Fuel Economy Regulations and Efficiency Technology Improvements in U.S. Cars Since 1975

by

Donald Warren MacKenzie
B.A.Sc. Chemical & Biological Engineering, University of British Columbia (2001)
S.M. Technology & Policy, Massachusetts Institute of Technology (2009)

Submitted to the Engineering Systems Division
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Engineering Systems: Technology, Management, and Policy
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2013

Author . . . .
Engineering Systems Division
May 18, 2013

Certified by . . . .
Kenneth Oye
Associate Professor of Political Science and Engineering Systems
Doctoral Committee Chair

Certified by . . . .
John Heywood
Professor of Mechanical Engineering and Sun Jae Professor, Emeritus
Thesis Supervisor

Certified by . . . .
Charles Fine
Chrysler Leaders for Global Operations Professor of Management
Doctoral Committee Member

Certified by . . . .
Christopher Knittel
William Barton Rogers Professor of Energy Economics
Doctoral Committee Member

Accepted by . . . .
Olivier L. de Weck
Professor, Aeronautics and Astronautics and Engineering Systems
Chair, Engineering Systems Division Education Committee

Fuel Economy Regulations and Efficiency Technology Improvements in U.S. Cars Since 1975

by

Donald Warren MacKenzie

Submitted to the Engineering Systems Division
on May 18, 2013, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Engineering Systems: Technology, Management, and Policy

Abstract

Light-duty vehicles account for 43% of petroleum consumption and 23% of greenhouse gas emissions in the United States. Corporate Average Fuel Economy (CAFE) standards are the primary policy tool addressing petroleum consumption in the U.S., and are set to tighten substantially through 2025. In this dissertation, I address several interconnected questions on the technical, policy, and market aspects of fuel consumption reduction.

I begin by quantifying historic improvements in fuel efficiency technologies since the 1970s. First, I develop a linear regression model of acceleration performance conditional on power, weight, powertrain, and body characteristics, showing that vehicles today accelerate 20-30% faster than vehicles with similar specifications in the 1970s. Second, I find that growing use of alternative materials and a switch to more weight-efficient vehicle architectures since 1975 have cut the weight of today's new cars by approximately 790 kg (46%). Integrating these results with model-level specification data, I estimate that the average fuel economy of new cars could have tripled from 1975–2009, if not for changes in performance, size, and features over this period. The pace of improvements was not uniform, averaging 5% annually from 1975–1990, but only 2% annually since then. I conclude that the 2025 standards can be met through improvements in efficiency technology, if we can return to 1980s rates of improvement, and growth in acceleration performance and feature content is curtailed.

I next test the hypotheses that higher fuel prices and more stringent CAFE standards cause automotive firms to deploy efficiency technologies more rapidly. I find some evidence that higher fuel prices cause more rapid changes in technology, but little to no evidence that tighter CAFE standards increase rates of technology change. I conclude that standards alone, without continued high gasoline prices, may not drive technology improvements at rates needed to meet the 2025 CAFE standards.

Finally, I discuss the political economy of state and federal fuel economy standards. I develop a simple model of automotive manufacturers' responses to alternative systems of fuel economy regulation, using it to demonstrate the importance of several factors determining industry support for nationwide fuel economy regulations.

Doctoral Committee Chair: Kenneth Oye
Title: Associate Professor of Political Science and Engineering Systems

Thesis Supervisor: John Heywood
Title: Professor of Mechanical Engineering and Sun Jae Professor, Emeritus

Doctoral Committee Member: Charles Fine
Title: Chrysler Leaders for Global Operations Professor of Management

Doctoral Committee Member: Christopher Knittel
Title: William Barton Rogers Professor of Energy Economics

Acknowledgments

My successful completion of this dissertation would not have been possible without the unwavering support of my family, colleagues, mentors, and sponsors.

Above all, I thank my wife, Kate, for her love, patience, and support over the last six years. I am fortunate to have a wonderful son in Duncan, who reminds me daily of what is most important. I thank my parents, Neil MacKenzie and Peggy LeBlanc, for their support of my education, and the examples they have set for me in all aspects of life.

The colleagues from the Sloan Automotive Laboratory and Engineering Systems Division who have helped me along the way are too numerous to name. I am particularly indebted to Stephen Zoepf, Valerie Karplus, David Keith, Lynette Cheah, Erik Wilhelm, and Jeff McAulay for the balanced doses of criticism, realism, and optimism that have suffused many a discussion.

The members of my doctoral committee have been outstanding mentors and wonderful to work with. As my advisor for the past six years, John Heywood has taught me the value of always remembering the broader context of my work, and the power of a single, pointed question. Ken Oye has been a great cheerleader, but never shy about telling me bluntly when I have said something foolish. Charlie Fine has pushed me to consider multiple perspectives, and to speak to multiple audiences. I was very fortunate to have Chris Knittel come to MIT when he did, as he has been a great teacher both inside and outside the classroom.

Many others at MIT and beyond have provided helpful insights over the years, helping me to ask the right questions and to find answers to them. In particular, I thank Jim Womack, Dan Whitney, Chris Magee, Qi Hommes, John DeCicco, and William Chernicoff. Kevin Green, Don Pickrell, and others at the Volpe Center have provided helpful feedback as well, particularly on Chapter 5 of this dissertation. I would never have found myself where I am today without the opportunities and guidance afforded to me by David Friedman, Michelle Robinson, Jason Mark, and Kevin Smith.

The logistics of graduate work are sometimes baffling, so I thank Beth Milnes of ESD for making it much easier. I also thank Janet Maslow and Karla Stryker-Currier for keeping the Sloan Lab running smoothly, facilitating travel, and helping to make sure meetings went off without a hitch.

Finally, I am grateful to all of those who have supported my research and studies financially these past six years. The William Asbjornsen Albert Memorial Fellowship brought me to MIT, and since then I have had benefited from the support of many benefactors, including the MIT Energy Initiative, Eni, Chevron, the Roos (1944) Fund for Alternative Energy, CONCAWE, and other supporters of John Heywood's Future Vehicles and Fuels research group. Ken Rose of CONCAWE deserves special thanks for his enthusiasm, interest, and deep engagement with our group's work.

Several chapters of this disseration are based on papers already published or presented at conferences. I thank my co-authors for blessing my inclusion of these materials in this dissertation.

Chapter 2 is based on a paper co-authored with John Heywood. I thank David Champion and his staff at *Consumer Reports* magazine for sharing the data underlying this chapter, and for discussing acceleration testing methodologies with me. I also thank Dan Edmunds of Edmunds.com for discussing acceleration testing practices.

Chapter 3 is based on a paper co-authored with Stephen Zoepf and John Heywood. I thank Audatex North America, and the American Chemistry Council for providing data that were critical to the work reported in this chapter.

Chapter 4 relies on vehicle attribute data generously provided by Jeff Alson and Aaron Hula at U.S. EPA, and I thank them for their help in obtaining and interpreting these data.

Contents

List of Figures

List of Tables

Chapter 1

Efficiency Technology, Design Tradeoffs, and Fuel Economy Regulations in U.S. Automobiles

Cars and light trucks in the United States account for nearly half (43%) of the country's petroleum consumption and one-quarter (23%) of its greenhouse gas (GHG) emissions (EIA, 2011; EPA, 2013a; Kromer and Heywood, 2007). Falling petroluem demand and rising domestic production have reduced our reliance on petroleum imports in recent years, but imports still account for almost half (46%) of U.S. petroleum supply (EIA, 2013). In addition, the U.S. transportation sector remains heavily dependent on petroleum and is consequently sensitive to fluctuations in the global price of oil. For these reasons, mitigating petroleum dependence and GHG emissions are seen as important policy goals. This is exemplified, for example, by the opening text of the Energy Independence and Security Act of 2007, which articulated the following purposes: *To move the United States toward greater energy independence and security, to increase the production of clean renewable fuels, to protect consumers, to increase the efficiency of products, buildings, and vehicles, to promote research on and deploy greenhouse gas capture and storage options, and to improve the energy performance of the Federal Government, and for other purposes.*

Broadly speaking, there are three principal strategies that can reduce GHG emissions and petroleum consumption from personal transportation. First, we can reduce the distance that vehicles are driven. Second, we can consume less energy per unit of distance traveled. Thirdly, we can reduce the petroleum or GHG intensity of our fuel mix. In the U.S., most policy efforts, particularly at the national level, are focused on reducing per-mile energy consumption. The Energy Policy and Conservation Act (EPCA) has since 1978 required manufacturers of new automobiles to ensure that their products meet a certain minimum level of fuel economy, on a sales-weighted average basis.[1] These Corporate Average Fuel Economy (CAFE) standards are to be set administratively at the "maximum feasible" level, taking into account factors that include "technological feasibility" and "economic practicability."

By definition, establishing sound fuel economy policy demands a clear understanding of what is technologically feasible and economically practical. This means, firstly, having a clear idea of how much technology might improve in the future. Secondly, it requires an understanding of what we are gaining (or, potentially, losing) by applying those technological improvements to reducing fuel consumption. Improving certain attributes valued by consumers — particularly acceleration performance, size, and features that add weight — tends to increase a vehicle's fuel consumption. As such, new technologies can be used either to reduce fuel consumption or to offset the fuel consumption penalties of other design changes. Understanding the relationships between efficiency technologies, fuel consumption, and these other vehicle attributes is essential to developing standards that are technologically feasible and economically practical.

The objective of this dissertation, in a broad sense, is to improve our understanding of what is technologically feasible and economically practical in the context of automotive fuel economy in the U.S. To do this, I explore the scope of historic fuel efficiency technology improvements in U.S. automobiles, trends in several key

---

[1] Throughout this dissertation, I will refer frequently to both *fuel economy* and *fuel consumption.* I use fuel consumption to refer to the amount of fuel that a vehicle consumes for a specified distance traveled, in units such as liters per 100 km or gallons per 100 miles. Fuel consumption therefore is the inverse of fuel economy, which is commonly expressed as miles per gallon. Fuel consumption in this sense should not be confused with the total quantity of fuel consumed by all vehicles combined; even as average fuel consumption has fallen over the years, the total quantity of fuel consumed has increased as the number of vehicles has grown and their collective distance traveled has increased.

attributes closely related to fuel consumption, and the relationships between fuel economy policies, the adoption of efficiency technologies, and vehicle designs.

The results of this dissertation will be relevant to policymakers implementing automotive fuel economy policy, and to those studying energy efficiency technology and standards more broadly. If fuel efficiency technology improves at rates determined exogenously to the policy environment, then policymakers must have a clear basis for predicting of how much technology will improve in the future. If, as is widely believed, technology changes more quickly in response to higher gasoline prices or stricter fuel economy standards, then sound policymaking demands an awareness of this fact and an estimate of the size of any such effect. Finally, policymakers concerned about economic practicability of fuel economy standards must understand the tradeoffs between fuel consumption and other attributes valued by consumers, as well as the underlying trends in those attributes. This dissertation addresses all of these issues.

1.1 Overview of Dissertation

This dissertation is organized into chapters addressing a sequence of interconnected topics related to automotive technology, vehicle attribute tradeoffs, and fuel economy policy. While the chapters build on and incorporate results from prior chapters, each one is written so as to be largely self-contained, and the interested reader should be able to approach them in a modular fashion.

Chapter 2 presents a set of linear regression models for estimating the acceleration performance of cars and trucks based on their power, weight, and key powertrain and body style characteristics. Key questions addressed in this chapter include:

- How can we concisely and conveniently estimate the acceleration performance of light-duty vehicles using commonly-reported attributes?
- Do vehicles today accelerate differently than vehicles in the past, even for the same level of power, weight, and key powertrain and body characteristics?

- How has the distribution of acceleration performance across new vehicles changed since the 1970s?

Chapter 3 examines changes in the weight of new cars in the U.S. since 1975, considering both weight-saving technologies and weight-increasing features and functionality improvements. This chapter addresses the following questions:

- How was the weight of the average new car in the U.S. reduced so precipitously in the late 1970s and early 1980s, only to rebound again in subsequent years?
- How much more would today's cars weigh, if not for the spread of weight-efficient architectures and materials since 1975?
- How much weight has the growth in size and the addition of safety, emissions, and comfort and convenience features added to the average new car since 1975?

Chapter 4 quantifies the improvements in efficiency technology in new cars since 1975, employing a system-level perspective that focuses on attributes relevant to consumers. The chapter relies on estimates of acceleration performance based on the methods reported in Chapter 2, and on estimates of weight-saving technologies reported in Chapter 3. The chapter addresses questions that include:

- Holding all else equal, by how much does a car's fuel consumption increase in response to a 1% increase in weight? A 1% reduction in fuel consumption?
- How much could the per-mile fuel consumption of new cars in the U.S. have been reduced since 1975, if functionality (including size, acceleration performance, and feature content) had remained unchanged over this time?
- How much of this potential has actually been realized as reductions in fuel consumption? How much has been needed to offset the effects of changes in acceleration, size, and feature content?
- Have these patterns been stable over time?

Chapter 5 empirically tests the hypothesis that tighter fuel economy standards will themselves spur faster technology change. This chapter employs model-level vehicle specifications and firm-level regulatory compliance data from 1978–2008 to address the following questions:

- Does the efficiency technology in a firm's fleet of new cars change more quickly when that firm is more tightly constrained by a CAFE standard?
- Does technology improve more quickly when gasoline prices are higher?

Chapter 6 discusses the political economy of fuel economy standards, and how it helps to explain recent rulemaking outcomes on automotive fuel consumption and GHG standards. The chapter presents a modeling framework to simulate the responses of automobile manufacturers to either uniform federal fuel economy standards, or "nested" standards in which a subset of states ("adopting states") adopt a tighter standard than the prevailing federal standard. Integrating results from the prior chapters, I present a simple implementation of this model for a single vehicle class, and address the following questions:

- How might firms respond to a stricter federal fuel economy standard, in terms of adoption of new technologies, giving up improvements in acceleration performance, and pricing of their products?
- How do firms' optimal strategies depend on the market size in adopting states, the relative stringency of standards in adopting states and federally, and the cost of developing multiple variants of their products?
- What are the implications of these results for firms' support of nationwide standardization of fuel consumption standards, and for the "leakage" of emissions from more strictly regulated regions to more loosely regulated regions?

Chapter 7 summarizes and integrates some of the key results of the dissertation, and reflects on their implications for meeting future fuel economy standards.

Chapter 2

Analyzing Acceleration Trends in U.S. Light-Duty Vehicles

This chapter is based on a paper jointly authored with John Heywood, published in Transportation Research Record: Journal of the Transportation Research Board, No. 2287*, pp. 122–131. ©National Academy of Sciences, Washington, D.C., 2012. Material in this chapter is reproduced with permission of the Transportation Research Board. None of this implies endorsement by TRB of a product, method, practice, or policy.*

The acceleration performance of light-duty vehicles has implications for the energy usage of those vehicles, their attractiveness to consumers, and how they are driven. Despite this importance, many investigators rely on correlations from the 1970s for estimating performance. This chapter presents a set of linear regression models for estimating acceleration times from 0–48, 0–97, and 72–105 km/h (0–30, 0–60, and 45–65 mph), based on engine power, vehicle weight, body style, and basic powertrain characteristics of more than 1000 vehicles tested by *Consumer Reports* magazine between 1975 and 2010. Importantly, the results include estimates of fixed effects for each year, capturing technological improvements not directly observed in the data set and making the models appropriate for estimating performance of vehicles from many different model years. Results indicate that contemporary vehicles are better able to transform engine power into acceleration performance than were vehicles in the past, yielding acceleration times 20–30% faster than comparable vehicles in the 1970s. Most of this improvement appears to have occurred before 1990, and the estimated effect is larger for 0–48 km/h acceleration than for higher-speed acceleration. One of the reported models was applied to historic sales and specification data for United States vehicles, and the results indicate that new vehicles in the U.S. today average 8.8 seconds from 0–97 km/h, 0.9 seconds (10%) faster than previously thought. Interestingly, the trends in 0–97 km/h acceleration times are consistent with exponential decay toward an asymptote, and today's vehicles are within one second of the estimated asymptotic acceleration time. The models reported in this chapter also will be applied to estimate the 0–97 km/h acceleration times of new cars as part of the work reported in Chapter 4.

2.1 Introduction

Understanding vehicle acceleration performance is important to transportation analysts and researchers for several reasons. First, acceleration can be traded off against other vehicle attributes, including fuel economy. All else being equal, larger improvements in acceleration performance over time mean smaller improvements in fuel economy, leading to higher energy consumption. Second, the acceleration performance of a vehicle can affect its utility to consumers, influencing purchase decisions. Finally, acceleration capabilities may influence how aggressively vehicles end up being driven, affecting in-use fuel consumption (Berry, 2010). The objective of this work was to develop an improved method for estimating vehicle acceleration performance using other vehicle attributes, and to quantify the annual improvements in acceleration performance that are due to factors beyond basic ones like increased power-to-weight ratio. The average acceleration performance of new vehicles sold in the United States has been improving steadily since the early 1980s, while fuel consumption has changed relatively little. Figure 2-1 shows the average 0–97 km/h (0–60 mph) acceleration times and average fuel consumption for new cars and light trucks in the U.S. since 1975, as reported by the EPA (2010). Despite the substantial reductions in acceleration times that are evident in these oft-cited numbers, this chapter will show that the actual rate of change has been even faster, and that the estimates in Figure 2-1 overstate acceleration times for contemporary U.S. vehicles by approximately 0.9

seconds, or 10%.

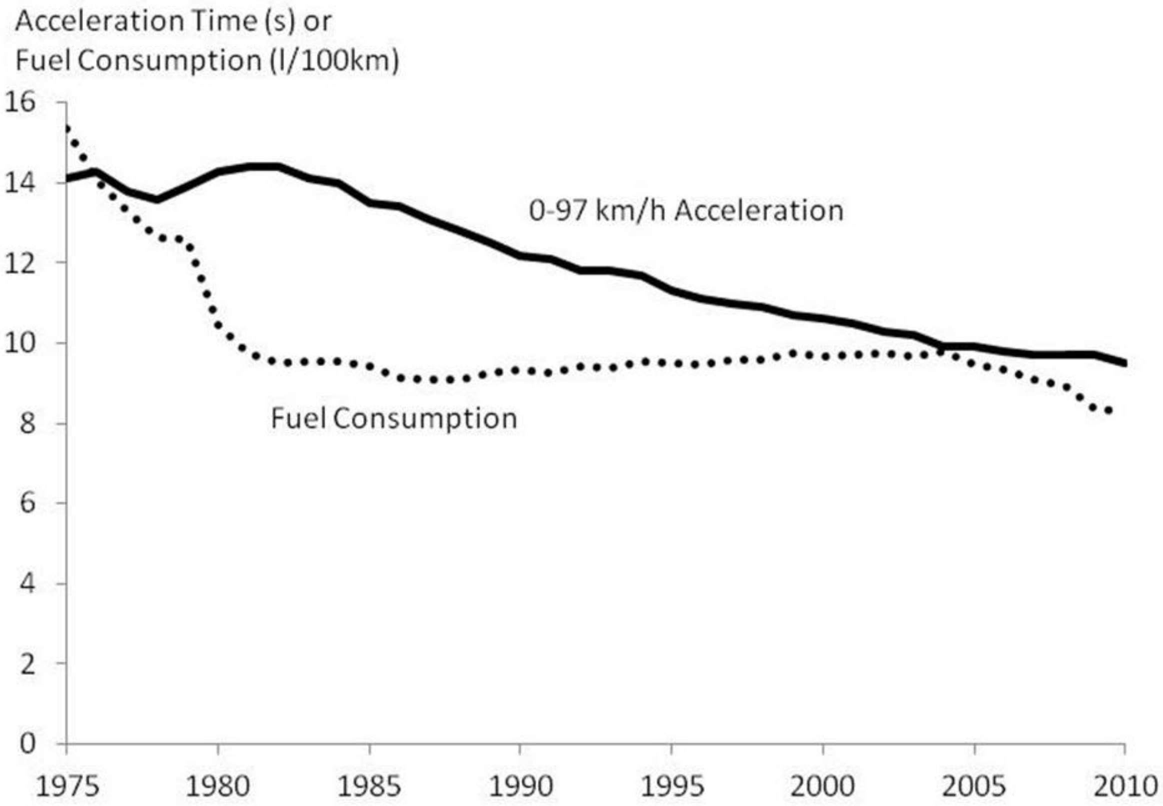

Figure 2-1: Average 0–97 km/h acceleration times and average unadjusted fuel consumption for new U.S. automobiles, 1975–2010, as reported by EPA (2010). The actual rate of reduction in acceleration times has been faster than shown here.

Acceleration performance is commonly reported as the time needed to accelerate between two speeds at wide-open throttle. Three common acceleration metrics are investigated in this chapter:

- Z48: Time to accelerate from 0–48 km/h (0–30 mph)
- Z97: Time to accelerate from 0–97 km/h (0–60 mph)
- P72105: Time to accelerate from 72–105 km/h (45–65 mph); P denotes "passing acceleration."

These measures are reported by *Consumer Reports* magazine based on their own testing, and are the same measures that were previously investigated for a more limited sample of vehicles by Santini and Anderson (1993). Despite the importance of acceleration performance, a key challenge to incorporating it into analytical work is that comprehensive databases containing standardized acceleration measurements are difficult or impossible to find, especially for the investigator who also requires reliable data on other vehicle attributes, sales volumes, and the like. For these investigators, detailed simulation of vehicle performance may be neither necessary nor practical. Instead, a convenient means to estimate vehicle acceleration performance based on other characteristics is desirable. Therefore, many authors rely on a correlation originally published in 1976 for estimating a vehicle's acceleration performance based on its power-to-weight ratio (Malliaris et al., 1976). This correlation continues to be used in academic papers (Knittel, 2011), reports (Greene, 2001; Bandivadekar et al., 2008), and government data like those summarized in Figure 2-1 (EPA, 2010). These government data are themselves used as the basis for analyses of acceleration trends (Lutsey and Sperling, 2005; An and DeCicco, 2007). One objective of this work was to develop an improved method for estimating acceleration performance, which would be more applicable to modern vehicles and more robust over time. A second objective was to quantify the changes in acceleration that are not captured by the power to weight variable that forms the basis of so many estimates.

Malliaris et al. (1976) estimated a model of 0–97 km/h acceleration time using the following form, where P is engine peak power, IWT is the inertia weight, and F and f are constants:

$$Z97 = F\left(\frac{P}{IWT}\right)^{-f} \tag{2.1}$$

They noted the importance of many factors other than the power to weight ratio for determining acceleration performance, such as drivetrain characteristics and the engine's torque curve. However, they argued, the power to weight ratio "is overwhelmingly influential and allows by itself and adequate description of the acceleration performance." They estimated different values of the parameters F and f for vehicles with manual transmissions and for those with automatics, using acceleration times reported in the automotive enthusiast literature (e.g. *Car & Driver* magazine) for model years 1974 and 1975.

Young (1991) updated the analysis of Malliaris, Hsia, and Gould, using similar sources of performance data for model years 1989-1990. She investigated several functional forms, including linear forms, and considered including engine displacement and axle ratio as additional explanatory variables. Young concluded that the best model was one with the same form as that advanced by Malliaris, Hsia, and Gould (Equation 2.1), though she recommended eliminating the distinction between automatic and manual transmissions.

