sallee, J. (2007, October). The incidence of tax credits for hybrid vehicles. Unpublished.

Searchinger, T., R. Heimlich, R. A. Houghton, F. Dong, A. Elobeid, J. Fabiosa, S. Tokgoz, D. Hayes, and T.-H. Yu (2008, February). Use of U.S. croplands for biofuels increases greenhouse gases through emissions from land-use change. *Science 29*(319), 1238–1240.

Shepardson, D. (2007, June 20). Auto lobby blitzes Senate. *Detroit News*.

Small, K. A. and K. V. Dender (2007). Fuel efficiency and motor vehicle travel: The declining rebound effect. *Energy Journal 28*(1), 25–51.

Solon, G. (1984, February). Estimating autocorrelations in fixed-effects models. NBER Technical Working Paper 32.

Stern, N. (2006). *Stern Review: The Economics of Climate Change*. Cambridge, MA: Cambridge University Press.

Tol, R. S. (2005, November). The marginal damage costs of carbon dioxide emissions: An assessment of the uncertainties. *Energy Policy 33*(16), 622–634. 2064-2074.

U.S. Department of Energy (2006, April). *Handbook for Handling, Storing, and Dispensing E85*.

U.S. Department of Transportation (2004, October). Automotive fuel economy manufacturing incentives for alternative fueled vehicles. Final Rule.

U.S. Department of Transportation (2008, April). Corporate average fuel economy for model years 2011–2015 passenger cars and light trucks. Preliminary Regulatory Impact Analysis.

U.S. Department of Transportation, U.S. Department of Energy, and U.S. Environmental Protection Agency (2002, March). Effects of the alternative motor fuels act CAFE incentives policy. Report to Congress.

U.S. Environmental Protection Agency (1990, April). Analysis of the economic and environmental effects of ethanol as an automotive fuel. Office of Mobile Sources Special Report.

U.S. Environmental Protection Agency (2000-2006). *Model Year 2000–2006 Fuel Economy Guide*.

Wald, M. L. (2005, November 4). Ford to subsidize ethanol pumps in Midwest. *The New York Times*.

Weiss, M. A. (1992, Spring). Skyscraper zoning: New York's pioneering role. *Journal of the American Planning Association 58*(2), 201–212.

Williams, B. (2008, April 17). Lucerne joins GM's flex-fuel lineup. *Automotive News*.

Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. MIT Press.

Yatchew, A. (1997). An elementary estimator of the partial linear model. *Economics Letters 57*, 135–143.

Yatchew, A. (1998, June). Nonparametric regression techniques in economics. *Journal* *of Economic Literature 36*(2), 669–721.

Yatchew, A. and J. A. No (2001, November). Household gasoline demand in Canada. *Econometrica 69*(6), 1697–1709.