Santini and Anderson (1993) made several improvements upon the methods of earlier investigators. First, they noted that the earlier functional forms placed unnecessary constraints on parameter values, by requiring that the exponent for power be the negative of the exponent for weight. They adopted a more general functional form, noting that the model shown in Equation 2.1 is a highly restricted form of their model, where CWT_i is the curb weight of vehicle i, D_i is its engine displacement, A_i is a surrogate for its frontal area, C_i is a dummy variable indicating that vehicle i is a car, V_i is a dummy variable indicating that it is a van, and X_i is a vector of dummy variables denoting various engine technology packages included on vehicle i:

$$\begin{aligned} ln(ACC_i) = \beta_0 + \beta_1 ln(P_i) + \beta_2 ln(CWT_i) + \beta_3 ln(D_i) + \beta_4 ln(A_i) + \beta_5 C_i ln(A_i) \\ + \beta_6 V_i ln(A_i) + \boldsymbol{\beta_X X_i} + \epsilon_i \end{aligned} \tag{2.2}$$

Santini and Anderson also argued that relying on the automotive enthusiast literature could be problematic because of inconsistency in the testing methods used. Instead, they used performance testing data reported by a single publication, *Consumer Reports*, for 107 vehicles from model years 1986–1988.

Santini and Anderson found that in addition to power and weight, important determinants of acceleration performance included engine displacement, transmission type, body type, and frontal area. They found that generally, the inclusion of specific engine technologies did not significantly affect acceleration performance after controlling for the major attributes listed above.

This chapter reports work that builds on Santini and Anderson's approach by employing a much broader data set spanning 1975–2010, while estimating fixed effects for year. These changes have two important consequences. First, they make the model more appropriate for estimating acceleration performance over multiple years. Second, the fixed effects can be interpreted as quantifying improvements in how effectively a vehicle transforms engine power into the acceleration of the vehicle's mass. The results indicate that for a given level of engine power and vehicle mass (and controlling for various powertrain and body characteristics) a contemporary vehicle delivers approximately 20-30% faster acceleration than a comparable new vehicle in 1977. The results also indicate that most of these gains occurred prior to 1990.

The remainder of this chapter is organized as follows: Section 2.2 describes the form of the linear regression model used to estimate the various measures of acceleration. Section 2.3 describes the data set that was used to fit the model. Section 2.4 contains the results of the model estimation, and discussion of the estimated parameter values. Section 2.5 discusses concerns over bias in the sample of vehicles selected for acceleration testing and applies the model in order to examine trends in the 0–97 km/h acceleration performance of U.S. vehicles since 1978. Section 2.6 summarizes some conclusions that can be drawn from the work.

2.2 Methodology

In this work, several model specifications were investigated, all of which use a general form similar to that advanced by Santini and Anderson (1993). In its most unrestricted form, the model used in this work is:

$$\begin{aligned} ln(ACC_{it}) = \beta_0 + \beta_1 ln(P_{it}) + \beta_2 ln(WT_{it}) + \beta_3 ln(D_{it}) + \beta_4 [ln(P_{it})]^2 + \beta_5 [ln(WT_{it})]^2 \\ + \beta_6 ln(P_{it}) ln(WT_{it}) + \beta_7 ln(P_{it}) ln(D_{it}) + \beta_8 TSpd_{it} \\ + \boldsymbol{\beta_T X_{it}^T} + \boldsymbol{\beta_{P,T}} ln(P_{it}) \boldsymbol{X_{it}^T} + \boldsymbol{\beta_E X_{it}^E} + \boldsymbol{\beta_{P,E}} ln(P_{it}) \boldsymbol{X_{it}^E} \\ + \boldsymbol{\beta_D X_{it}^D} + \boldsymbol{\beta_B X_{it}^B} + \boldsymbol{\beta_Y X_{it}^Y} + \epsilon_{it} \end{aligned} \tag{2.3}$$

In the above model, ACC is an acceleration metric, P is engine peak power, WT is vehicle weight (both curb weight and inertia weight were investigated), D is engine displacement, and $TSpd$ is the number of transmission speeds (defined as zero for continuously variable transmissions). $\boldsymbol{X^T}$ is a set of dummy variables for transmission type, $\boldsymbol{X^E}$ a set of dummy variables for engine type, $\boldsymbol{X^D}$ a set of dummy variables for drive type, and $\boldsymbol{X^B}$ a set of dummy variables for body style. $\boldsymbol{\beta_T}$, $\boldsymbol{\beta_E}$, $\boldsymbol{\beta_D}$, and $\boldsymbol{\beta_B}$ are vectors of fixed effects capturing the average effects of dummy variables $\boldsymbol{X^T}$, $\boldsymbol{X^E}$, $\boldsymbol{X^D}$, $\boldsymbol{X^B}$, respectively. The term $\boldsymbol{X^Y}$ is a set of dummy variables equal to 1 for year t and 0 for all other years. Thus, $\boldsymbol{\beta_Y}$ represents a vector of fixed effects estimating acceleration performance in each year relative to a base year, similar to the approach employed by Knittel (2011) in estimating technological progress for U.S. automobiles. Additional terms capture interaction effects of power with weight, displacement, transmission type, and engine type. The last term, ϵ, is an error term representing random variation due to factors not captured by the independent variables. The subscript i is an index for each vehicle model observation, while t denotes the model year of the vehicle.

Several restricted versions of the above model were also estimated, and results are also reported for a simplified model that includes only engine power and weight.

2.3 Data

The database used in the present work includes approximately 1,500 vehicles that were offered for sale in the U.S. between 1975 and 2010 and were tested by *Consumer*

Reports. Personal communications with testing directors at both *Consumer Reports* and Edmunds.com suggested that inconsistency in testing methods, as noted by Santini and Anderson (1993), continues to be an issue today, especially among enthusiast publications which compete to report the most aggressive performance numbers. As such, this work relies on testing by a single publication to eliminate the effects of this testing variation. The database includes a variety of engineering attributes and performance metrics, including:

- Curb weight
- Engine peak power, displacement, and type (naturally aspirated gasoline, turbodiesel, etc.)
- Transmission type and number of speeds
- Drive type (rear-, front-, four- or all-wheel drive)
- Body style (sedan, SUV, etc.)
- Acceleration performance from 0–48, 0–97, and 72–105 km/h (0–30, 0–60, and 45–65 mph)

The *Consumer Reports* database is not a random sample of vehicles offered in the U.S. Instead, its membership is determined by the decisions of *Consumer Reports*' staff. Nevertheless, the average weight and power of the *Consumer Reports* sample tracks reasonably well with the averages for all new vehicles sold in the U.S., as shown in Figure 2-2. Figure 2-2 does suggest that before 1990, *Consumer Reports* was somewhat biased toward testing lighter, less powerful vehicles. Personal communication with *Consumer Reports*' testing director indicated that *Consumer Reports* tends to test vehicles and configurations that they expect will be high-volume sellers, and which have recently undergone a redesign or refresh. The bias toward recently redesigned models may contribute to the lower average weight and power in the *Consumer Reports* sample in the earlier years. When power and weight were declining quickly, overall market averages would have lagged changes among vehicles that had been redesigned more recently. The issue of bias in the tested sample is discussed further in Section 2.5.

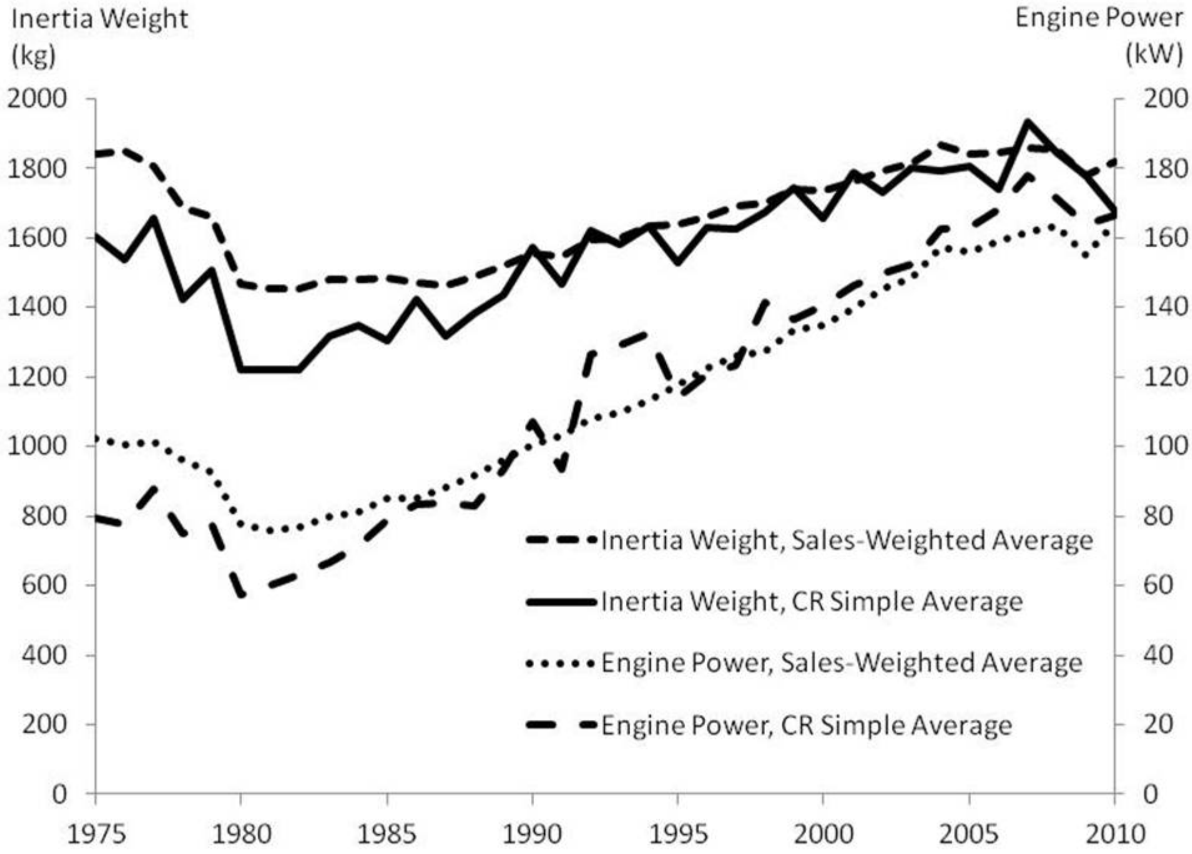

Figure 2-2: Inertia weight and engine power trends since 1975. Shown are simple averages for the vehicles tested by *Consumer Reports* (CR), which were used in fitting the model reported in this paper. Also shown are sales-weighted averages for new U.S. light-duty vehicles as reported by EPA (2010).

In light of its large size, the database was divided into a calibration set and a holdback set. The holdback set was used to evaluate the robustness of different model specifications to changes in the data and to assess the ability of each model specification to make out-of-sample predictions. Twenty percent of the observations were randomly assigned to the holdback set, while the remaining 80 percent of the observations were used to fit the models.

2.4 Results

Multiple model specifications were estimated, but in the interest of brevity, only one specification is reported in detail here. Other specifications distinguished between automatic transmissions and automated manuals (no significant difference), between all- and four-wheel drive (no significant difference), or included a dummy variable for convertibles (convertibles were slightly faster than standard cars from 0–48 and 0–97 km/h). Parameter estimates were generally stable across the different specifications. In all cases, the additional variables improved the model's adjusted r-squared by less than 0.001. A model specification using inertia weight instead of curb weight returned marginally worse (by 0.003–0.006) adjusted r-squared values. All model specifications performed well at predicting acceleration performance in the holdback data set.

Regression results are presented in Table 2.1 and Table 2.2 for one of the investigated specifications. Table 2.1 lists the parameter estimates and standard errors for the engineering and design attributes, including power, weight, and powertrain and body characteristics. Table 2.2 summarizes the estimated fixed effects for year, which represent the expected difference between the log of acceleration time in each year and its value in a base year, all else being equal. The base year was defined as 1977, the first year for which all three acceleration metrics were available.

Table 2.1 includes estimates for the square of the power term, which in all three cases indicate that the sensitivity of acceleration to engine power decreases as power increases. For example, a 1% increase in power in a 75 kW (100 hp) vehicle is expected to produce a 0.70% reduction in 0–97 km/h acceleration time, whereas a 1% increase in power on a 300 kW (400 hp) vehicle is expected to produce just a 0.58% reduction in the 0–97 km/h time. Also investigated, but not reported here, were model specifications that included squared weight terms, and terms for power interacted with weight, displacement, and engine type and transmission type. With the exception of the squared term for power, none of these proved to be significant or to improve model fit. The estimated coefficients for power, weight, displacement, and manual transmissions are similar to those estimated by Santini and Anderson

(1993). The coefficients estimated here suggest that 0–48 km/h acceleration is a little more sensitive to peak power and less sensitive to curb weight and displacement than was indicated by the results of Santini and Anderson. They also indicate that the 0–97 km/h and 72–105 km/h acceleration times are less sensitive to peak power, curb weight, and displacement than indicated by their results.

2.4.1 Effects of Body Style

Light trucks are estimated to deliver marginally slower acceleration than cars at low speeds, and significantly slower acceleration at higher speeds. This is consistent with trucks suffering larger aerodynamic losses due to their higher drag coefficients and larger frontal areas. While the aerodynamic losses may not be important at low speeds, they can become considerably more important at higher speeds.

The estimated coefficients for different body types are smaller than those found by Santini and Anderson (1993). In most cases, but not all, they are directionally the same. Santini and Anderson (1993) estimated the effect of vehicle body type using a dummy for vehicle type interacted with the logarithm of frontal area in m^2, as shown here for vans:

$$BodyStyleEffect = \beta ln(FrontalArea)Van \tag{2.4}$$

The data set used in this work did not include frontal area, so effects were estimated for vehicle type dummy variables without including frontal area:

$$BodyStyleEffect = \beta' Van \tag{2.5}$$

Thus, the coefficient estimates from this work are more appropriately compared with the logarithm of the frontal area from Santini and Anderson's results, appropriately interacted with their dummy variables for body type. For model year 2008, the logarithm of the frontal area in m^2 ranged from 0.8–1.2 for cars, from 1.0–1.6 for pickups and SUVs, and from 1.2–1.8 for vans. Based on these ranges, the combined effects of size and body type from Santini and Anderson's work suggest the following:

- Pickups and SUVs have similar 0–48 km/h acceleration to cars, while vans have slightly faster acceleration times,
- Vans do slightly worse, and SUVs and pickups do considerably worse than cars accelerating from 0–97 km/h, and
- SUVs, vans, and pickups all have considerably worse acceleration than cars from 72–105 km/h.

In contrast, the results of this work suggest that SUVs, pickups, and vans all have somewhat slower acceleration than cars, with the effect less pronounced for SUVs. In addition, the magnitudes of the estimates reported here are smaller than the combined size and body type effects reported by Santini and Anderson.

2.4.2 Effects of Drivetrain Characteristics

The parameter estimates reported in Table 2.1 indicate that a manual transmission delivers approximately 8% faster acceleration from 0–48 km/h and 4–5% faster acceleration from 0–97 km/h and 72–105 km/h than an automatic transmission. Interpreting the coefficients for continuously variable transmissions (CVTs) demands caution, because CVTs were defined as having zero speeds. Thus, the effects for CVTs must be compared against the combined effects of transmission type and number of transmission speeds for automatic or manual transmissions. For example, although the coefficient for CVTs is estimated at -0.185 according to the results for 0–97 km/h acceleration in Table 2.1, the expected 0–97 km/h acceleration time for a vehicle with a CVT would only be about 3% faster than an identical vehicle equipped with a 5-speed automatic transmission:

$$ln(Z97_{CVT}) - ln(Z97_{auto-5}) = \beta_{CVT} - \beta_{TSpd}TSpd = -0.185 - (-0.031) \cdot 5 = -0.03 \tag{2.6}$$

The results indicate that turbocharged and supercharged gasoline vehicles accelerate faster than naturally aspirated engines, all else being equal. However, caution is again required because a boosted engine typically has a smaller displacement than a naturally aspirated engine delivering the same peak power. Because the models investigated here separately control for displacement, boosting and downsizing an engine while maintaining peak power incurs two offsetting effects on predicted acceleration: a decrease in acceleration time due to boosting and an increase due to smaller displacement. For example, consider a vehicle with an engine that has been downsized by 30% and turbocharged so as to maintain the original peak power. Assuming that vehicle weight remains unchanged, the results in Table 2.1 predict 1%, 2%, and 4% net reductions in the acceleration times from 0–48 km/h, 0–97 km/h, and 72–105 km/h, respectively.

The results indicate that hybrid electric powertrains do not deliver significantly different acceleration performance than conventionally powered vehicles with the same peak power. (*Consumer Reports* lists combined system power for hybrids.) Hybrid designations were made by *Consumer Reports*, and this analysis did not distinguish between different types of hybrids. There were 20 hybrids in the data set: 8 Toyota and Lexus, 3 Ford and Mercury, 1 Nissan, 5 Honda, and 3 GM (2007 Vue Greenline, 2009 Malibu, 2008 Tahoe).

Naturally aspirated diesels, which are not found in the data set after 1982, delivered significantly slower acceleration than similar gasoline vehicles. Turbodiesels are estimated to deliver similar to slightly faster performance than conventional gasoline vehicles.

Rear-wheel and front-wheel drive vehicles are estimated to deliver similar acceleration performance, with rear-drive vehicles delivering slightly faster passing acceleration. Four- and all-wheel drive vehicles deliver faster acceleration up to 48 km/h, which may be due to reduced wheel spin in high-powered vehicles. However, this advantage is reversed at higher speeds, consistent with increased driveline losses and the reduced importance of wheel slip as a limiting factor at higher speeds.

2.4.3 Effects of Time

The fixed effects estimated for each year can be interpreted as measuring how much more effectively a vehicle transforms engine power into acceleration for a vehicle of a given mass, controlling for various vehicle characteristics. Knittel (2011) used year fixed effects in a similar way, interpreting them as a measure of vehicle technological improvements in a broad sense. In the current context, these improvements may include better aerodynamics, reduced tire rolling resistance, and increased efficiency of powertrain components downstream of the engine. Such improvements would reduce the engine power devoted to overcoming losses, freeing up more power to accelerate the vehicle. Another possibility is that improvements in tire technology have reduced wheel slip, a hypothesis that is consistent with the result (discussed below) that 0–48 km/h acceleration has shown larger relative improvements than 0–97 km/h and 72–105 km/h acceleration. Regardless of the particular sources of the improvements, it does appear that engineers today can obtain higher performance per unit of power (and weight, etc.) than they could in the past.

For small values, the year fixed effects are approximately equal to a percentage change in acceleration time. For example, a year fixed effect of -0.01 indicates that a vehicle is expected to deliver approximately a 1% faster acceleration time than a comparable vehicle in the base year. For larger values, nonlinearities become significant, but the fixed effects can be transformed into a ratio of expected acceleration time for a vehicle relative to a comparable vehicle in the base year. Assuming all independent variables other than the year fixed effect to be equal in year t and in some base year, and normalizing the fixed effect to be zero in the base year, one can subtract Equation 2.3 for the base year from Equation 2.3 for year t to obtain:

$$ln(ACC_t) - ln(ACC_{Base}) = \beta_{Y,t} - \beta_{Y,Base} = \beta_{Y,t} \tag{2.7}$$

Where $\beta_{Y,t}$ is the fixed effect for year t. Equation 2.7 can be rearranged to yield the ratio of the expected acceleration time for a given vehicle in year t to the expected acceleration time of a similar vehicle in the base year:

$$\frac{ACC_t}{ACC_{Base}} = e^{\beta_{Y,t}} \tag{2.8}$$

The estimated fixed effects of year were approximately the same in all model specifications, but varied between the different measures of acceleration, as shown in Figure 2-3.The figure plots expected acceleration times in each year relative to a base year of 1977, using the fixed effects from Table 2.2. Several features visible in the figure are worthy of attention.

First, the expected acceleration performance of a vehicle, conditioned on engine power, vehicle weight, and various other attributes, is considerably faster today than it was in the 1970s. As discussed previously, this can be interpreted as technological progress that has improved the ability of vehicles to squeeze more useful performance from a given level of power.

Second, the rate of change in acceleration performance, conditioned on other attributes, has not been uniform. Rapid improvements through the 1980s were followed by more gradual changes since 1990. This is generally consistent with findings that the overall rate of technical improvement in U.S. light-duty vehicles was most rapid in the early 1980s and has slowed down in more recent years (Knittel, 2011).

Third, the estimates suggest that 0–97 and 72–105 km/h acceleration may have deteriorated between the late 1970s and early 1980s, while lower-speed acceleration performance (i.e. 0–48 km/h) did not. However, the statistical significance of the year fixed effects through 1985 is marginal at best, so it is difficult to conclude with any confidence that acceleration performance (conditioned on other attributes) actually deteriorated during this time.

Finally, the relative improvements in the 0–48 km/h acceleration have been larger than those in the 0–97 and 72–105 km/h acceleration. It is not possible to say from the available data why this is the case, but several plausible explanations could be investigated if a richer data set were available. First, the difference may be due to improvements in throttle response, which would improve acceleration“off the line.” This would have a larger relative effect on the 0–48 km/h acceleration than on the higher speed acceleration measures. Second, improvements in tire technology that reduce wheel slip would be expected to have a larger effect at lower speeds, where wheel slip is more likely to be a limiting factor. Finally, market forces might have driven a greater emphasis on acceleration performance at lower speeds than at higher speeds. These demands could be met, for example, by altering gear ratios to favor low-speed performance. However, changing the ratios in lower gears would not necessarily improve acceleration at higher speeds, and could even compel tradeoffs that reduce performance at higher speeds.

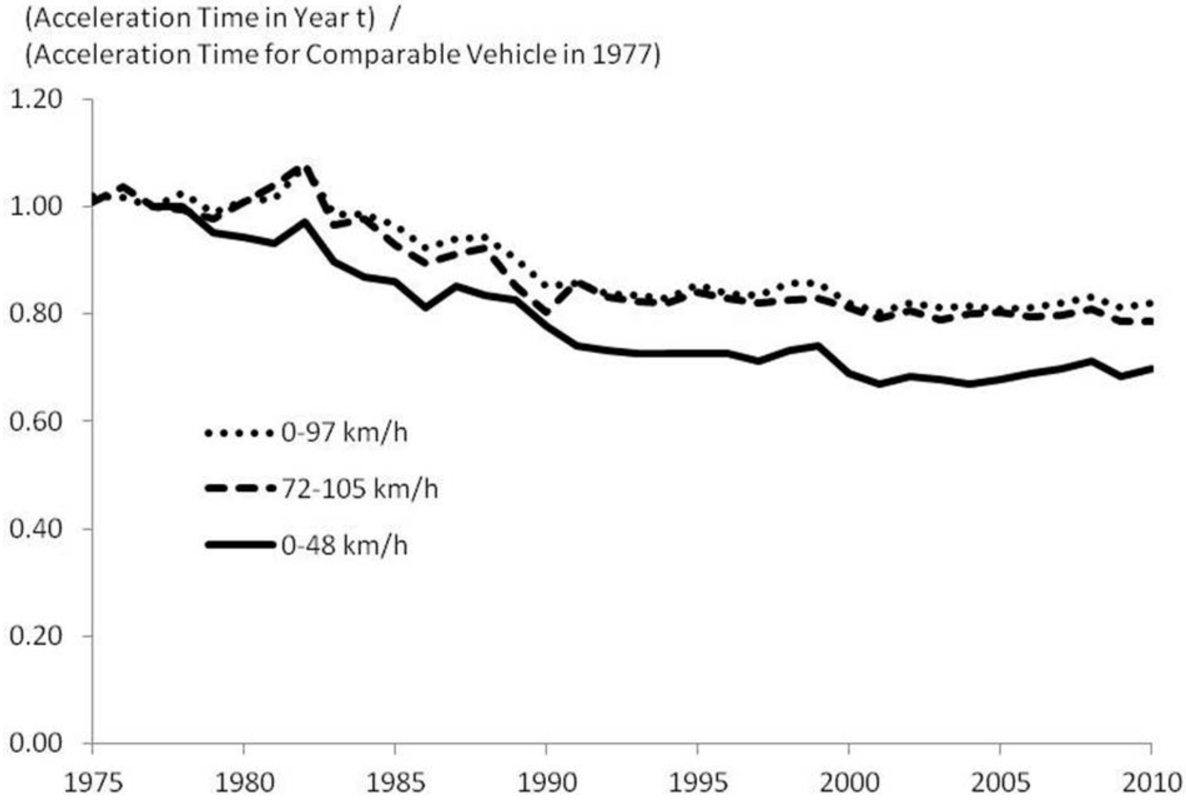

Figure 2-3: Ratio of expected acceleration time for a vehicle in each year to that of a comparable vehicle (i.e. same power, weight, transmission type, etc.) in 1977. Relative reductions in 0–48 km/h acceleration times have been greater than those in 0–97 km/h and 72–105 km/h acceleration times.

2.4.4 Alternative Model Specifications

Practitioners who wish to estimate acceleration performance of vehicles may not have access to all of the variables included in the model specification reported above. For example, the U.S. Environmental Protection Agency reports vehicle weight according to inertia weight class, rather than curb weight, so users of of EPA data may wish to have a model of acceleration performance that relies on inertia weight. Other users may have access only to basic data including weight, power, and model year. For the convenience of such readers, two more sets of models are reported here: one in which inertia weight replaces curb weight in the models reported earlier; and another which relies only on weight, power, and model year. Table 2.3 and Table 2.4 summarize the results of regressions that use inertia weight in place of curb weight. Table 2.5 and Table 2.6 summarize the results of the regressions that use only power, curb weight, model year, and class (car vs. light truck). The simplest models (using only power, weight, and class) should not be used to infer the effects of power and weight on acceleration performance, due to the risk of omitted variable bias. For the same reason, considerable caution is warranted if the simple specification is used to make out-of-sample predictions.

2.5 Applying the Model to New Data Sets

The predictive ability of each model specification was assessed using the holdback data set. The prediction errors were generally similar across different model specifications, and increased slightly as the adjusted r-squared values fell. Applying the models to the holdback data suggested that there were no surprises associated with making predictions from any of the model specifications. Any of the specifications, including those reported here, appear to be appropriate for predicting acceleration performance.

2.5.1 Representativeness of the *Consumer Reports* Sample

The sample of vehicles tested by *Consumer Reports* is not randomly selected from the population of vehicles available on the market, raising concerns about possible bias in the estimates of the regression coefficients, and the applicability of the models to vehicles outside the sample. To address these concerns, propensity scores were estimated for the likelihood of a vehicle being included in *Consumer Reports*' testing program, conditional on a variety of vehicle characteristics. The estimated propensity scores were then incorporated into the regression analyses using two approaches outlined by Schafer and Kang (2008).

Propensity scores were estimated using a logit model of the probability of a vehicle being included in the *Consumer Reports*' testing program. The model included as predictor variables vehicle class, power, weight, powertrain characteristics, manufacturer, model year, and various interactions among these. There was good overlap between the propensity scores in the population and those in the *Consumer Reports* sample.

Next, dummy variables were defined for the deciles of propensity scores, and were included as additional regressors in the models. The dummy variables were generally insignificant, and the estimates of the coefficients and year fixed effects remained essentially unchanged relative to the models without the propensity scores.

In an alternative approach, the regressions were weighted by the inverses of the propensity scores. The rationale for this approach is that it can weight each observation in the sample by the number of vehicles that it represents in the full population. This procedure led to increases in the estimated sensitivity of acceleration times to power and to weight, and to slight changes in the coefficient estimates for powertrain types, body types, and other vehicle characteristics. The weighting did not change the general trends or levels of the year fixed effects, but did increase their volatility from year to year.

Finally, both the weighted and unweighted regression models were applied to the holdback data set. In all cases, the weighted model returned larger average errors than the unweighted model (including when errors were averaged by the inverse propensity score). Based on these findings, it is recommended that the coefficient estimates from the unweighted regression model be used, even for making out-of-sample predictions.

2.5.2 Implications of Results for Estimates of U.S. Vehicle Performance

The results of this work suggest that the acceleration performance of new U.S. vehicles has been improving more quickly than previously thought. Figure 2-4 shows the sales-weighted average 0–97 km/h acceleration times calculated by applying these results to a comprehensive database of vehicle attributes and sales volumes spanning 1978–2009, as well as the average acceleration calculated by applying the more typical method of Malliaris et al. (1976). Also shown are the average acceleration values reported by EPA (2010) for 1975-2010. The latter agree very closely with the results obtained using the method of Malliaris et al. (1976), suggesting that there are only small discrepancies between the database used here and that used by EPA. Between 2006 and 2009, the average acceleration calculated using the Malliaris et al. (1976) method was approximately 1 second, or 11%, greater than the average of 8.8 seconds calculated using the model reported in this work. Between 1982 and 2009, the estimated average 0–97 km/h acceleration time of new U.S. vehicles decreased from 16.6 seconds to 8.8 seconds. Over the same period, the average 0–48 km/h acceleration time decreased from 5.5 seconds to 3.2 seconds, and the average 72–105 km/h passing acceleration time fell from 10.9 seconds to 5.6 seconds.

Reductions in 0–97 km/h acceleration times occurred within high-performance and low-performance vehicles alike. Figure 2-5 shows how 0–97 km/h acceleration times have changed since 1978 for the median vehicle as well as for vehicles at the fastest (5^{th}percentile) and slowest (95^{th}percentile) ends of the market.

Two features of Figure 2-5 are especially striking. First, 95% of vehicles sold today achieve a level of acceleration performance that beats the average from 1992, and would have put them in the top 5% in 1985. To put this more concretely, consider

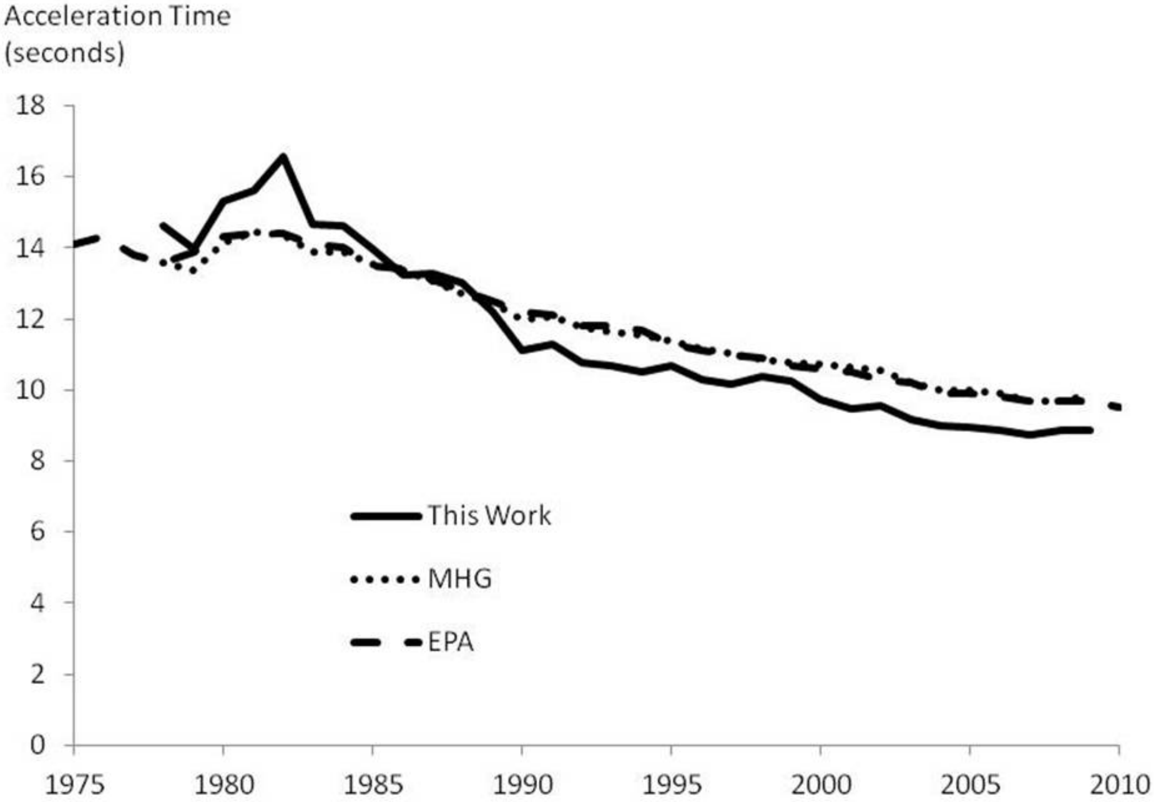

Figure 2-4: Sales-weighted average 0–97 km/h acceleration times calculated by applying method of Malliaris et al. (1976) and the model reportid tn this chapter to vehicle attributes and sales data for 1978–2009. Also shown are the averages reported by EPA (2010) for 1975–2010.

two cars that have been owned by the author: the 1987 Mazda RX–7 and the 2007 Honda Fit. The former is a venerable sports car, which was faster from 0–97 km/h than 94% of its contemporaries. The latter is the archetype of the modern economy car, and is slower from 0–97 km/h than 97% of vehicles sold in 2007. Yet, the actual acceleration times of these vehicles are nearly the same: 10.5 seconds for the Mazda, versus 11.2 seconds for the Honda. [1]

Second, the chart shows that although acceleration times have been getting faster, the rate of change has been declining. In fact, the chart appears to suggest that

---

[1] This is hardly unique. The 1985 RX–7, Nissan 300ZX, and Toyota Supra all had 0–97 km/h times of 11.0 seconds. The 2009 Fit, along with the 2009 Toyota Yaris and 2008 Nissan Versa, all had 0–97 km/h times between 10.9 and 11.1 seconds.

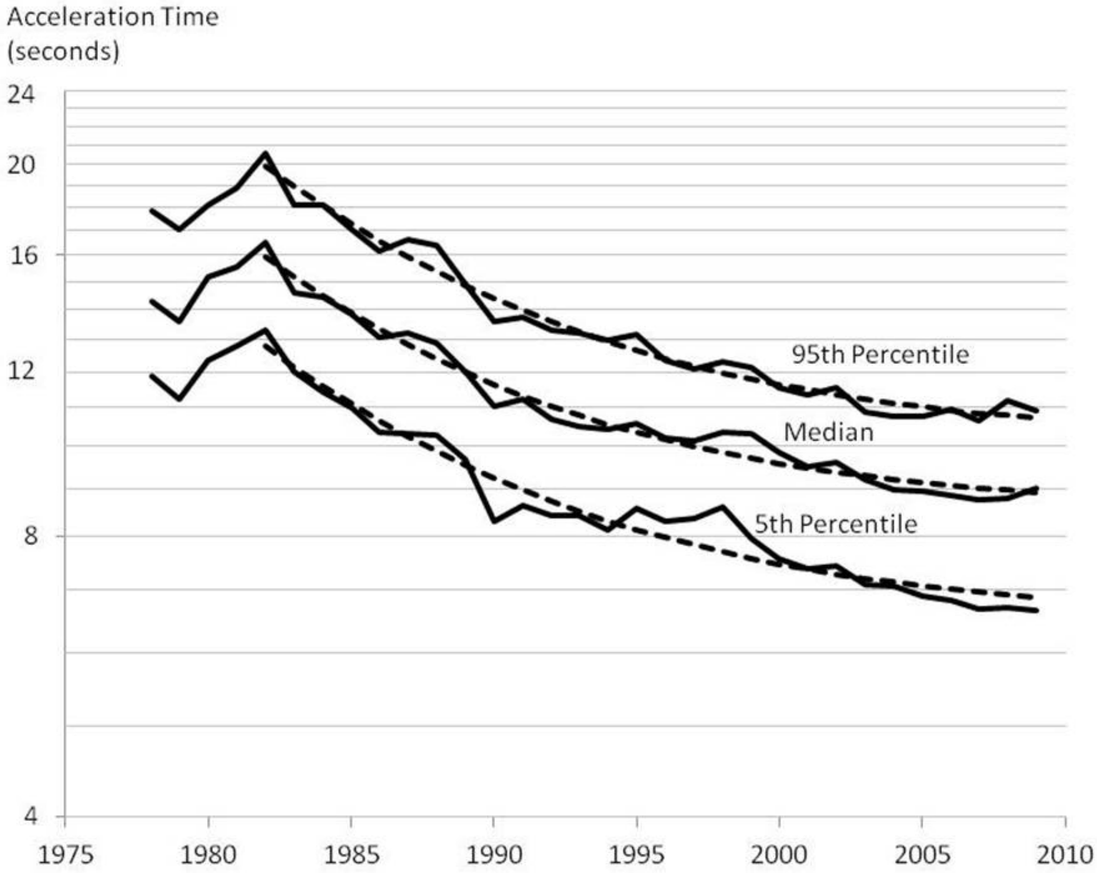

Figure 2-5: Median, 5^{th}percentile, and 95^{th}percentile times for acceleration from 0–97 km/h, as estimated using the model reported here, for 1978–2009. Also shown are curves fitted for the years 1982–2009. Reductions in 0–97 km/h acceleration times have been observed across the whole market, and trends are consistent with decay toward an asymptote.

acceleration performance may be asymptoting. A model of exponential decay toward an asymptote captures both the asymptotic acceleration level and the rate of approach toward that level:

$$Z97_t = ae^{b(t-1980)} + c \tag{2.9}$$

Parameter c in Equation 2.9 represents the estimated asymptotic performance level, while parameter b captures the average rate at which acceleration performance has been approaching this level, and parameter a is a constant. These parameters were estimated using least-squares estimation for the years 1982–2009, and the curves fitted in this manner for the median, 5th percentile, and 95th percentile performance levels have been added to Figure 2-5. The fitted parameters suggested, firstly, that the rate of decay, b, is fairly stable regardless of whether vehicles are high-performance, low-performance, or in the middle of the pack. In addition, the estimated asymptotic performance levels ranged from 6.1 seconds for vehicles in the 5th percentile to 10.1 seconds for vehicles in the 95th percentile. It is interesting to note that even high-performance vehicles are today within one second of their estimated asymptotic values. This is, of course, far from proof that reductions in acceleration times are going to stop any time soon, but it at least suggests that Americans' thirst for power in their cars may in fact be quenchable.

2.6 Conclusions

In this chapter, models were fitted for estimating the acceleration performance of light-duty automobiles, based on testing conducted by *Consumer Reports* between 1975 and 2010. A flexible functional form was adopted and various specifications were estimated, which controlled for vehicle attributes including powertrain characteristics and body type. Power and weight were extremely important in determining acceleration performance, consistent with findings by previous investigators. Other attributes including displacement, powertrain characteristics, and body type have smaller but still significant effects on acceleration performance. Judging by their performance on a holdback data set, the primary model specifications reported here appear to be valid for making predictions of acceleration performance. These estimation methods will be applied in Chapter 4 to support work characterizing the tradeoffs between acceleration performance and fuel consumption, and quantifying improvements in efficiency technology since 1975.

The change in the relationship between acceleration performance and observed attributes was also estimated. All else being equal, new vehicles today achieve approximately 20–30% faster acceleration times than 1970s-vintage vehicles with the same observed attributes. This improvement occurred mainly in the 1980s, and was larger for lower-speed acceleration than for higher-speed acceleration. Prior to 1990, performance improvements were driven both by increases in power and by improvements in how effectively that power is used to accelerate a vehicle. Since 1990, however, most of the performance improvements appear to be attributable to changes in power, weight, and other variables that are tracked directly in the *Consumer Reports* data.

The acceleration performance of vehicles sold in the U.S. today is even faster than is indicated by commonly reported numbers. New vehicles sold in the U.S. between 2006 and 2009 have an estimated average 0–97 km/h acceleration time of 8.8 seconds, about 0.9 seconds faster than the averages reported by U.S. EPA. Interestingly, however, the changes in 0–97 km/h acceleration performance since 1982 fit very well with a model of exponential decay toward an asymptote. This pattern holds up for vehicles from all performance segments, all of which appear at present to be within one second of their respective asymptotic values. This suggests that further work may be warranted to investigate whether consumers' appetites for higher performance are indeed becoming satiated.

Table 2.1: Regression Results for Logs of Acceleration Times. P is in kW, CWT is in kg, D is in Liters, and Continuously Variables Transmissions (CVTs) are Defined as Having Zero Speeds. Significance: + = 0.1 * = 0.05 ** = 0.01 *** = 0.001

Variable	$ln(Z48)$	$ln(Z97)$	$ln(P72105)$
Intercept, β_0	1.236 ***	2.070 ***	2.609 ***
	(0.370)	(0.298)	(0.365)
$ln(P)$	-0.815 ***	-1.088 ***	-1.386 ***
	(0.139)	(0.112)	(0.138)
$[ln(P)]^2$	0.046 **	0.044 ***	0.074 ***
	(0.014)	(0.012)	(0.014)
$ln(CWT)$	0.483 ***	0.665 ***	0.620 ***
	(0.037)	(0.031)	(0.037)
$ln(D)$	-0.171 ***	-0.121 ***	-0.101 ***
	(0.020)	(0.016)	(0.020)
$TSpd$	-0.035 ***	-0.031 ***	-0.022 ***
	(0.005)	(0.004)	(0.005)
Transmission Types			
Manual	-0.081 ***	-0.045 ***	-0.052 ***
	(0.010)	(0.008)	(0.010)
CVT	-0.109 ***	-0.185 ***	-0.168 ***
	(0.032)	(0.026)	(0.032)
Engine Types			
Turbocharged Gasoline	-0.071 ***	-0.066 ***	-0.081 ***
	(0.013)	(0.011)	(0.013)
Supercharged Gasoline	-0.041	-0.059 *	-0.073 *
	(0.034)	(0.028)	(0.034)
Diesel	0.122 ***	0.116 ***	0.158 ***
	(0.027)	(0.022)	(0.027)
Turbodiesel	0.029	-0.029	-0.086 **
	(0.029)	(0.024)	(0.030)
Hybrid Electric	-0.023	-0.023	-0.036
	(0.026)	(0.022)	(0.026)
Drive Types			
Rear-Wheel Drive	-0.011	0.005	0.023 **
	(0.008)	(0.007)	(0.008)
All / 4-Wheel Drive	-0.024 *	0.019 +	0.031 **
	(0.012)	(0.010)	(0.012)
Body Styles			
Wagon	0.029 **	0.019 *	0.022 *
	(0.011)	(0.009)	(0.010)
SUV	0.019	0.019 +	0.038 **
	(0.013)	(0.011)	(0.013)
Van	0.024 +	0.042 ***	0.063 ***
	(0.014)	(0.012)	(0.014)
Pickup	0.026	0.031 *	0.056 ***
	(0.016)	(0.013)	(0.016)
Adjusted R2	0.8782	0.9408	0.9180

Table 2.2: Estimated Fixed Effects of Year on Logs of Acceleration Times, Normalized to Zero in 1977. Note that 0–48 km/h acceleration data were not available for 1975–1976.

Year	$ln(Z48)$	$ln(Z97)$	$ln(P72105)$
1975		0.02	0.00
1976		0.01	0.03
1977	0.00	0.00	0.00
1978	0.00	0.02	-0.01
1979	-0.05	-0.01	-0.02
1980	-0.07	0.00	-0.01
1981	-0.08	0.01	0.02
1982	-0.04	0.07	0.07
1983	-0.11	-0.02	-0.04
1984	-0.14	-0.02	-0.03
1985	-0.14	-0.03	-0.07
1986	-0.21	-0.08	-0.10
1987	-0.16	-0.06	-0.09
1988	-0.18	-0.06	-0.08
1989	-0.18	-0.09	-0.15
1990	-0.24	-0.15	-0.20
1991	-0.30	-0.15	-0.14
1992	-0.30	-0.17	-0.17
1993	-0.31	-0.17	-0.18
1994	-0.31	-0.18	-0.18
1995	-0.31	-0.15	-0.16
1996	-0.31	-0.17	-0.17
1997	-0.33	-0.17	-0.18
1998	-0.30	-0.15	-0.17
1999	-0.30	-0.15	-0.18
2000	-0.36	-0.19	-0.20
2001	-0.39	-0.21	-0.22
2002	-0.37	-0.19	-0.20
2003	-0.38	-0.20	-0.22
2004	-0.39	-0.20	-0.21
2005	-0.38	-0.20	-0.20
2006	-0.37	-0.20	-0.22
2007	-0.35	-0.20	-0.22
2008	-0.34	-0.18	-0.20
2009	-0.37	-0.20	-0.23
2010	-0.36	-0.19	-0.23

Table 2.3: Regression Results for Logs of Acceleration Times. P is in kW, IWT is in kg, D is in Liters, and Continuously Variables Transmissions (CVTs) are Defined as Having Zero Speeds. Significance: + = 0.1 * = 0.05 ** = 0.01 *** = 0.001

Variable	$ln(Z48)$	$ln(Z97)$	$ln(P72105)$
Intercept, β_0	1.513 ***	2.382 ***	2.801 ***
	(0.382)	(0.317)	(0.378)
$ln(P)$	-0.711 ***	-0.958 ***	-1.274 ***
	(0.142)	(0.117)	(0.140)
$[ln(P)]^2$	0.037 *	0.032 **	0.064 ***
	(0.015)	(0.012)	(0.015)
$ln(IWT)$	0.395 ***	0.561 ***	0.540 ***
	(0.038)	(0.032)	(0.037)
$ln(D)$	-0.132 ***	-0.070 ***	-0.059 **
	(0.020)	(0.017)	(0.020)
$TSpd$	-0.033 ***	-0.030 ***	-0.021 ***
	(0.005)	(0.005)	(0.005)
Transmission Types			
Manual	-0.091 ***	-0.056 ***	-0.061 ***
	(0.010)	(0.009)	(0.010)
CVT	-0.106 **	-0.185 ***	-0.170 ***
	(0.033)	(0.028)	(0.033)
Engine Types			
Turbocharged Gasoline	-0.062 ***	-0.055 ***	-0.071 ***
	(0.013)	(0.011)	(0.013)
Supercharged Gasoline	-0.028	-0.041	-0.057
	(0.035)	(0.030)	(0.035)
Diesel	0.129 ***	0.124 ***	0.164 ***
	(0.027)	(0.023)	(0.027)
Turbodiesel	0.069 *	0.024	-0.040
	(0.030)	(0.025)	(0.030)
Hybrid Electric	0.001	0.009	-0.006
	(0.027)	(0.023)	(0.027)
Drive Types			
Rear-Wheel Drive	-0.010	0.007	0.025 **
	(0.008)	(0.007)	(0.008)
All / 4-Wheel Drive	-0.019	0.025 *	0.036 **
	(0.012)	(0.010)	(0.012)
Body Styles			
Wagon	0.034 **	0.025 **	0.028 **
	(0.011)	(0.009)	(0.011)
SUV	0.033 *	0.036 **	0.052 ***
	(0.013)	(0.011)	(0.014)
Van	0.046 **	0.069 ***	0.086 ***
	(0.014)	(0.012)	(0.014)
Pickup	0.036 *	0.043 **	0.065 ***
	(0.016)	(0.014)	(0.016)
Adjusted R^2	0.8727	0.9344	0.9138

Table 2.4: Estimated Fixed Effects of Year on Logs of Acceleration Times, When Using Inertia Weight. Normalized to Zero in 1977. Note that 0–48 km/h acceleration data were not available for 1975–1976.

Year	$ln(Z48)$	$ln(Z97)$	$ln(P72105)$
1975		0.03	0.01
1976		0.02	0.03
1977	0.00	0.00	0.00
1978	-0.01	0.02	-0.01
1979	-0.05	-0.01	-0.02
1980	-0.07	0.01	0.00
1981	-0.08	0.01	0.03
1982	-0.04	0.07	0.08
1983	-0.11	-0.01	-0.04
1984	-0.14	-0.02	-0.03
1985	-0.15	-0.03	-0.07
1986	-0.21	-0.08	-0.11
1987	-0.16	-0.06	-0.09
1988	-0.18	-0.05	-0.07
1989	-0.18	-0.09	-0.14
1990	-0.24	-0.15	-0.20
1991	-0.29	-0.14	-0.13
1992	-0.30	-0.16	-0.17
1993	-0.32	-0.17	-0.18
1994	-0.32	-0.18	-0.19
1995	-0.31	-0.16	-0.16
1996	-0.31	-0.17	-0.17
1997	-0.33	-0.17	-0.18
1998	-0.30	-0.14	-0.17
1999	-0.30	-0.15	-0.17
2000	-0.36	-0.19	-0.20
2001	-0.38	-0.21	-0.21
2002	-0.37	-0.19	-0.20
2003	-0.38	-0.20	-0.22
2004	-0.39	-0.20	-0.21
2005	-0.38	-0.20	-0.20
2006	-0.36	-0.19	-0.21
2007	-0.35	-0.19	-0.21
2008	-0.34	-0.18	-0.20
2009	-0.37	-0.20	-0.23
2010	-0.36	-0.19	-0.23

Table 2.5: Regression Results for Logs of Acceleration Times, Using Power, Weight, Year, and Class Only. P is in kW, CWT is in kg. Significance: + = 0.1 * = 0.05 ** = 0.01 *** = 0.001

Variable	$ln(Z48)$	$ln(Z97)$	$ln(P72105)$
Intercept, β_0	0.364 *	1.222 ***	0.978 ***
	(0.151)	(0.118)	(0.141)
$ln(P)$	-0.579 ***	-0.799 ***	–0.775 ***
	(0.016)	(0.013)	(0.016)
$ln(CWT)$	0.536 ***	0.686 ***	0.645 ***
	(0.027)	(0.021)	(0.026)
Light Truck	-0.013	0.026 ***	0.054 **
	(0.009)	(0.007)	(0.009)
Adjusted R^2	0.8442	0.9286	0.9036

Table 2.6: Estimated Fixed Effects of Year on Logs of Acceleration Times, When Using Power, Weight, Year, and Class Only. Normalized to Zero in 1977. Note that 0–48 km/h acceleration data were not available for 1975–1976.

Year	$ln(Z48)$	$ln(Z97)$	$ln(P72105)$
1975		0.00	-0.02
1976		0.00	0.02
1977	0.00	0.00	0.00
1978	0.01	0.03	0.01
1979	-0.04	0.00	0.00
1980	-0.07	0.01	0.02
1981	-0.06	0.01	0.04
1982	-0.01	0.08	0.09
1983	-0.13	-0.04	-0.06
1984	-0.08	0.01	0.00
1985	-0.14	-0.04	-0.09
1986	-0.16	-0.06	-0.10
1987	-0.10	-0.04	-0.09
1988	-0.10	-0.02	-0.05
1989	-0.15	-0.09	-0.16
1990	-0.18	-0.13	-0.20
1991	-0.24	-0.13	-0.14
1992	-0.23	-0.13	-0.16
1993	-0.23	-0.14	-0.17
1994	-0.25	-0.15	-0.18
1995	-0.21	-0.11	-0.14
1996	-0.23	-0.13	-0.16
1997	-0.24	-0.13	-0.16
1998	-0.21	-0.11	-0.16
1999	-0.19	-0.10	-0.15
2000	-0.26	-0.14	-0.18
2001	-0.28	-0.16	-0.19
2002	-0.28	-0.15	-0.18
2003	-0.28	-0.16	-0.20
2004	-0.30	-0.16	-0.20
2005	-0.29	-0.17	-0.20
2006	-0.28	-0.17	-0.21
2007	-0.25	-0.16	-0.20
2008	-0.24	-0.15	-0.20
2009	-0.27	-0.17	-0.22
2010	-0.27	-0.18	-0.24

Chapter 3

Analyzing Weight Trends in U.S. Passenger Cars

This chapter is based on a paper jointly authored with Stephen Zoepf and John Heywood. Reproduced by permission from a paper accepted for publication in a forthcoming issue of the International Journal of Vehicle Design ©Inderscience Publishers.

After a precipitous drop from 1976–1982, the weight of U.S. passenger cars has grown steadily. This chapter examines multiple conflicting influences on vehicle weight in two categories: technological changes that reduce vehicle weight; and improvements in functionality that, *ceteris paribus*, add to vehicle weight. The widespread adoption of unibody construction, lightweight materials, and smaller engines have been offset by growth in vehicle size and feature content. The best estimates from the work reported here indicate that new features and functionality would have added at least 250 kg (550 lbs) to the weight of the average new car between 1975 and 2009, if not for offsetting improvements in technology. Over the same period, it is estimated that alternative materials, more weight-efficient vehicle architectures, and reduced engine sizes have taken 790 kg (1,700 lbs) out of the weight of the average car. These observable influences do not explain the full extent of the drop and subsequent growth in weight, suggesting that substantial non-observed technological improvements were made from 1976–1982, and that unobserved improvements in areas such as crashworthiness and NVH have added substantially to vehicle weight in the past two decades.

3.1 Introduction

Between the late 1970s and the early 1980s, something remarkable happened to the cars sold in the United States: they shed more than 450 kg (1000 lbs) from their average curb weight. Even more remarkably, this 25% weight reduction occurred in only six years (and most of it in just four). The first goal of this work was to take a retrospective look at the specific changes that permitted the dramatic weight reductions that occurred between 1976 and 1982, decomposing the overall weight change into constituent changes in technology and vehicle functionality.

Today, after nearly a decade of rising fuel prices and with Corporate Average Fuel Economy (CAFE) standards poised to jump sharply, the auto industry is once again looking to cut dramatically the weight of new vehicles. Ford has stated a goal of cutting 340 kg (750 lbs) off its vehicles by 2020, and announced plans to reduce the weight of the F-150 pickup by 320 kg (700 lbs) in its 2014 redesign. Renault and PSA Peugeot Citroen have established a goal of cutting 200 kg (440 lbs) by 2018, while Hyundai planned in 2010 to cut its average vehicle weight by 10% (150 kg or 330 lbs) over five years. These reductions, on the order of 30–40 kg per year, amount to approximately 2–3% of the initial vehicle weight being removed each year. Previous assessments have suggested that plausible targets for weight reduction are on the order of 20% over 25 years, or 30% after accounting for secondary weight savings Cheah (2010). This amounts to about 1.2% of base vehicle weight reduced each year, or about 15–25 kg per year. A second objective of this work, therefore, was to judge how aggressive the 2–3% per year weight reduction goals are, in the context of what has been achieved historically.

Following its rapid decline, the average weight of a new U.S. car remained essentially unchanged until the late 1980s, at which point it began to rise steadily, a trend that persisted until quite recently. A final objective of this work was to examine why weight increases resumed despite continued growth in the use of weight-saving technologies.

The general patterns described above were observed for both cars and light trucks.

Figure 3-1: Average weights of new U.S. vehicles since 1975 (EPA, 2012)

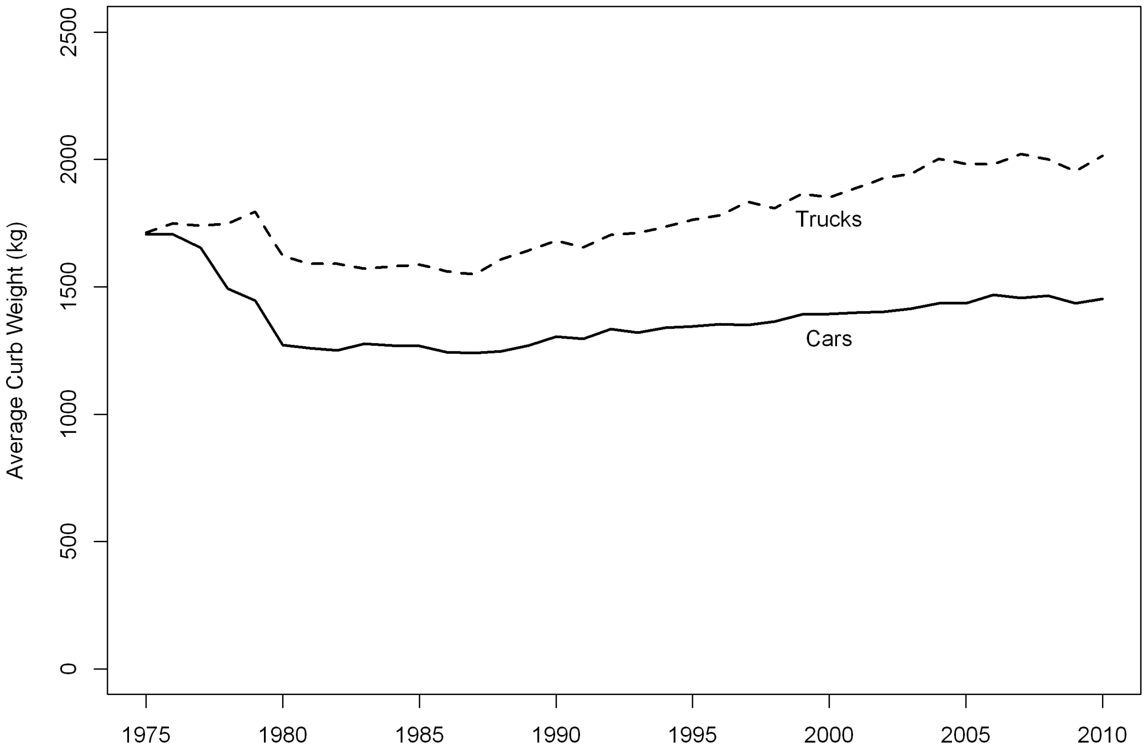

However, cars showed a more dramatic weight decrease in the early years, while trucks underwent a steeper increase in the later years. The focus of this work is on cars exclusively rather than light trucks, for several reasons. First, cars showed a more dramatic decline in weight than did trucks in the late 1970s – early 1980s, so they provide a more interesting subject for studying how rapid weight reductions can be realized. Second, some key data sources were more readily available for cars than for trucks. Finally, cars have remained fairly well defined as a class of vehicles over time, whereas the types of roles fulfilled by trucks have changed considerably over the last 35 years, in particular with the advent of the minivan, SUV, and crossover.

All else equal, a vehicle with a lower weight (and commensurate reductions in power to maintain constant performance) will consume less energy per kilometer traveled. The magnitude of this reduction is significant: Cheah (2010) finds that a 10% reduction in weight results in a 6–7% reduction in fuel consumption. Changes in curb weight are brought about by two fundamentally opposing forces. Improvements in vehicle capability, such as higher performance, larger size or carrying capacity, and greater levels of equipment add weight to a vehicle. Advances in materials, design and manufacturing technologies remove weight from a vehicle. Therefore manufacturers must carefully balance content added to vehicles against investments in weight-saving technology during the course of product development.

This analysis incorporates estimates of the above sources of weight increase and decrease that are observable (vehicle size, architecture, features, engine size, material content) using a bottom-up analysis that estimated the contributions to changes in average weight of various technologies, features, and vehicle size classes. First, the effects of weight-saving technologies on the weight of an individual car were estimated from model-level data and from the literature. Second, the weight-increasing effects of adding features and changing car classes were estimated. Finally, these individual-level estimates were aggregated into overall effects on average vehicle weight based on the changes in market penetration of the various technologies and features.

3.2 Weight reduction: Technological improvements

Weight-reducing effects were estimated for both major architectural changes and for the incremental replacement of traditional steel and iron with lighter and stronger alternative materials. A broader definition of weight-reducing technology would include myriad other advances in engineering, design, and manufacturing practices that permit materials to be used more effectively in building vehicles, but these sorts of technologies were not represented in the data used in this work.

3.2.1 Major architectural changes

New cars in the United States underwent large architectural shifts between 1975 and 1990 that contributed substantially to reductions in weight. In 1975, about half the cars on the market in the U.S. used unitized body (unibody) construction, and fewer than one in 10 were front-wheel drive. By 1990, 85% were front-wheel drive and 95% used unibody construction. (EPA, 2012)

Unibody construction Unibody construction reduces weight by eliminating the traditional frame and integrating its structural functions into the vehicle's body shell. Data compiled by Audatex North America indicate that the overwhelming majority of cars offered in the U.S. since 1975 have used either unibody or body-on-frame construction. In addition, a small number of cars have used space frame construction, which employs a 3-dimensional structure of welded tubes to which non-structural body panels are attached, primarily in low production high performance cars. A few others have used a mixed, unibody-on-frame construction that incorporates elements of both unibody and body-on-frame construction.

Estimates of the weight savings from unibody construction vary widely. Dupnick and Graham (1996) suggested a weight difference of more than 450 kg (1000 lbs) between unibody and body-on-frame cars, while a 1970s case study from Ford attributed only 87 kg (192 lbs) of weight reduction to the switch from body-on-frame to unibody (Gutherie, 1978).

In this work, the weight reduction due to switching from body-on-frame to unibody construction was estimated by creating matched sets of unibody cars and comparable body-on-frame cars using Mahalanobis matching. Size, transmission, drive, and model year data were obtained from a database maintained by the U.S. Environmental Protection Agency (U.S. EPA). Data on construction type by model and year were provided by Audatex North America, and were merged with the EPA database. The matched sets of vehicles were created by matching unibody cars with body-on-frame cars that had the same transmission type and drive type, similar interior volume (within 5 cubic feet or 0.14m^3), and were of similar vintage (within 2 model years).

The two matched groups had identical fractions of transmissions and drive types, and were very well balanced on model year and interior volume. The difference between these groups indicated that the average unibody car weighs 280 kg (616 lbs) less than a body-on-frame car with the same drive type, transmission type, and size, from the same model year, would weigh. A similar analysis indicated that the average space frame car weighs 156 kg (344 lbs) less than a comparable unibody car would, and that cars using unibody-on-frame construction do not differ significantly in weight from comparable unibody cars. These results are summarized in Table 3.1.

A potential source of bias in these estimates is the possibility that vehicles employing more advanced construction techniques may also tend to make greater use of alternative materials. Such as bias would lead to inflated estimates of the weight savings associated with advanced construction techniques, but data on materials composition were not sufficient to support an exploration of this possibility.

Table 3.1: Estimated weight effects of different construction types.

Comparison	Relevant Group	Estimated Difference (kg)	Standard Error (kg)
Construction			
Unibody vs. Body-on-Frame	Unibody cars	-280	5
Space Frame vs. Unibody	Space frame cars	-156	19
Unibody-on-Frame vs. Unibody	Unibody-on-Frame cars	-39	35
Drive			
FWD vs RWD	FWD cars	-296	6

Front-wheel drive A second major architectural change in U.S. cars is the transition from rear-wheel drive to front-wheel drive. Compared with rear-wheel drive, front-wheel drive yields both a direct weight reduction in the drivetrain, and an indirect weight reduction due to improved packaging of the drivetrain. Eliminating the need for a tunnel running the length of the vehicle allows for greater interior space, and so the exterior dimensions of the vehicle (and thus weight) can be reduced while maintaining interior space.

The weight effect of front-wheel drive relative to rear-wheel drive was estimated by matching front-wheel drive vehicles with rear-wheel drive vehicles that had the same transmission type and construction type, similar interior volume (within 5 cubic feet or 0.14m^3), and were of similar vintage (within 2 model years). The matched groups had identical fractions of different construction types and transmission types, and were very similar in average interior volume and model year. Based on the different between these groups, it was estimated that a front-wheel drive car weighs an average of 296 kg (653 lbs) less (standard error: 6 kg) than a rear-drive vehicle with the same transmission type, construction type, interior volume, and model year.

Engine cylinder counts During the period of this analysis engine technology has matured in numerous ways, allowing manufacturers to extract more performance from a given engine mass. These changes are partly reflected in materials analysis discussed later: newer engines more commonly use aluminum blocks and heads than in 1975, and ancillary equipment (intake manifolds and accessories) are increasingly manufactured from composite materials.

This analysis attempts to quantify the improved power density, less materials changes, by assessing the mass impact of changes in engine cylinder counts as a proxy for materials-neutral power density. Nearly all production cars during this timeframe use 4, 6 or 8 cylinder engines—balanced configurations for which most production and reliability issues have been resolved. A few vehicles use 3, 5, 10, 12 and 16 cylinder engines in limited production.

Using a matching process that holds constant for vehicle model, model year, body style and transmission type, this analysis identified an average decrease in weight of 64 kg (142 lbs) by decreasing from 8 to 6 cylinders, and an average decrease of 67 kg (147 lbs) by decreasing from 6 to 4 cylinders.

3.2.2 Alternative (lighter / stronger) materials

Traditional low-carbon steel and iron now make up less than half the weight of a new vehicle, as they are increasingly displaced by alternatives such as high-strength steel, aluminum, plastics, and magnesium. Because the substitution of alternative materials into a vehicle's design is strongly dependent on the demands of the specific application in question, estimating a single figure for the amount of weight saved by these materials is difficult. Nevertheless, it is helpful to generate some rough approximations based on the properties of different materials and reports in the literature. Cheah (2010) and Wohlecker et al. (2006) provide relationships for estimating the weight ratios of parts made with alternative materials to those made with conventional materials, for a variety of generic load cases. These provide a useful starting point for estimating the weight reduction potential of various alternative materials. In addition, rules of thumb and case studies of vehicle designs using alternative materials have been reported by a variety of authors. The weight-saving characteristics assumed for key materials are summarized in Table 3.2 and more detailed discussions of the assumptions for each material are found in the following sections.

Table 3.2: Assumed weight savings for alternative materials.

Material	**Relative Weight Savings**[1]	**Weight Reduction Potential, P**[2]
Conventional Steel	0%	1.0
Iron	0%	1.0
High-Strength Steel	23%	1.3
Aluminum	45%	1.8
Magnesium	60%	2.5
Plastics & Composites	50%	2.0

High-strength steel Rule-of-thumb relationships like those mentioned above, when used with typical values for materials properties, indicate that parts made from high-strength steel (HSS) can be expected to weigh between 0 and 25% less than a conventional steel part, depending on the application. Salonitis et al. (2009) estimated a 10–30% weight reduction from using advanced high-strength steels, and Roth et al. (1998) reported an advanced steel unibody weighing 25% less than conventional unibodies. Das et al. (1997) assumed that high-strength steels could reduce weight by 50% relative to conventional steels, but the rationale for this high value was unclear. A particular challenge in estimating a weight reduction potential for high-strength steel is that there is such a broad range of grades available, with widely varying properties. For purposes of this work, with its focus on assessing weight reduction to date, it was assumed that each kg of high-strength steel replaced 1.3 kg of conventional steel (a 23% weight reduction).

Aluminum Rules of thumb based on generic load cases suggest that substituting aluminum for conventional steel can reduce weight by up to 70%, with a 50% reduction predicted in many applications. The trade press has noted that the greatest concentration of automotive aluminum use is in engines, and that aluminum engine blocks weigh half as much as iron blocks (Murphy, 2006). Stodolsky et al. (1995) estimated that in engine applications, aluminum reduced cylinder head weight by 50% and block weight by 40%. They also reviewed a number of studies and concluded that substituting aluminum for steel in the body reduces weight by about 40–47%, even when “the design of the vehicle is not completely optimized for aluminum manufacture.” Mayer et al. (1994) concluded that a 45% reduction in weight for the body-in-white was possible by substituting aluminum for steel in a BMW 3-series. Das et al. (1997) assumed that substituting aluminum for steel and cast iron delivers a 45% weight reduction, while Carle and Blount (1999) estimated a 40% reduction in weight relative to steel in automotive body applications. Although generic load cases suggest that replacing steel with aluminum can reduce weight by as much as 70%, most of the (considerable) literature on the topic suggests that a value of around 45% is more realistic(i.e. that each kg of aluminum can replace 1.8 kg of conventional iron or steel). The latter was the value assumed in this work.

Magnesium Magnesium still represents a very small fraction (0.3% in 2009) of automotive materials usage, and fewer estimates of its weight reduction potential have been reported. Based on generic load cases, it is estimated that magnesium can reduce weight by up to 70% compared with conventional steel or iron. Luo (2002) calculated savings as high as 80% for some wrought magnesium alloys. Das et al. (1997) assumed that substituting magnesium for steel and cast iron would deliver a 67% weight reduction. In this work, it was assumed that each kg of magnesium replaced 2.5 kg of conventional steel or iron (a 60% reduction).

---

[1] Fraction of weight saved by replacing conventional steel or iron with alternative material

[2] kg conventional material displaced per kg alternative material

Plastics & composites Estimating weight reduction potential for plastics and composites is particularly difficult because of the wide range of materials included in this category. However, some rough calculations with typical ranges of values for materials properties indicate that weight reductions in excess of 80% could be possible, relative to conventional steel or iron. For example, Luo (2002) estimated a weight reduction potential of 35–70% for polycarbonate/ABS based on generic load cases. Das et al. (1997) assume a 30–60% weight reduction from substituting composites for steel. The American Chemistry Council (ACC, 2011) has estimated that each kg of plastics and composites replaces 2–3 kg of other materials (a 50–67% reduction). A report commissioned by Plastics Europe (Pilz et al., 2005) concluded that each kg of plastic replaces an average of 1.5 kg of heavier material (a 33% reduction in weight), but found reductions of up to 75% in some components. In the present analysis, it was assumed that each kg of plastic or composite has displaced 2 kg of traditional steel or iron (a 50% weight reduction).

3.2.3 Other technological improvements

In addition to major architectural changes and increased use of alternative materials, several other sources of technological improvements may have contributed to weight reductions in new cars. Improvements in manufacturing processes and technologies, tighter tolerances, and the like may enable the production of parts that are more precisely designed to use material only where it is needed.

Concurrently, improvements in engineering methods and the availability of tools such as finite element analysis and computer-aided engineering may have enabled the design of parts and systems that make more effective use of materials, better optimizing component designs and interactions. There is clearly also the possibility of synergies between these two sets of improvements. The contributions of these tools to overall weight reduction is hard measure directly, but we will return to them later as a possible means to explain gaps between observed weight changes and those expected from materials substitution and architectural changes.

3.3 Weight growth: Functionality improvements

While the use of weight-saving technologies has steadily grown, it has been offset (and at times, more than offset) by increases in the deployment of weight-increasing features and a shift toward heavier (larger) car classes. In this section, the weight differences between various car classes are summarized, as are the subsystem weights associated with a variety of emissions, safety, and comfort & convenience features.

3.3.1 Weights by car class

In this work, the effects on weight of shifts in the vehicle size mix were estimated using data found in the U.S. EPA's annual Fuel Economy Trends Report. Figure 3-2 shows the average curb weight (average inertia weight minus 136 kg or 300 lbs) for each of six classes of cars tracked by EPA (weights for large, midsize, and small wagons tracked closely with those of the corresponding sedan classes, and are omitted from the figure for clarity). Several features visible in the figure are worth highlighting. First, weight decreases in the early years were much greater among the larger car classes than among the smaller classes. Second, since about 1980, the weights of compacts, subcompacts, and two seaters have been quite similar. Finally, the overall spread between classes has been shrinking over time, first due to weight decreases in the larger classes and more recently driven by weight increases in the smaller classes.

3.3.2 Feature content

One of the most readily observable changes to vehicles during the past four decades has been the widespread addition new features. For the purposes of this analysis, features included not only optional equipment (e.g. air conditioning) but also safety and emissions equipment such as airbags and catalytic converters that are required by regulation in new passenger cars.

This analysis is unable to capture all improvements in these areas—noise, vibration, and harshness (NVH), for example, has dramatically improved in new vehicles as a result of balance shafts, sound insulating materials and active noise cancelation.

Figure 3-2: Average weights of new U.S. cars, by class. (EPA, 2012)

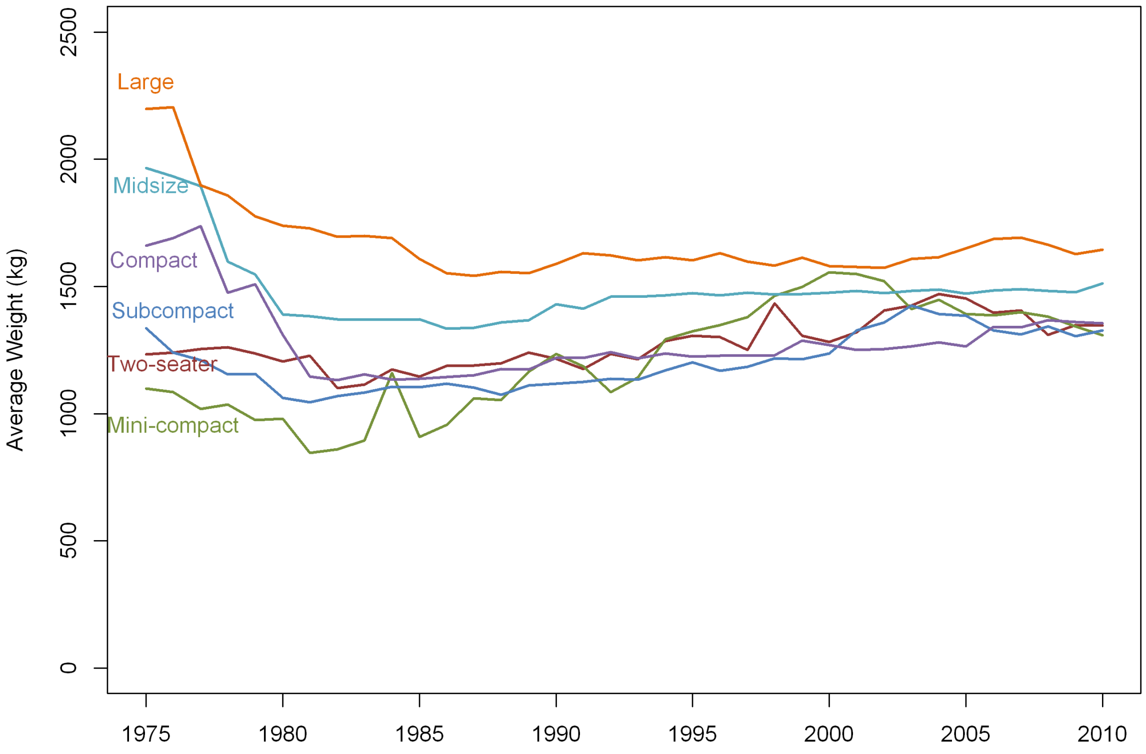

Other metrics such as reliability and body rigidity have also improved. This analysis does not attempt to quantify such metrics.

In total, features were estimated to add a total of 109 kg (240 lbs) to the average 1975 passenger car. In 2010, the estimated contribution grows to 223 kg (62 kg safety, 25 kg emissions, 136 kg comfort/convenience; a total of 491 lbs). These estimates do not include the contribution of secondary weight, discussed in the following section.

3.4 Secondary weight

Secondary weight represents the notion that for every unit of weight added to (or removed from) a vehicle, the supporting systems and structures also grow (or shrink) so that structural integrity and braking, acceleration and handling performance can be maintained. Typically, the secondary weight is expressed as some percentage an initial (primary) weight change. In the work reported here, secondary weight enters as the form of a percentage multiplier on bottom-up component analyses. In reality, the addition or removal of secondary weight may be discontinuous, as in the case where a discrete number of existing engines or transmissions are available for inclusion in a particular vehicle model. Moreover, secondary weight effects may vary depending on the subsystem in which the primary weight reduction occurs. Nevertheless, it is not uncommon to use a single secondary weight multiplier, and this approach is believed to be adequate for the type of analysis presented here, which relies mainly on fleet-wide data.

Cheah (2010) reviewed more than twenty published studies of secondary weight and identified estimates ranging from 23–129%, with a mean value of 79.6%. For the purposes of this study the secondary weight was assumed to be 80% of the primary weight added or removed. This secondary weight coefficient was applied to the bottom-up analyses of features and materials, in which the initial estimates of weight change were generated from component-level data. However, the secondary weight multiplier was not applied for mix shifting or architectural changes, since the weight effects of these changes had already been assessed at the whole-vehicle level.

3.5 Estimating fleet-level contributions of technologies and functionality

The general approach to estimating the effects of technologies and functionalities on average vehicle weight is captured in the question, "What would be the average weight of vehicles from some base year, if those vehicles instead had some future year's combination of technologies and features?" The effects of changes in size mix, features, major architectural technologies, and alternative materials were estimated sequentially and added up as outlined in this section. In all cases, the changes in size mix, features, and technologies are measured relative to a base year of 1975.

3.5.1 Mix shifting

The first step in estimating weight changes was to account for shifts in the mix of size classes between 1975 and each future year t up to and including 2009. Nine car classes based on body style and interior volume are tracked by U.S. EPA and were used as the basis of the mix shifting analysis in this work. For each year t, a weight value was calculated to represent the average weight of cars in that year, if the average weight within each class had been the same as in 1975 but the market share of each class were the same as the actual share in the year t.

$$W_t^s = \sum_i S_{it} \cdot W_{i,1975}$$

In the above expression, $W_{i,1975}$ is the average weight of cars in class i in 1975, S_{it} is the market share of class i in year t, and W_t^s is interpreted as the average weight in year t, adjusted only for changes in shares of various car classes since 1975.

The deployment rates of other major technologies considered in this analysis also tend to be correlated with car class. As a result, some changes in the market share of front-wheel drive, unibody construction, and different engine sizes would be expected to occur due to mix shifting. Therefore, the market share expected based on the change in mix was estimated as follows:

$$S_{jt}^s = \sum_i S_{it} \cdot S_{ij,1975}$$

In the above expression, S_{it} is the market share of class i in year t, $S_{ij,1975}$ is the market share of technology j within class i in 1975, and S_{jt}^s is interpreted as the market share of technology j in year t, adjusted only for the changes in class mix between 1975 and year t.

3.5.2 Architectural changes

After accounting for changes in weight and major technologies due to mix shifting, the next step was to estimate the aggregate weight effects of changes in major vehicle architectures. To do this, the per-vehicle weight effect of each architectural change (δ_j) was multiplied by the difference in market share of that architecture between the year t and 1975 (where the latter share has first been adjusted for mix shifting between 1975 and year t, as noted in the preceding section). The values obtained were then summed over all architectural changes:

$$W_t^{s,a} = W_t^s + \sum_j (S_{jt} - S_{jt}^s) \cdot \delta_j$$

The above expression, $W_t^{s,a}$ is interpreted as the average weight in year t, after adjusting for changes in both mix and the prevalence of major architectures between 1975 and year t. Specifically, this approach was used to estimate the weight effects of the front-wheel drive and unibody transitions, as well as shifts in the prevalence of 4-, 6-, and 8-cylinder engines. As discussed in section 3.4, the weight differences associated with these architectural changes were estimated using vehicle-level data, and do not require further adjustment to account for secondary weight effects.

3.5.3 Safety, comfort, & convenience features

The next step was to estimate the weight effects of growth in the adoption of safety, comfort, and convenience features. Since the weights of the features were estimated based on teardown data for the associated components or subsystems, the secondary weight multiplier was included to estimate the effect of these features on overall vehicle weight:

$$W_t^{s,a,f} = W_t^{s,a} + (1+\sigma) \sum_k (S_{kt} - S_{k,1975}) \cdot \delta_k$$

In the above expression, S_{kt} is the take rate of feature k in year t, $S_{k,1975}$ is the take rate of the feature in 1975, δ_k is the weight associated with the feature based on teardown analysis, σ is the secondary weight factor, and $W_t^{s,a,f}$ is interpreted as the average weight in year t, based on adjusting average weight in 1975 for changes in mix, major architectural changes, and features.

The component-level masses used in these analyses of feature mass are limited in their ability to embody technological improvements in the features themselves. The study of safety features referenced, (DOT, 2004) incorporates at best two point values for the mass of a given feature, while comfort and convenience feature weights are based on a single teardown of four MY2009 production vehicles.

3.5.4 Alternative materials

The final step in estimating weight changes was to incorporate the effects of changes in alternative materials usage between 1975 and each future year. This was done by answering the question: "By how much would the weight of the average car change if the materials mix from 1975 were replaced with the materials mix from some future year t?" This question was addressed in several steps, premised on an assumption that all materials substitution occurs among the materials in set C, which comprises conventional steel and iron, high-strength steel (HSS), aluminum, magnesium, and plastics.

The first step was to ask, "What would be the weight of iron and steel in a car if all of the steel, iron, high-strength steel (HSS), aluminum, magnesium, and plastics in it were replaced by only conventional steel and iron?" This estimate was made using the following expression:

$$W_t^{conv\ only} = W_t^{s,a,f} \sum_{\mathrm{l \in C}} f_{l,1975} \cdot P_l$$

In the above expression, $W_t^{m,a,f}$ is the average weight adjusted for changes in size mix, major architectures, and features, as explained in the preceding sections; $f_{l,1975}$ is the weight fraction of material l in 1975; P_l is that material's weight reduction potential (i.e. the weight of conventional material replaced per unit weight of material used, equal to 1 for conventional steel and iron and greater than 1 for weight-saving alternative materials); and set C includes those materials assumed to be subject to substitution. Other materials, such as glass, zinc, and fluids, were explicitly excluded at this stage.

The second step was to ask, "How would the resulting weight estimate change if the conventional steel and iron were replaced with HSS, aluminum, magnesium, and plastics in the proportions used in year t?" The new weight using these proportions of alternative materials was estimated as:

$$W_t^{adv} = W_t^{conv\ only} \frac{\sum_{\mathrm{l} \in \mathrm{C}} f_{lt}}{\sum_{\mathrm{l} \in \mathrm{C}} (f_{lt} \cdot P_l)}$$

Third, the weight of miscellaneous materials (the set D, including glass, zinc, fluids, etc), was assumed to be unaffected by the materials substitution, and was added to the estimated weight of steel, iron, HSS, aluminum, magnesium and plastics:

$$W_t^{s,a,f,m} = W_t^{adv} + W_t^{s,a,f} \sum_{\mathrm{l} \in \mathrm{D}} f_{lt}$$

Finally, the weight was adjusted for the secondary weight effects resulting from the materials substitution:

$$W_t = W_t^{s,a,f,m} + (W_t^{s,a,f,m} - W_t^{s,a,f})(1 + \sigma)$$

The above weight, W_t, constitutes an estimate of the average car weight in year t, obtained by starting from the average weight in 1975 and adjusting for changes in size mix, major architectures, feature content, and materials composition between 1975 and year t. The logic for estimating the weight effects of alternative materials is summarized in Figure 3-3.

3.6 Examining weight trends since 1975

By combining all estimates of weight-reducing technologies and weight-increasing functionality improvements as shown in Figure 3-3, an estimate was developed for

Figure 3-3: Representation of modeling logic for weight-reducing materials.

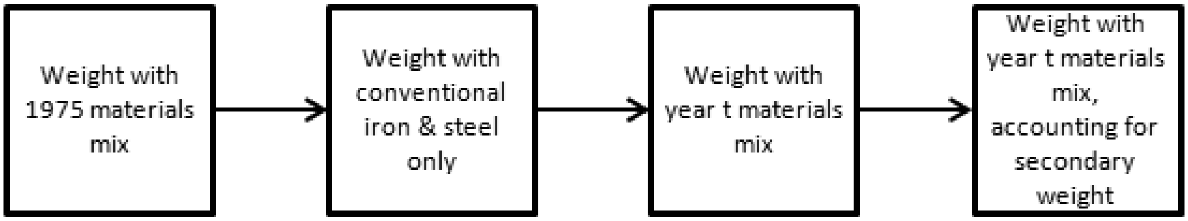

the net change in weight year over year. The results of these individual contributions are reported here, and in section 3.7 they are aggregated up and compared with actual net changes in average passenger car weight.

3.6.1 Sources of weight increase

All else equal, adding functionality to a vehicle in the form of greater size or feature content will generally tend to increase weight. As will be shown in this section, changes in the size mix have added approximately 50–100 kg (110–220 lbs) to the average new car in recent years. The addition of new safety, comfort, and convenience features (and associated secondary weight increases) has added approximately 200 kg (440 lbs) to the average new car since 1975.

Figure 3-4 shows the mix of car classes, weighted by sales, reported by EPA (2012). In the late 1970s, there were noticeable shifts from the compact segment to subcompacts, and to a lesser degree from large sedans to midsize. Over the subsequent years, these trends were gradually reversed. More recently, subcompacts have once again regained share from compacts, as have midsize cars from large cars.

Figure 3-5 shows the estimated weight effects of safety, emissions, and comfort & convenience features for the average new car since 1975, including secondary weight effects. Comfort and convenience features have added more weight than either safety or emissions features, with the most significant contributions coming from near-universal application of air conditioning and automatic transmissions. Greater detail, including the market shares and weight effects of each specific feature considered, can be found in Zoepf (2011).

Figure 3-4: Historical sales mix of passenger cars by class.

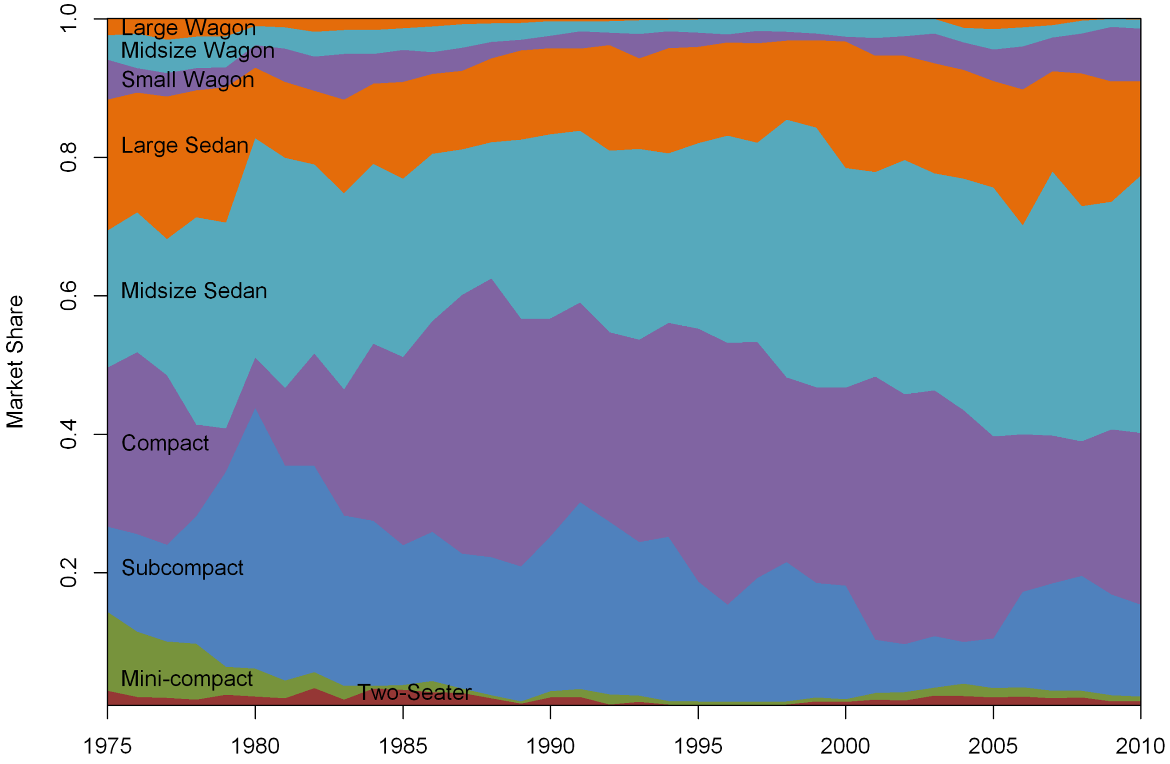

The estimated cumulative weight effects of additional features and mix shifting are summarized in Figure 3-6. There was little change in the weight of features between 1975 and 1980, but since 1980 the weight of features has steadily increased, and was estimated to account for approximately 200 kg (440 lbs) of additional vehicle weight in 2009, compared with 1975. The effect of mix shifting since 1975 peaked at an increase of 106 kg (233 lbs) in 2005, before falling back to +54 kg (120 lbs) by 2009 as the longstanding trend toward larger cars was reversed. The results in Figure 3-6 indicate that if not for the effects of weight-saving technologies, mix shifting and the adoption of new features would have increased the weight of the average car by

Figure 3-5: Weight effects of changes in since 1975, including secondary weight.

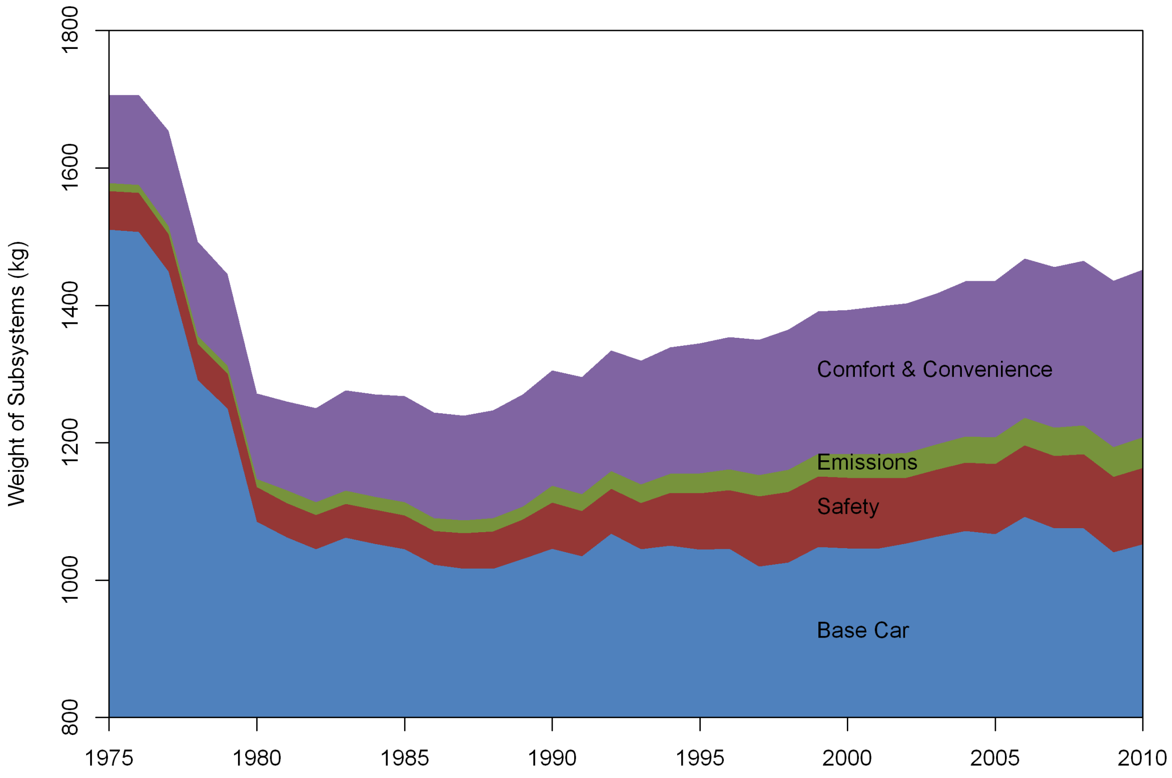

approximately 250 kg (550 lbs) between 1975 and 2009.

3.6.2 Sources of weight reduction

Technological improvements, including greater use of front-wheel drive, unibody construction, lighter and stronger materials, and a shift from 8-cylinder to 4- and 6-cylinder engines, have all tended to reduce the average weight of new cars since 1975.

Figure 3-7 shows the fractions of cars using each construction type from 1975 through 2010. The shares in Figure 3-7 are based on data compiled by Audatex North America, and are based on available models rather than sales. About half of the cars offered for sale in the U.S. in 1975 used unibody construction, but this share

Figure 3-6: Cumulative weight effects of feature changes and mix shifting.

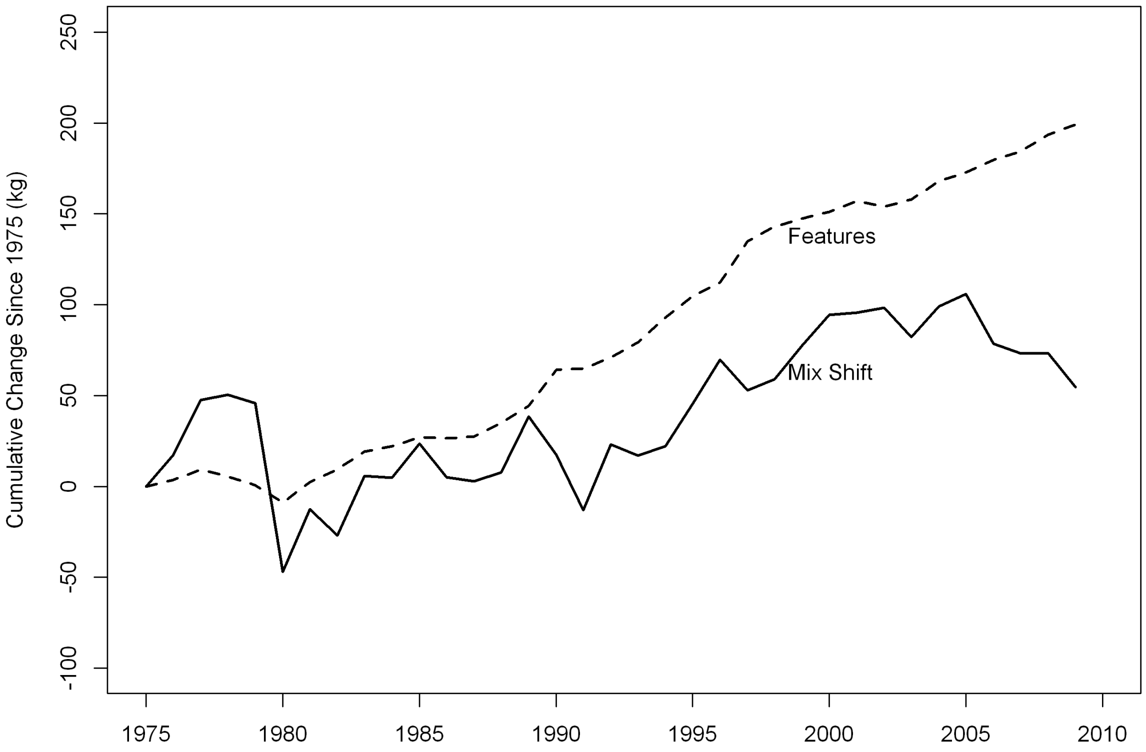

rose steadily to more than 90% by 1990, and has changed only slightly since then.

Figure 3-8 shows the market share of front-wheel drive in U.S. cars from 1975 to 2010, as reported by EPA (2012). These values are sales-weighted. Less than 10% of new cars employed front-wheel drive in 1975. As with unibody construction, this share rose rapidly through 1990, but has remained fairly steady since then.

Figure 3-9 shows the sales-weighted shares of 4-, 6-, and 8-cylinder engines by year, which together account for the overwhelming majority of all engines in U.S. cars. Eight-cylinder engines suffered a huge loss in market share between the late 1970s and early 1980s, followed by slower declines in later years. After peaking in the mid-1980s, the share of 4-cylinder engines declined slightly and then held steady before recovering as fuel prices rose after 2005.

Figure 3-7: Availability of major construction types in U.S. cars since 1975.

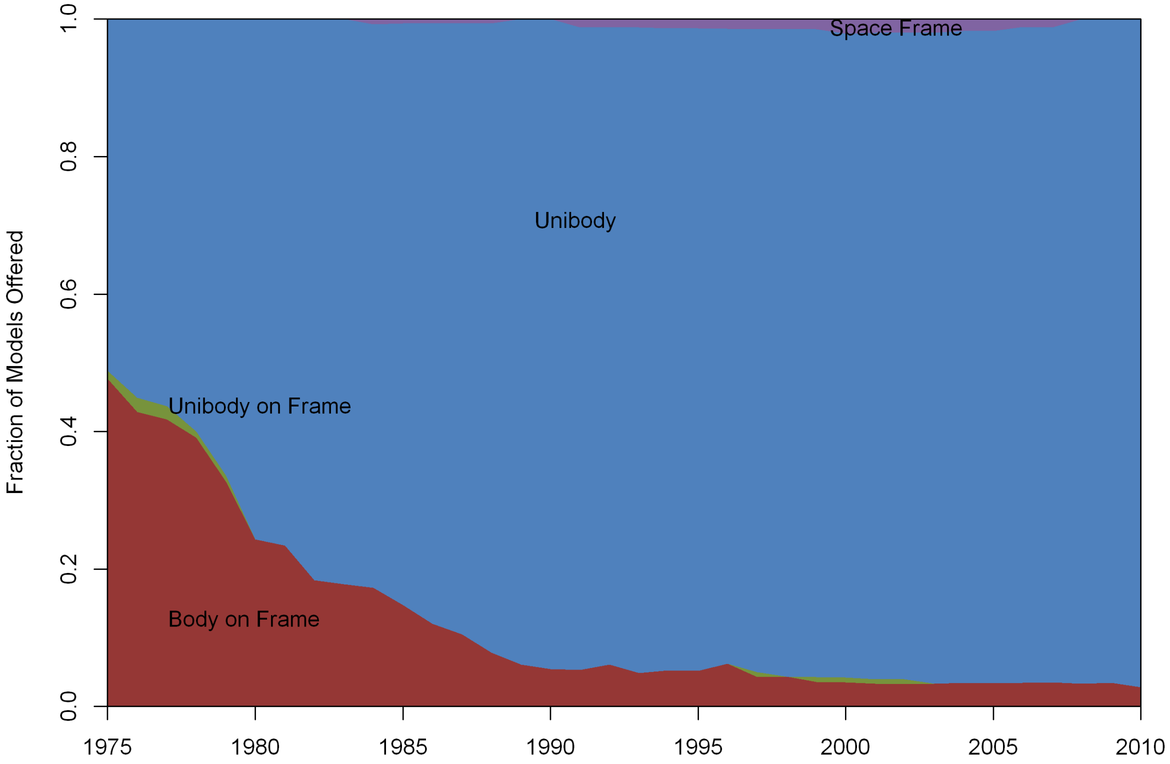

The materials content of cars was compiled from several related sources, primarily various editions of the Transportation Energy Data Book (TEDB) published by Oak Ridge National Laboratory. The scope of the materials analyses and the original sources of the data varied from year to year, but there was substantial overlap between the various time series. Table 3.3 summarizes the years included in each time series, and which of these were incorporated into the time series reported in this work. In years for which no data could be found, materials content was linearly interpolated.

The materials content data warrants additional discussion. Two potential issues are the changing scope used over the years (U.S.-built car, domestic car, domestic

Figure 3-8: Sales-weighted penetration of front-wheel drive in passenger cars.

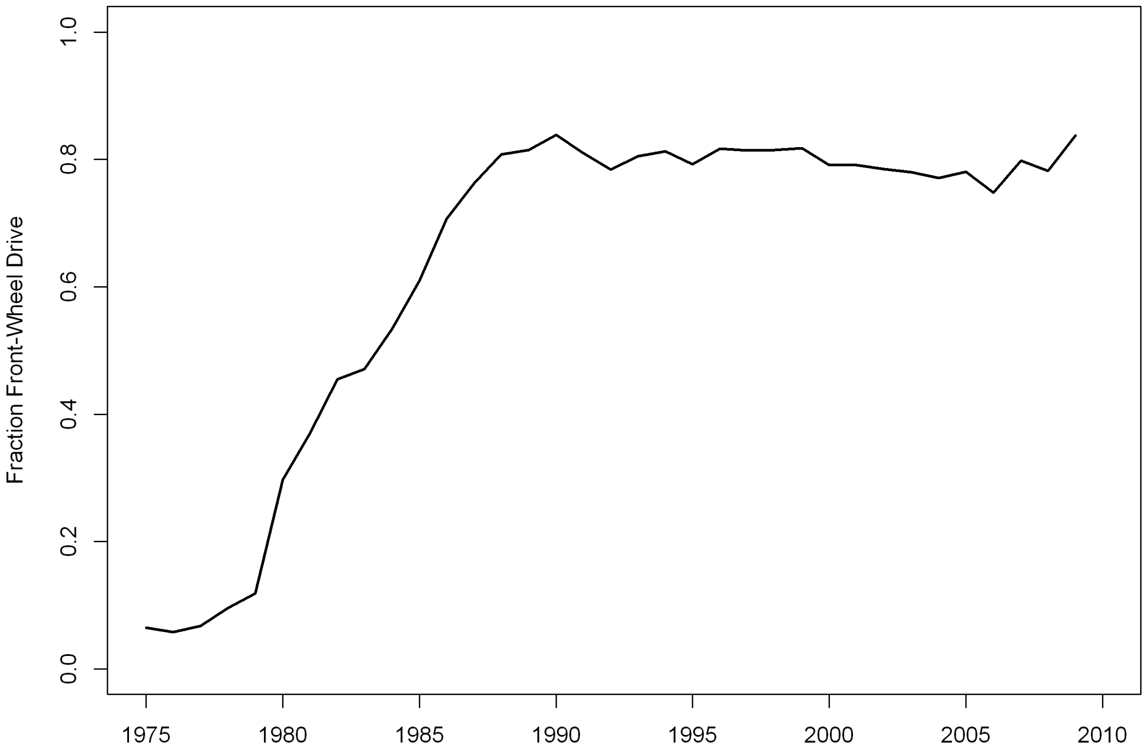

Table 3.3: Sources of material content data used in this analysis.

Series	Source	Scope	Years Available	Years Used
(1)	TEDB, Ed. 6 (citing Ward's)	U.S.-built car	1975–1981	1975–1981
(2)	TEDB, various (citing Ward's)	Domestic car	1978, 1984, 1985, 1992, 1993, 1994, 1996, 1997	1984, 1985, 1992, 1993, 1994, 1996, 1997
(3)	TEDB, various (citing American Metal Market)	Domestic car	1977, 1978, 1985, 1987, 1990, 1998, 1999, 2001, 2003, 2004	1987, 1990
(4)	TEDB, various (citing Ward's)	Domestic light vehicle	1995–2009	1998–2009
(5)	American Chemistry Council	Domestic light vehicle	1987–2005	None

light vehicle), and the differing original sources of the data. Fortunately, there is substantial overlap between the different data sets, and in general the weight fractions

Figure 3-9: Market shares of 4-, 6- and 8-cylinder engines in new U.S. passenger cars.

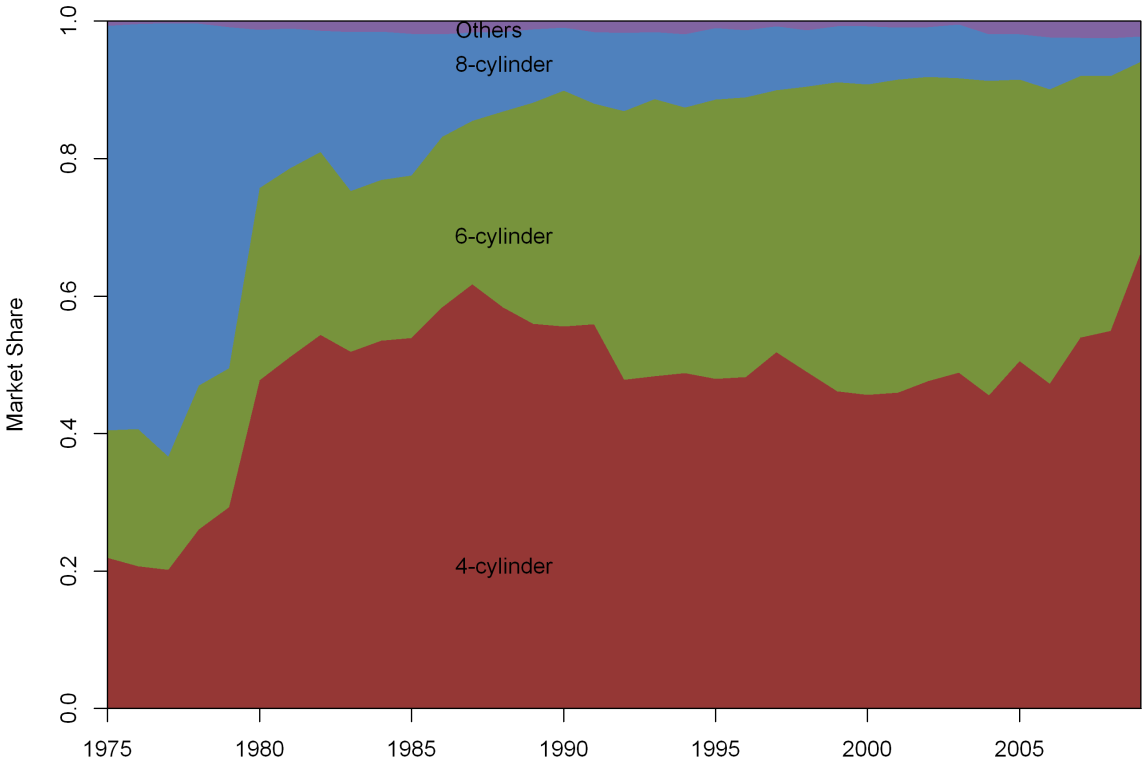

of each of the key materials are very similar in the different data series. This helps to mitigate concerns over inconsistencies in the data over time. Moreover, personal communication with Ward's has confirmed that the American Chemistry Council is the original source of the Ward's data.

An additional issue is that these data are based on domestic vehicles, and not imports. This is a legitimate concern, but in the absence of any data on the materials composition of imported cars, the materials composition time series reported here were assumed to apply to all cars sold in the U.S.

Figure 3-10 charts the evolution of the materials composition of U.S. cars based on the composite data set described above. The share of conventional iron and steel

Figure 3-10: Materials composition of new U.S. cars since 1975.

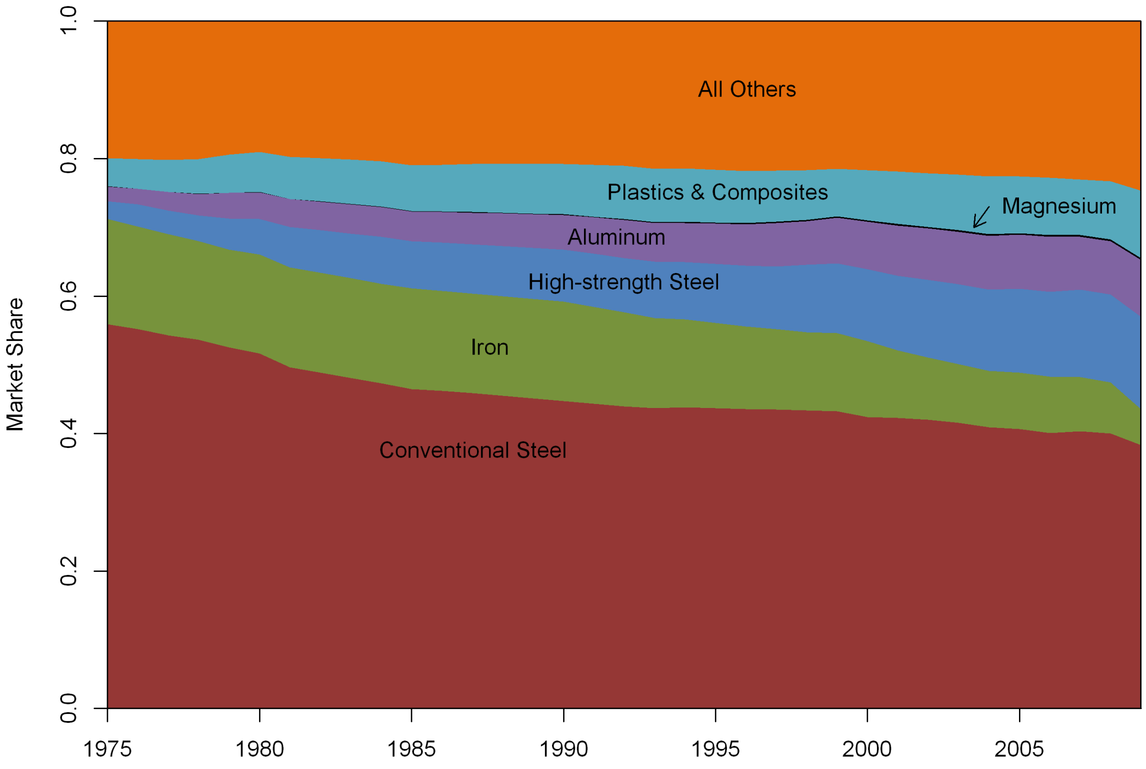

in overall vehicle weight has shrunk by a third since 1975, with HSS, aluminum, and plastics rising to fill in the gap. Plastics & composites account for twice as large a share of weight as they did in 1975, while aluminum content has quadrupled and HSS content quintupled. Although magnesium's share has grown by a factor of 10, it still only accounted for about 0.3% of vehicle weight in 2009. Miscellaneous materials (rubber, glass, fluids & lubricants, etc.) have risen modestly, from 20% to 25% of overall vehicle weight.

The estimated effects of the various technological improvements on average car weight are summarized in Figure 3-11. Each series corresponds to the estimated cumulative weight change relative to 1975. All of these technological changes contributed significantly to weight reduction, especially in the early years. Three distinct phases are evident in the results show in Figure 3-11. Between 1975 and 1982, average weight dropped by 52 kg, or about 3% of the average 1975 car weight, each year. From 1982–1990, the pace slowed considerably, to 26 kg per year, or about 2% of the average 1982 car weight.[1] Since 1990, weight reductions have largely been driven by alternative materials substitution, as the market has become saturated with front-wheel drive and unibody cars and the shift away from 8-cylinder engines slowed considerably. Between 1990 and 2009, annual weight reduction was 11 kg per year, or roughly 1% of the weight of the average 1990 car. All told, the combined effects of technology are estimated to have reduced the weight of the average new car in the U.S. by 790 kg (1,700 lbs) versus 1975.

3.7 Estimating average weight

The effects of growing feature content, mix shifting, and weight-reducing technologies were summed up as outlined in section 3.5 (Estimating Fleet-Level Contributions of Technologies and Functionality). The resulting weight estimates for each year are plotted in Figure 3-12, along with the actual average weight of new cars as reported by U.S. EPA. The estimates developed here capture the general trends observed in the weight of the average new car between 1975 and 2000: namely, a large and abrupt decrease in the late 1970s–early 1980s, followed by a gradual reversal and increase. After 2000, actual average weight continued to climb, while the estimates of this work declined slightly, driven by growth in alternative materials. As with any change in new vehicles, changes in the application of features or weight-saving technologies will take several decades to fully permeate the in-use fleet, due to the dynamics of fleet turnover.

Although the results reported here reflect the general trends in average car weight, two key discrepancies are apparent in Figure 3-12. First, the estimates developed in

---

[1] 26 kg is about 2% of the weight of the average 1982 car, even though 52 kg is only 3% of the weight of the average 1975 car, because the average 1982 car was so much lighter than the average 1975 car

Figure 3-11: Estimated contributions of technological improvements to weight reduction in new U.S. cars, since 1975.

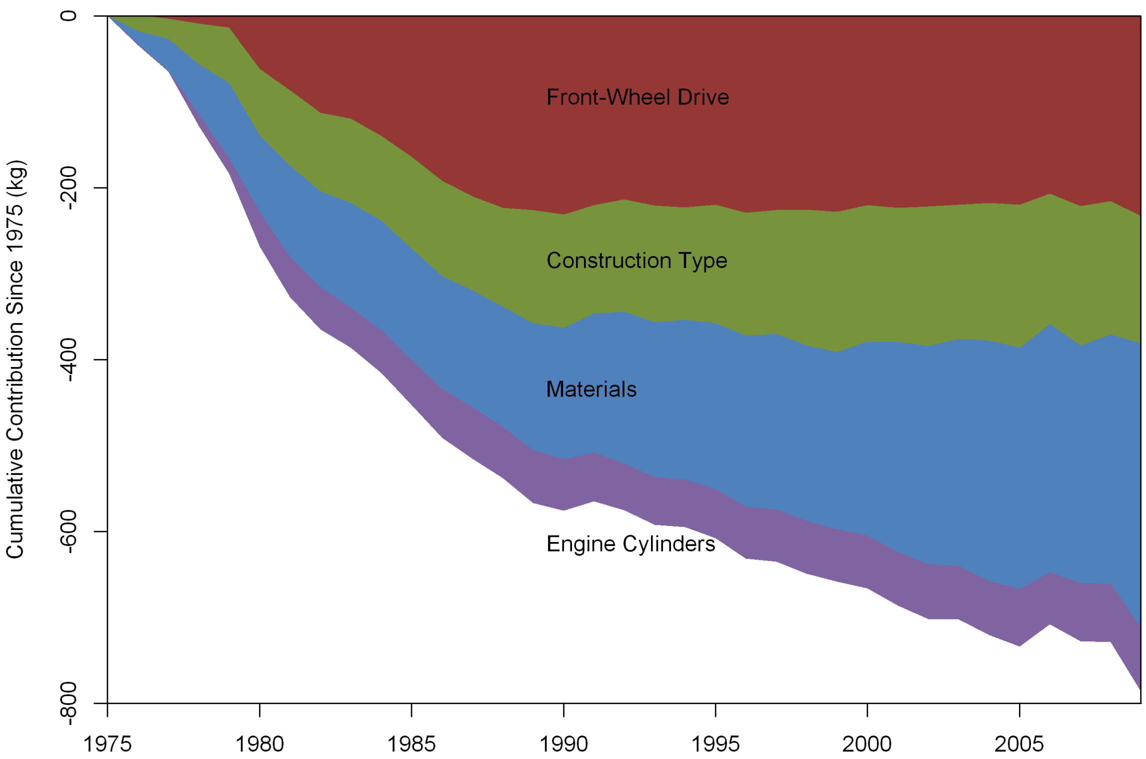

this work are too high in the early years. Second, the estimates continue to fall while actual weights remained constant, and have not risen as quickly as actual weights have in recent years. These discrepancies are highlighted by a residual term, also plotted in Figure 3-12 and calculated as the difference between the actual average weight and the average weight predicted by this analysis. That the predicted values for weight would not perfectly agree with the true averages should not be surprising, given the disparate and often highly aggregate nature of the data used to generate the estimates. Nonetheless, patterns in the actual and estimated weight trends may be due to a variety of causes, and deserve further discussion.

It appears that some substantial source of weight reduction has been omitted from the estimated values in 1976–78, which suggests that one of two things may have been happening at that time. There may have been additional technological improvements that served to reduce weight. Some of these were discussed in section 3.2, and might include improvements in design, engineering, and manufacturing practices. Another possibility is that design tradeoffs were made that sacrificed functionality in order to save weight, but that these sacrifices were not captured in the data on feature content or size mix. The latter explanation seems less plausible, given that there was so little change in weight due to feature content or size reductions during this period. Presumably, if design tradeoffs were being made to save weight, they would not be restricted to only those attributes that happen to be excluded from our data set.

A second feature of Figure 3-12 is that the estimated weight is approximately constant after 1988, but actual weight increased during this period. This suggests that additional features or functionality, not captured in the data on feature content, may have been added to cars during this period. Some possibilities were discussed in section 3.5, and include things like NVH improvements. It is also possible that the estimated weight reductions from technology in this latter period were too high. However, in order to fully explain the actual change in weight during this time, the weight savings from technology improvements (largely alternative materials during this period) would need to be essentially zero. Even if the true weight savings from alternative materials were lower than assumed, it is unlikely that they would be zero, since there would be no incentive to adopt more expensive materials if they didn't save weight (or equivalently, increase strength per unit weight).

An additional possible explanation for the discrepancies between the actual and predicted weights is that certain parameters assumed to be constant in this analysis may have actually varied over time. For example, it is plausible that the secondary weight factor might have been higher in earlier years and lower in more recent years. Alternatively, the weight savings from front-wheel drive or unibody construction may have declined over time. While these factors would tend to push the estimated weights closer to the actual weights (lower in the early years, higher in the later years), none of them on its own is adequate to resolve the discrepancies between the actual and estimated average weights.

Figure 3-12: Actual and estimated average weights of new U.S. cars since 1975, and differences between these (residuals).

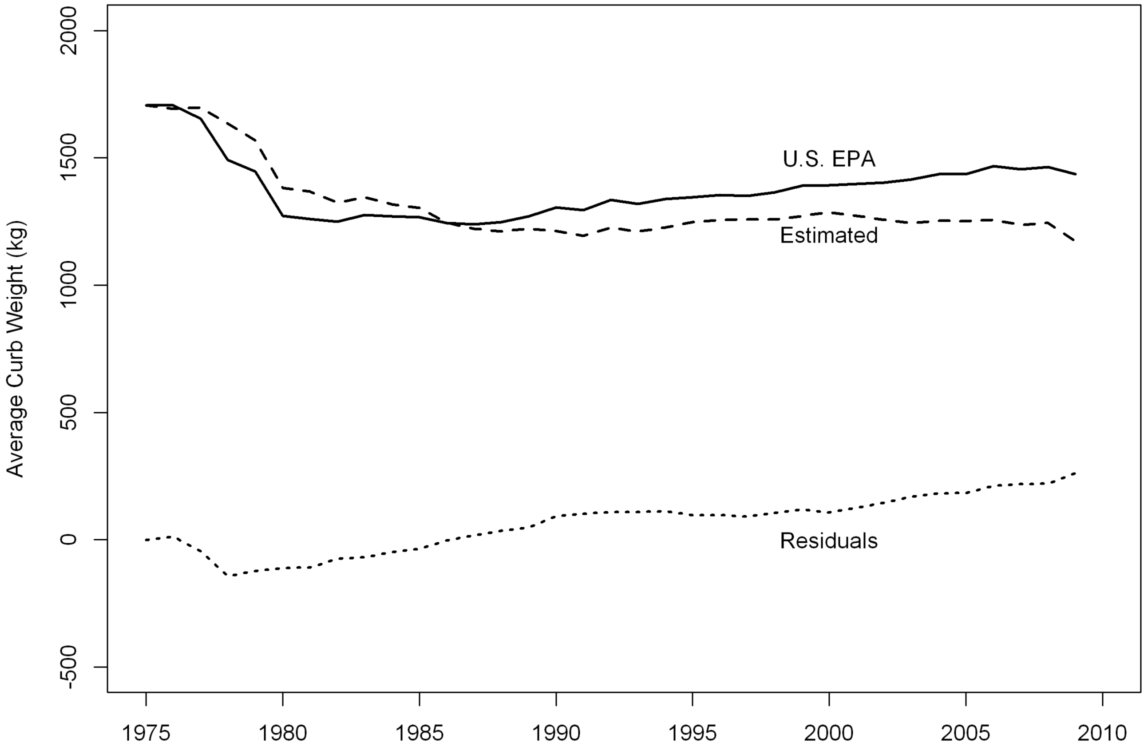

3.8 Conclusions

In the work reported here, a bottom-up analysis was developed to explain as fully as possible the rapid decline and subsequent steady increase in the average weight of new U.S. cars since 1975. The weight effects of various technologies and features were estimated using reviews of available literature and analysis of model-level data. Weight trends across the fleet of new cars were estimated by aggregating up these vehicle-level data using fleet-wide penetration rates.

The best estimates from this work indicate that new features and functionality would have added at least 250 kg (550 lbs) to the weight of the average new car between 1975 and 2009, if not for offsetting improvements in technology. Over the same period, it is estimated that alternative materials, more weight-efficient vehicle architectures, and reduced engine sizes have taken 790 kg (1,700 lbs) out of the weight of the average car. These results will be revisited in the next chapter, as they are central to that chapter's work characterizing the magnitude of efficiency technology improvements in new U.S. cars.

The switch to front-wheel drive, that from body-on-frame to unibody construction, and increased use of alternative materials all contributed significantly to historic weight reductions, and observed changes were able to explain about 80% of the weight reduction that occurred between 1975 and 1982. In later years, accounting for observed features and technology replicated the observed weight increases of the 1990s, but did not capture the continued weight increases since 2000.

Several explanations are possible for the observed patterns. First, there might be substantial unobserved technology improvements that could contribute to time-varying residuals. For example, the use of alternative materials — and especially combinations of such materials — requires substantial capabilities in forming, joining, and design that would change over time. Alternatively, time-varying engineering emphasis on reducing weight that could have been omitted from the data set, particularly concentrated in the early years. Such a theory is generally supported by anecdotal evidence from industry experts at the time (Lutz, 2011) who claim that CAFE regulations and economic conditions forced radical redesigns of domestic vehicles in the early 1980s. Such a shift in engineering emphasis could result in the discontinuity of the application of technology or diminishing marginal returns in unobserved weight-reducing technologies that are reflected in the residual error observed here. This effect could potentially be modeled using a non-fixed secondary weight multiplier that is higher in early years. Finally, our inability to accurately predict weight changes during late phases of the 1975–2009 period considered could indicate a shift in focus to improve the feature and functionality of vehicles in non-observed ways such as crashworthiness or reductions in NVH.

As noted at the outset, many automakers have recently announced a renewed focus on reducing the weight of their vehicles, aiming to cut vehicle weight by 2–3% per year. Such goals are often ambiguous: do they refer to gross weight[2] reductions, accounting for the weight-savings potential of all new lightweighting technologies? Or so they refer to net weight reductions, accounting also for the growth in new features and functionalities that seem to inexorably drive weight up? The results presented in this chapter suggest that the former would be difficult enough, and the latter very challenging indeed. Gross reductions of 30–40 kg, or 2–3%, each year fall in between the rates of weight reduction realized in the late 1970s–early 1980s and those realized later in the 1980s. Some of the technologies available in the 1970s and 1980s — most notably unibody construction and front-wheel drive — are now found on almost all new cars, limiting their potential for delivering further weight reductions. However, additional weight reductions might still be found through greater use of alternative materials, and if processes can be developed that make space frame construction more practical for high-volume models. Moreover, one-third of new light truck models in the U.S. still use body-on-frame construction and one-quarter use rear-wheel drive, so the potential for weight reduction among light trucks may be somewhat greater than among cars.

In light of the many unresolved questions in this area, an update to this work within the next few years will be particularly telling. Will automakers realize their weight reduction goals? If so, will they do so primarily through observable materials or structural changes, or will these reductions come through more subtle design and process changes? More importantly, however, will be whether vehicle weight actually declines, or whether new weight-saving technologies are offset by the continued growth in vehicle features and functionality witnessed over the past three decades.

---

[2] the term "gross" is used here not in the sense of gross vehicle weight, but to distinguish it from net changes

Chapter 4

A Broader View of Efficiency Technology Gains in U.S. Cars

Quantitative measurements of historic improvements in fuel efficiency technology help to illuminate the feasibility of future fuel economy standards. Past investigations have produced widely varying estimates of this rate of improvement, though all seem to indicate that fuel consumption reductions implied by the 2025 CAFE standards cannot be met solely through technological improvements at historic rates. In this chapter, I estimate that holding all else equal, a 1% increase in weight increases a car's fuel consumption by 0.69%, and a 1% reduction in 0–97 km/h acceleration time increases fuel consumption by 0.44%. These tradeoff parameters are combined with technological improvements documented in the preceding two chapters to yield a new, more comprehensive measure of technological improvements. When accounting for all of these sources of improvement, I conclude that the per-mile fuel consumption of new cars in the U.S. could have been reduced by 5% per year from 1975–1990, if acceleration, features, and functionality had remained at their 1975 levels. Approximately 80% of this potential was realized as actual reductions in fuel consumption. Between 1990–2009, technological improvement averaged just 2.1% per year, only 34% of which was realized as actual fuel consumption reductions. To meet the 2025 CAFE standards for cars without sacrificing capabilities that consumers have come to expect, technology must improve quickly enough to reduce fuel consumption by 4.3% per year for 14 years — considerably faster than has occurred since 1990, but consistent with the pace of improvements observed between 1975 and 1990.

4.1 Introduction

The Federal government has recently finalized Corporate Average Fuel Economy (CAFE) standards for light-duty vehicles in the years 2017–2025. The standards will increase by 3.8–4.7% annually for cars, and by 2.5–4.9% annually for trucks, during these eight years. Combined with previously-announced increases for the years 2011–2016, the new rules are expected to yield fuel economy increases of 80% and 60% for cars and light trucks, respectively, over 14 years (EPA and NHTSA, 2010, 2012).

An important question for policymakers, researchers, and analysts is how ambitious these targets — equivalent to reductions in per-mile fuel consumption[1] averaging 3.4–4.3% per year for 14 years — really are, relative to historic rates of technology improvement in light-duty vehicles. If required rates of fuel consumption decrease are within historic rates of technology change — and provided that there are similar levels of efficiency improvements yet to be exploited — it would suggest that the standards can be met through business-as-usual technology improvements (Lutsey and Sperling, 2005). This might mean foregoing additional gains in other vehicle attributes such as features or performance, but would not require giving up levels of the attributes that consumers have already come to expect. If, on the other hand, the required fuel consumption reductions are outside the range of historic technology improvements, it suggests that new technologies will need to be deployed more rapidly than has been done historically, or that other vehicle attributes will have to be "taken back" from their current levels.

---

[1] Throughout this dissertation, and in this chapter in particular, I refer frequently to *fuel consumption.* I use fuel consumption to refer to the amount of fuel that a vehicle consumes for a specified distance traveled, in units such as liters per 100 km or gallons per 100 miles. Fuel consumption therefore refers to the inverse of fuel economy, which is commonly expressed as miles per gallon. Fuel consumption in this sense should not be confused with the total quantity of fuel consumed by all vehicles combined; even as average fuel consumption has fallen over the years, the total quantity of fuel consumed has increased as the number of vehicles has grown and their collective distance traveled has increased.

The constraint of maintaining functionality at current levels has been central to establishing technical feasibility and economic practicability in recent rulemakings on automotive greenhouse gas and fuel economy standards. The Final Rule setting standards for 2017–2025 states that:

> [T]hese rules should not have a significant effect on the relative availability of different size vehicles in the fleet. The agencies' analyses used a constraint of preserving all other aspects of vehicles' functionality and performance, and the technology cost and effectiveness estimates developed in the analyses reflect this constraint. — EPA and NHTSA (2012)

That is to say, the feasibility of recent fuel economy standards has been built on a premise of maintaining constant acceleration performance and other measures of functionality.

There are two objectives to the work reported in this chapter. The first objective is to quantify the historic rates of technology improvement in U.S. cars. Technology improvement is measured as how much the per-mile fuel consumption of the average new car could have been reduced over time, if not for changes in acceleration, size, features, and other attributes valued by consumers. The second objective is to characterize concisely the tradeoff between acceleration performance and fuel consumption. This tradeoff model will be integral to the product portfolio simulations pursued in Chapter 6.

Before proceeding, it is helpful to contrast two major views of technology improvement in motor vehicles (or in any sophisticated, energy-consuming system). The first view can be thought of as bottom-up, in that it focuses on the contributions of individual subsystems, vehicle loads, and "widgets" that can be adopted to incrementally improve efficiency. In contrast, there is a top-down view that focuses on the services and attributes that are provided to the user of the vehicle. This contrast can also be thought of as one between inputs (what goes into building the system) versus outputs (how well the system performs).

Much work, both retrospective and prospective, employs the bottom-up view. For example, the modeling system underlying the CAFE rulemaking process employs this view, evaluating the cost effectiveness of potential future technology packages. For example, works by EPA (2013b) and Zoepf and Heywood (2012) study the deployment of individual technologies. Kasseris (2006) developed projections of future vehicle efficiency based on assumed values for road loads and efficiency parameters in various subsystems — many of which were informed by their historic trends. One challenge when working with the bottom-up view is the potential interdependence of the rates of adoption of multiple technologies at the same time. Zoepf and Heywood (2012) examined the rates of adoption of individual features and engine technologies, but it is much harder to know how quickly multiple vehicle technologies can be integrated into the production system simultaneously.

The top-down view provides a simpler perspective on these subsystem interactions, by focusing on the efficiency with which the overall system delivers key services and utility attributes. In the case of vehicles, this may mean characterizing the fuel consumption of a vehicle for a given level of size, acceleration capability, and feature content. An and DeCicco (2007) reviewed a number of studies from the engineering literature which have attempted to quantify past improvements in automotive technology, or technical efficiency in their parlance, using what amounts to a top-down perspective. They found widely varying estimates ranging from 1.0–3.8% per year (though mostly clustered between 1.5–2.2% per year). An and DeCicco noted that:

> Because the full range of features that interact both physically and economically with fuel economy cannot be observed with publicly available data, fully characterizing automotive technical efficiency trends is probably not possible. However, at least some portion of the trend can be quantified using attributes that are readily observable (such as size or mass) or calculable from public data. — An and DeCicco (2007)

Attempting to get closer to this goal, they define a performance-size-fuel economy index, which they interpret as "represent[ing] the ratio of moving a spatial carrying capacity a unit distance with a given performance capability per unit of fuel consumed.”

An and DeCicco (2007) emphasized the importance of shifting the focus from engineering metrics like ton-miles per gallon to more consumer-centric attributes like size and performance. Focusing on attributes that are most directly relevant to consumer utility will provide a more complete picture of all of the technological improvements occurring over time. Nonetheless, An and DeCicco chose to focus on power-to-weight ratio as their measure of performance, rather than acceleration time, although the latter is arguably more directly related to consumers’ driving experience. As shown in Chapter 2, the relationship between acceleration performance and power-to-weight ratio has changed over time, as newer vehicles tend to deliver faster acceleration than do older vehicles with comparable power and weight. A second downside to An and DeCicco’s approach is that they make a strong assumption about the relationship between size, power-to-weight ratio, and fuel economy. Namely, their performance-size-fuel economy index implies that for any given technology level, a proportional increase in size or power-to-weight ratio will be met with an equal and opposite proportional decrease in fuel economy. However, there is no theoretical or empirical reason to believe that this 1:1:1 tradeoff must hold.

Knittel (2011) showed that the tradeoffs between engine power, vehicle weight, and fuel economy are not in fact 1:1. He also concluded that the fuel economy of new U.S. cars could have been increased by approximately 65% between 1980 and 2006, if not for increases in weight and power. His approach was to empirically estimate the logarithm of fuel economy as a function of the logarithms of weight (w), power (hp), torque (tq), selected covariates (X), and a set of year fixed effects (T_t):

$$\ln mpg_{it} = T_t + \beta_1 \ln w_{it} + \beta_2 \ln hp_{it} + \beta_3 \ln tq_{it} + \mathbf{X}'_{it}\mathbf{B} + \epsilon_{it} \tag{4.1}$$

In Equation 4.1, the β coefficients represent the tradeoffs between fuel economy and the various independent variables, expressed as the elasticity of fuel economy with respect to each attribute. The year fixed effects, T_t, are interpreted as the cumulative change in technology between some base year and year t. Specifically, the change in technology is expressed as the change in the expected value of log fuel economy conditional on weight, power, torque, and covariate values. Holding power, weight, torque, and covariate values constant, the difference in expected fuel economy between the base year 0 and some future year t would be:

$$\ln mpg_{it} - \ln mpg_{i0} = T_t \tag{4.2}$$

In Equation 4.2, T_0 is normalized to zero. Rearranging and taking the exponential, we see that the ratio of expected fuel economy in year t to the that in year 0, if power, weight, torque, and covariates were unchanged, is:

$$\frac{mpg_{it}}{mpg_{i0}} = e^{T_t 2} \tag{4.3}$$

Knittel's methodology offers both advantages and disadvantages relative to the prior work. It is attractive because it allows the tradeoff parameters between power, weight, and fuel consumption to be estimated as parameters of the model, rather than assuming these values *a priori*. However, a shortcoming is that by focusing on engineering attributes (power and weight) rather than consumer attributes (such as acceleration, size, and features) his definition of technology change may not capture all sources of improvement. Critically, since he conditions on power, his estimated effects for technology do not capture any improvements in how effectively vehicles turn engine power into acceleration performance.[3] As shown in Chapter 2, these improvements have been substantial. Similarly, by conditioning on weight, his specification will not capture technology improvements that allow newer vehicles to weigh less than similar vehicles in the past. As shown in Chapter 3, modern weight-saving technologies would have allowed a reduction in weight of approximately 650 kg, or 40%, for the average new car between 1975 and 2009, if not for increases in size and features over that

---

[2] For small values of T_t, $e^{T_t} \approx 1 + T_t$, so T_t represents the fractional increase in expected fuel economy, holding other attributes constant. However, as T_t increases, this approximation no longer holds.

[3] The relationship between power and other performance metrics, like towing capability or ability to hold speed on an uphill grade, may also have changed, though these are not investigatedi this work

time. The fuel consumption reductions resulting from such a change in weight are substantial, but are not incorporated into Knittel's estimates of technology change.

The objective of this chapter is to develop a more comprehensive estimate of how much automotive technology has improved since 1975. The approach taken draws heavily on the empirical approach of Knittel, while embracing the philosophy of An and DeCicco by focusing on attributes most relevant to consumers. The methodological approach is described in the next section. Section 4.3 briefly describes the data used in this chapter. Section 4.4 presents the resulting estimates of technology change and compares them with values found in the literature, and the final section offers some conclusions against the backdrop of the required increases in Corporate Average Fuel Economy through 2025.

4.2 Methodology

This chapter attempts to answer the question, "How much could the per-mile fuel consumption of new U.S. cars have been reduced between 1975 and 2009, if not for changes in vehicle size, performance, features and functionality?" Effectively, we are asking what would happen if we built the fleet of cars sold in 1975, using modern materials, designs, and powertrain technologies. Actually building these vehicles is clearly impractical, but engineering simulations offer one possible solution, as they might enable the application of contemporary technologies to the design of a vehicle with 1975-level performance, capacity, comfort, and safety specifications. While vastly simpler than actually building a vehicle, such simulations are nevertheless labor-intensive to develop and calibrate, and questions would likely persist over the representativeness of the specific vehicle model or models chosen for study.

Estimation of a simplified econometric model based on observed vehicle characteristics offers an alternative approach that is both tractable and can incorporate data on all vehicle models. This is the approach taken by Knittel (2011). As discussed in Section 4.1, Knittel's models may not capture all of the sources of technology improvement in new vehicles. Ideally, such a model would estimate the expected level of fuel consumption for a vehicle, conditional on all related attributes. These related attributes might include acceleration, towing, and handling capabilities, passenger and cargo capacity, some measure or measures of comfort and ride quality, and the presence of various convenience, emissions, and safety features.

In this work, detailed, model-level data on comfort, convenience, and safety features were not available, so an alternative, two-step methodology was adopted instead. First, technology change from a top-down perspective was estimated as the change in expected fuel consumption since 1975, holding weight, 0–97 km/h (0–60 mph) acceleration time, vehicle interior volume, and selected covariates constant. Next, the resulting estimates of technology change were adjusted to account for the fuel consumption benefits of weight-saving technologies, based on the bottom-up analysis presented in Chapter 3. The weight analysis is intended to capture the weight reductions that would have occurred in the average new car if size, features, and functionality had remained at 1975 levels. The basic empirical model for implementing the first step was:

$$\ln gpm_{it} = T_t + \beta_1 \ln IWT_{it} + \beta_2 \ln Z97_{it} + \beta_3 \ln VOL_{it} + \mathbf{X}'_{it}\mathbf{B} + \epsilon_{it} \tag{4.4}$$

In Equation 4.4, gpm_{it} is the fuel consumption of car model i in year t in gallons per mile, VOL_{it} is its interior volume in m^3, $Z97_{it}$ is its 0–97 km/h acceleration time in seconds, IWT_{it} is its inertia weight in kg, and $\mathbf{X}_{it}$ is a vector of dummy variables indicating whether the vehicle has a manual transmission or all-wheel or 4-wheel drive, whether it is a two-seater or a wagon body style. Also included were terms for the interactions of manual transmissions with year, all-wheel or 4-wheel drive with year, and two-seater or wagon body styles with interior volume. As in Equation 4.1, T_t is a set of year fixed effects representing the expected reduction in log fuel consumption for cars in each year t relative to some base year, if the other attribute levels had remained unchanged. The year fixed effects T_t will therefore be interpreted as the improvement in technology of the average new car between the base year and year t.

Some specifications also included dummy variables for powertrain type[4] engine specific power quintiles.[5] The reason for including specific power is to control for the possibility that more sophisticated engine technologies tend to be correlated with heavier vehicle weight or higher performance. Specific power (the ratio of engine peak power to displacement, often measured in kW/liter) is commonly used as a measure of the technical sophistication of an engine (Chon and Heywood, 2000). If more sophisticated engines tended to be used in heavier or higher-performance vehicles, then ignoring differences in technology could lead to biased estimates of the coefficients on weight or acceleration performance (β_1 and β_2). Dummy variables for quintile were used to introduce specific power for two reasons. First, since engine specific power has generally increased over time, quintiles were used to provide a measure of specific power relative to other vehicles in the same model year. Controlling for the absolute level of specific power would bias the estimates of technological improvement downwards, since specific power increases over time are themselves a part of the overall technological improvement that we want to come out in the year fixed effects. A second reason for using dummy variables for specific power quintile is that it does not impose any assumption on the particular form of the relationship between specific power and fuel consumption.

The second step involves estimating the expected fuel consumption of a new car in year t relative to year 0. This is done as in Equation 4.3, but with an additional adjustment to account for the weight reduction that would have occurred if not for changes in size, features, and functionality:

$$\frac{gpm_t^{potential}}{gpm_0} = e^{T_t}\left(\frac{IWT_t}{IWT_0}\right)^{\beta_1} \tag{4.5}$$

In Equation 4.5, IWT_0 is the inertia weight of an average new car in some base year, and IWT_t is the estimated inertia weight of a similar car using weight-reducing

---

[4] dummy variables were created for turbocharged gasoline, supercharged gasoline, naturally aspirated diesel, turbodiesel, and hybrid electric powertrains, with naturally aspirated gasoline engines representing the base case.

[5] The first quintile consists of those vehicles with the lowest one-fifth of engine specific power values in their model year, the second quintile includes those vehicles with engine specific power values in the second fifth in their model year, etc.

technologies characteristic of year t.

4.3 Data

The data used to estimate the model in this chapter were obtained from the U.S. Environmental Protection Agency. The data include interior volume, inertia weight, body style, and powertrain characteristics for all cars and trucks offered for sale in the U.S. between 1975 and 2009. However, only the data for cars were used, since the weight analysis in Chapter 3 was performed only for cars. The 0–97 km/h acceleration times were estimated for these cars using the methods reported in Chapter 2.

4.4 Results

4.4.1 Model Estimation Results

Table 4.1 contains results of the estimation of several different specifications of the general model provided in Equation 4.4. Multiple model specifications were explored in order to examine the effects of including different sets of control variables on the estimates of tradeoff parameters and technological progress. In each case, the coefficient on a continuous variable represents the partial derivative of log fuel consumption with respect to that variable. Since logs of all continuous variables are used, each of these coefficients represents the estimated elasticity of fuel consumption with respect to the corresponding variable.[6] Model 1 employs a similar specification to that used by Knittel (2011). The remaining model specifications in Table 4.1 explore the effects of weight, acceleration performance, size, body style, powertrain type, and engine specific power on fuel consumption, and provide estimates of the degree of technological improvement (the year fixed effects) for each year since 1975.

Model 1 is included to provide a baseline comparison with the results of Knittel

---

[6] For example, in Model 5, the coefficient on log of inertia weight ($\ln IWT$) indicates that the elasticity of fuel consumption with respect to inertia weight is 0.686. That is to say, according to this model, the expected change in fuel consumption resulting from a 1% increase in inertia weight is an increase of 0.686%.

(2011), and its specification is similar to that of Model 2 in Knittel's work. The reported effects are similar in magnitude but opposite in sign to the results of Knittel, because Knittel used log of fuel economy (in miles per gallon) as the dependent variable, while the present work uses log of fuel consumption as the dependent variable. For Model 1, the difference in the fixed effect estimates between 1980 and 2006 is equal to 0.503, which is very close to the results reported by Knittel (0.512 between 1980 and 2006 in Knittel's Model 2).

Model 2 controls for acceleration time instead of power, which leads to two notable changes. First, it increases the sensitivity of fuel consumption to weight. This is to be expected, since increasing weight at constant power should lead to both higher fuel consumption and slower acceleration. To hold acceleration constant while increasing weight, an increase in power is required, which reduces fuel consumption beyond what is expected from the weight change alone. The second notable change when controlling for acceleration performance rather than power is that the magnitude of the estimated technology changes (captured in the year fixed effects) increases. This is also to be expected, given the finding in Chapter 2 that newer vehicles can extract better acceleration performance from the same weight and power than could older vehicles.

Model 3 represents a further shift toward controlling for attributes more closely linked to consumer utility, introducing terms to control for all-wheel drive and for body styles (the presumed basic body style is sedan/coupe). Not surprisingly, all-wheel drive is associated with higher fuel consumption, but this effect has been shrinking over time. Model 3 also drops the variables identifying different powertrain types. The reasoning behind this decision is that shifts toward more inherently efficient powertrain technologies are themselves a part of the overall process of technology change, so it is desirable to capture their contributions to overall efficiency in the year fixed effects.

Model 4 introduces the dummy variables for specific power quintiles. The coefficients on the specific power quintile variables decrease with increasing quintiles. This indicates that within a given year, vehicles with higher engine specific power tend to have lower expected fuel consumption, conditional on their other attributes. In addition, accounting for the specific power quintiles changes the coefficient estimates on weight and 0–97 km/h acceleration time. Since specific power is commonly used as a proxy for the technological sophistication of the engine, this suggests that the relative sophistication of a vehicle's engine (compared to others in the same model year) is correlated with weight and acceleration performance.

In Model 5, the variables for powertrain type (diesel, hybrid, etc.) are reintroduced, in order to check whether the coefficient estimates on other variables are robust to their inclusion. Notably, the coefficient estimates on weight and acceleration performance change when the model includes powertrain type. This suggests that weight and acceleration also tend to be correlated with powertrain type, and that we may obtain biased estimates of the coefficients on weight and acceleration if we omit the variables indicating powertrain type. The estimates of technology change, reflected in the year fixed effects, also decrease in magnitude when powertrain type is included in the model. This latter result is consistent with the idea that growth in these powertrains (diesels, boosted gasoline engines, and hybrids) constitutes an increase in the technical efficiency of vehicles. Thus, when powertrain type is not explicitly controlled for, the efficiency effects of shifts to these powertrain technologies will be captured in the fixed effects. However, when powertrain type is controlled for, efficiency gains stemming from the shifts in powertrain type will not be represented in the estimates of technology improvement.

Models 6–8 are used to investigate the effects of size, expressed as interior volume, on fuel consumption. Size data were missing for 1975 and 1976, so these models were based only on the years 1977–2009. Model 6 is identical to Model 4, except that it excludes observations from 1975–76. Excluding these years has little effect on any of the coefficient estimates or year fixed effects (the year fixed effects are shifted toward zero in Model 6, since the base year for this model is 1977 instead of 1975, as in Model 4).

Model 7 introduces terms for interior volume, and Model 8 also includes terms for powertrain type. Including the size terms in Model 7 has little effect on the estimates of the weight and acceleration coefficients, or on the estimated year fixed effects (representing technology improvements). Adding in the powertrain variables creates a change similar to that seen between Model 4 and Model 5: a larger coefficient on weight, a smaller (in magnitude) coefficient on acceleration, and lower estimates for the year effects. The estimated coefficients on interior volume are negative, indicating that larger size is associated with slightly lower fuel consumption. Although much of the effect of size on fuel economy is through its effect on weight, and weight is controlled for separately, it is not entirely clear why larger size would be associated with higher fuel economy. This may be a result of the preferential application of efficiency technologies to larger cars (which also tend to be more expensive). However, while opposite in sign to what might be expected, the effect of size is a small one and of little practical significance.

Model 5 is the preferred specification among those investigated here, forming the basis of the technology potential calculated in the next section and the tradeoffs between acceleration and fuel consumption in the work reported in Chapter 6. Comparing Models 2–5, it is apparent that failing to account for body style, powertrain type, or engine specific power introduces biases into the estimates of the tradeoffs between fuel consumption, weight, and acceleration. Comparing Models 6–8 with the other specifications indicates that including interior volume as a regressor has little effect on the estimated tradeoffs between fuel consumption, weight, and acceleration time, or on the estimated improvements in technology. Thus, Model 5 is preferred because it provides estimates of technology improvement across all years (unlike Models 6–8) and because it avoids the apparent bias that would be introduced by omitting either engine specific power or powertrain type from the model specification (as in Models 2–4). Model 5 also has the greatest explanatory power (measured as the highest adjusted R^2 value) among all of the models investigated here, although the differences were relatively small. Finally, it should be reiterated that since Model 5 controls for powertrain type, its estimates of technology improvement omit the gains due to the introduction of diesel and hybrid powertrains, and can thus be regarded as slightly conservative in this respect.

Table 4.1: Results of estimating regression models of car fuel consumption as a function of weight, size, power, acceleration performance, and related attributes. Year fixed effects represent improvements in technology over time. Standard errors for each estimate are listed in parentheses.

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
(Intercept)	-7.985^{***}	-7.027^{***}	-6.153^{***}	-5.637^{***}	-6.594^{***}	-5.613^{***}	-5.786^{***}	-6.824^{***}
	(0.028)	(0.034)	(0.036)	(0.040)	(0.038)	(0.041)	(0.043)	(0.041)
$\ln IWT$	0.521^{***}	0.730^{***}	0.649^{***}	0.604^{***}	0.686^{***}	0.592^{***}	0.625^{***}	0.717^{***}
	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
$\ln HP$	0.289^{***}							
	(0.003)							
$\ln Z97$		-0.401^{***}	-0.499^{***}	-0.559^{***}	-0.438^{***}	-0.567^{***}	-0.561^{***}	-0.433^{***}
		(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Two-seater			0.038^{***}	0.036^{***}	0.049^{***}	0.034^{***}	-0.057	-0.038
			(0.002)	(0.002)	(0.002)	(0.002)	(0.035)	(0.031)
Wagon			0.014^{***}	0.019^{***}	0.010^{***}	0.019^{***}	0.084^{***}	0.085^{***}
			(0.001)	(0.001)	(0.001)	(0.002)	(0.015)	(0.013)
Manual Transmission	-0.095^{***}	-0.132^{***}	-0.144^{***}	-0.146^{***}	-0.134^{***}	-0.151^{***}	-0.153^{***}	-0.141^{***}
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Manual$\cdot$ (Years since 1975)	0.003^{***}	0.003^{***}	0.003^{***}	0.003^{***}	0.003^{***}	0.003^{***}	0.003^{***}	0.003^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

Continued on next page

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
4WD/AWD			0.095^{***}	0.103^{***}	0.081^{***}	0.103^{***}	0.096^{***}	0.073^{***}
			(0.007)	(0.007)	(0.006)	(0.007)	(0.007)	(0.006)
4WD/AWD · (Years since 1975)			-0.002^{***}	-0.002^{***}	-0.002^{***}	-0.002^{***}	-0.002^{***}	-0.002^{***}
			(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\ln Vol$							-0.087^{***}	-0.100^{***}
							(0.006)	(0.006)
Two-seater $\cdot \ln Vol$							0.102	0.063
							(0.098)	(0.088)
Wagon$\cdot \ln Vol$							-0.037^{**}	-0.043^{***}
							(0.012)	(0.011)
Naturally Aspirated Diesel	-0.237^{***}	-0.196^{***}			-0.203^{***}			-0.201^{***}
	(0.003)	(0.003)			(0.003)			(0.003)
Turbodiesel	-0.293^{***}	-0.290^{***}			-0.296^{***}			-0.299^{***}
	(0.006)	(0.006)			(0.005)			(0.005)
Supercharged Gasoline	-0.035^{***}	-0.042^{***}			-0.028^{***}			-0.028^{***}
	(0.006)	(0.006)			(0.006)			(0.006)
Turbo Gasoline	-0.038^{***}	-0.052^{***}			-0.040^{***}			-0.046^{***}
	(0.002)	(0.002)			(0.002)			(0.002)
Hybrid Electric	-0.341^{***}	-0.351^{***}			-0.363^{***}			-0.358^{***}

Continued on next page

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
	(0.011)	(0.011)			(0.011)			(0.010)
Specific Power Quintile 2				−0.010***	−0.020***	−0.007***	−0.008***	−0.019***
				(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Specific Power Quintile 3				−0.027***	−0.037***	−0.026***	−0.026***	−0.037***
				(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Specific Power Quintile 4				−0.038***	−0.050***	−0.036***	−0.038***	−0.051***
				(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Specific Power Quintile 5				−0.065***	−0.061***	−0.063***	−0.066***	−0.060***
				(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
1976	−0.063***	−0.091***	−0.096***	−0.102***	−0.095***			
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)			
1977	−0.096***	−0.097***	−0.098***	−0.099***	−0.098***			
	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)			
1978	−0.094***	−0.127***	−0.140***	−0.148***	−0.133***	−0.050***	−0.048***	−0.033***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)
1979	−0.098***	−0.143***	−0.164***	−0.176***	−0.153***	−0.078***	−0.076***	−0.052***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)	(0.004)	(0.004)	(0.003)
1980	−0.153***	−0.190***	−0.217***	−0.229***	−0.201***	−0.131***	−0.126***	−0.098***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)

Continued on next page

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
1981	−0.204***	−0.244***	−0.276***	−0.288***	−0.254***	−0.191***	−0.185***	−0.150***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)
1982	−0.242***	−0.256***	−0.290***	−0.298***	−0.263***	−0.201***	−0.196***	−0.160***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)	(0.004)	(0.004)	(0.003)
1983	−0.276***	−0.326***	−0.372***	−0.387***	−0.340***	−0.290***	−0.284***	−0.236***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)
1984	−0.300***	−0.353***	−0.401***	−0.419***	−0.369***	−0.322***	−0.316***	−0.264***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)
1985	−0.323***	−0.381***	−0.431***	−0.451***	−0.400***	−0.355***	−0.348***	−0.294***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)
1986	−0.355***	−0.430***	−0.468***	−0.493***	−0.450***	−0.397***	−0.390***	−0.345***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
1987	−0.362***	−0.432***	−0.466***	−0.492***	−0.454***	−0.397***	−0.389***	−0.349***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
1988	−0.388***	−0.452***	−0.486***	−0.513***	−0.476***	−0.417***	−0.409***	−0.369***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
1989	−0.403***	−0.482***	−0.520***	−0.551***	−0.508***	−0.455***	−0.447***	−0.402***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
1990	−0.419***	−0.523***	−0.568***	−0.604***	−0.551***	−0.508***	−0.501***	−0.445***

Continued on next page

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)
1991	−0.436***	−0.535***	−0.582***	−0.618***	−0.564***	−0.523***	−0.516***	−0.458***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)
1992	−0.455***	−0.563***	−0.614***	−0.652***	−0.592***	−0.557***	−0.549***	−0.486***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)
1993	−0.469***	−0.582***	−0.635***	−0.675***	−0.613***	−0.580***	−0.573***	−0.507***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)
1994	−0.481***	−0.597***	−0.652***	−0.693***	−0.628***	−0.598***	−0.591***	−0.522***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.004)
1995	−0.518***	−0.622***	−0.677***	−0.717***	−0.651***	−0.623***	−0.615***	−0.546***
	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.004)
1996	−0.532***	−0.641***	−0.696***	−0.738***	−0.671***	−0.643***	−0.636***	−0.566***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
1997	−0.539***	−0.649***	−0.708***	−0.751***	−0.682***	−0.657***	−0.650***	−0.577***
	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
1998	−0.556***	−0.655***	−0.715***	−0.757***	−0.687***	−0.663***	−0.656***	−0.582***
	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
1999	−0.559***	−0.659***	−0.721***	−0.765***	−0.693***	−0.671***	−0.663***	−0.587***
	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)

Continued on next page

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
2000	−0.565***	−0.683***	−0.748***	−0.794***	−0.717***	−0.700***	−0.692***	−0.611***
	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
2001	−0.587***	−0.710***	−0.781***	−0.829***	−0.747***	−0.735***	−0.728***	−0.642***
	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
2002	−0.594***	−0.712***	−0.782***	−0.831***	−0.750***	−0.738***	−0.732***	−0.646***
	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
2003	−0.612***	−0.734***	−0.806***	−0.859***	−0.774***	−0.765***	−0.758***	−0.669***
	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
2004	−0.618***	−0.739***	−0.815***	−0.868***	−0.781***	−0.774***	−0.767***	−0.676***
	(0.004)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
2005	−0.630***	−0.753***	−0.832***	−0.887***	−0.796***	−0.793***	−0.787***	−0.691***
	(0.004)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
2006	−0.656***	−0.777***	−0.854***	−0.909***	−0.819***	−0.815***	−0.808***	−0.714***
	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
2007	−0.669***	−0.787***	−0.867***	−0.922***	−0.830***	−0.829***	−0.822***	−0.725***
	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
2008	−0.681***	−0.797***	−0.878***	−0.933***	−0.840***	−0.840***	−0.833***	−0.735***
	(0.004)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
2009	−0.706***	−0.832***	−0.917***	−0.976***	−0.877***	−0.883***	−0.876***	−0.772***

Continued on next page

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
Manufacturer Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.918	0.918	0.901	0.904	0.922	0.891	0.892	0.914
Adj. R^2	0.918	0.917	0.901	0.903	0.922	0.891	0.892	0.914
Num. obs.	29829	29829	29829	29829	29829	27855	27841	27841

Statistical significance of t-tests on coefficient estimates: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

4.4.2 Overall Estimates of Technology Improvement

Figure 4-1 shows the actual average fuel consumption of new U.S. cars, along with several “potential fuel consumption” series, each representing the estimated level of per-mile fuel consumption if a particular set of attributes had remained constant at 1975 levels. The estimated year fixed effects from Model 1 (which are very similar to those of Knittel) are used to generate the red series in Figure 4-1. This series represents the expected fuel consumption, given estimated technology improvements, if power and weight had remained at 1975 levels. The estimated year fixed effects from Model 5 represent the expected fuel consumption if acceleration, weight, and the fraction of cars that were wagons, two-seaters, all-wheel drive, and unconventional powertrains had all remained at 1975 levels. This scenario is represented by the green series in Figure 4-1. The blue series in Figure 4-1 represents the expected fuel consumption if acceleration, fraction of cars that were two-seaters, wagons and all-wheel drive, and the content of safety, emissions, and comfort & convenience features had remained constant since 1975.

Figure 4-1 highlights the vast improvements in automotive technical efficiency that have been made since 1975. If power and weight had remained unchanged, per-mile fuel consumption could have been reduced by approximately 50% between 1975 and 2009. Between 1980 and 2006, the potential reduction is estimated to be about 40%, consistent with the results of Knittel (2011). If acceleration and weight had remained unchanged, per-mile fuel consumption could have been reduced by nearly 60% between 1975 and 2009. In other words, improvements in the ability to turn power into acceleration performance contributed the equivalent of a 16% fuel consumption reduction over this period. If acceleration, features, and functionality had remained constant, per-mile fuel consumption could have been reduced by approximately 70% between 1975 and 2009. In other words, new weight-saving technologies cut the average new car’s inertia weight by about 35%, and thus contributed the equivalent of a 25% reduction in per-mile fuel consumption over this period. Between 1975 and 2009, the actual fuel consumption of the average new car was reduced by 50%.

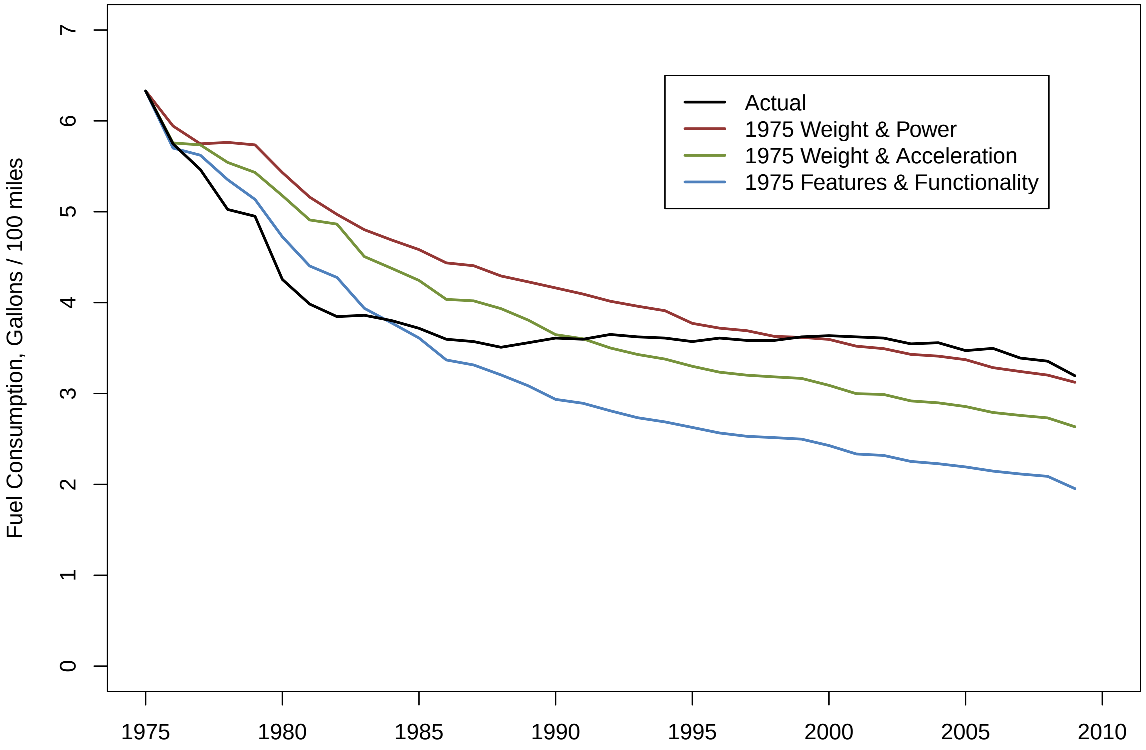

Figure 4-1: Actual fuel consumption of average new U.S. car, and potential fuel consumption if various attributes had remained at 1975 levels.

Although the improvements in technical efficiency since 1975 have been impressive, they have not occurred consistently over time. Between 1975 and 1990, the potential reduction in fuel consumption averaged 5% per year. That is to say, per-mile fuel consumption could have been reduced by 5% annually over this period if not for changes in acceleration, features, and functionality of new cars. Between 1990 and 2009, however, the average rate of change was just 2% per year. This result is consistent with the findings of Knittel (2011) that technology changed more rapidly in the 1980s than in subsequent years.

4.4.3 Sources and Sinks for Technology Gains

So far in this chapter, we have seen that improvements in technology have come from a number of sources: reductions in fuel consumption for a given level of power and weight; improvements in acceleration time, even for the same level of weight and power; and the introduction of vehicle architectures and materials that permit vehicle weight to be reduced while maintaining functionality. In this section, the relative contributions of each of these sources are first compared with one another over time. Then, we look at the major "sinks" for technology: the major attribute changes to which the efficiency improvements were applied.

Figure 4-2 shows the annual contribution of each major technology source to the overall potential reduction in fuel consumption. The red series represents the potential reduction in per-mile fuel consumption relative to the preceding year, that could have been realized if power and weight had remained unchanged. This is calculated as the year-over-year (percentage) change in the value of the red series in Figure 4-1. The green series represents the additional fuel consumption reduction that might have been realized each year if power had been reduced to maintain acceleration performance. This is calculated as the difference between the year-to-year changes in the green series in Figure 4-1, and the year-to-year changes in the red series in the same figure. The blue series represents the additional reduction that could have been realized if all weight-saving technologies had gone to reducing weight, rather than offsetting increased feature content. This is calculated as the difference between the year-to-year changes in the blue series in Figure 4-1, and the year-to-year changes in the green series in the same figure.

The red and blue series in Figure 4-2 move together; they both are higher in the earlier years and lower in more recent years. This is consistent with the intuition that when automakers are seeking to make efficiency improvements, they will encounter diminishing marginal returns in any one technology area and will seek to equalize the marginal costs of efficiency improvement across multiple technology areas. The green series is much more volatile than the other sources of technology improvement. The

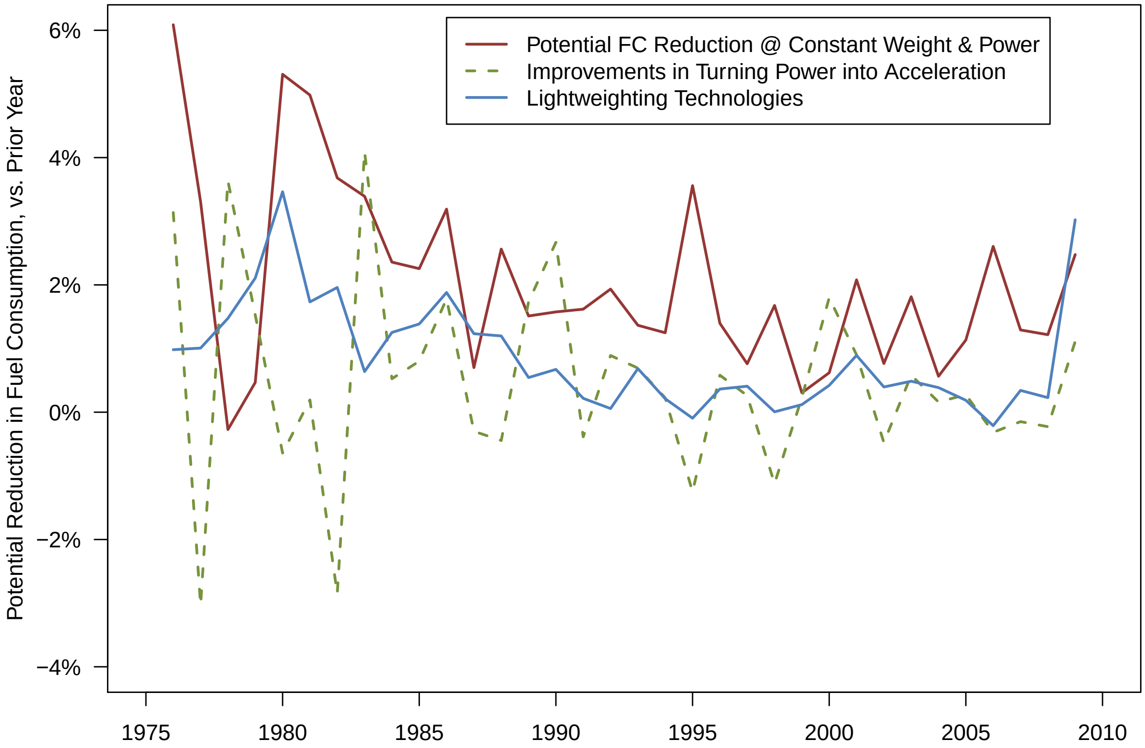

Figure 4-2: Contributions to potential reductions in per-mile fuel consumption from different technology areas.

volatility likely follows from the volatility evident in Figure 2-3, which may itself be due at least in part to the fact that the acceleration analysis reported in Chapter 2 relies on only a sample of vehicles from each year. The volatility in the green series makes it difficult to determine whether there is any correlation between the green series and the other series.

Having assessed the contributions to overall efficiency improvement from various technology sources, we now turn to the question of where the efficiency improvements have gone. To what design goals have the efficiency improvements been applied? To address this question quantitatively, Bandivadekar et al. (2008) introduced a variable that they called Emphasis on Reducing Fuel Consumption (ERFC), which they defined as the ratio of the actual fuel consumption reduction realized over a certain period, and the potential reduction that could have been realized if size and acceleration performance had remained constant. Adapting that concept to the present context, ERFC is redefined here at the ratio of the actual reduction in fuel consumption of the average new car over some n-year interval, and the potential reduction that could have been achieved if acceleration performance, features, and functionality had remained unchanged:

$$ERFC = \frac{gpm_t - gpm_{t+n}}{gpm_t - gpm_t \frac{gpm_{t+n}^{potential}}{gpm_t^{potential}}} \tag{4.6}$$

In Equation 4.6, gpm denotes actual average fuel consumption, and $gpm^{potential}$ denotes potential fuel consumption as calculated according to Equation 4.5. Figure 4-3 shows the ERFC over five-year intervals from 1975–2005 and over the four-year interval from 2005–09. Also shown are the annual average gasoline prices over the same period. Between 1975 and 1980, ERFC exceeded 100%, indicating that per-mile fuel consumption decreased by more than would have been expected at constant acceleration, features, and functionality. This suggests that either (1) actual technological improvement was greater over this period than has been estimated here, or (2) there was some pull-back in the levels of other attributes that enabled the larger decrease in fuel consumption. In fact, there was a slight decrease in acceleration times between 1975 and 1980, and a slight reduction in the weight associated with average size and feature content. These two effects appear to have approximately canceled one another out, so we can conclude that if it was a pull-back in attributes that enabled ERFC to exceed 100% over this period, it occurred in attributes other than acceleration performance and the features and size tracked in Chapter 3.

Between 1980 and 1985, ERFC fell to approximately 50%, and fell further in subsequent years, as gasoline prices remained low. Between 1995 and 2000, ERFC was negative, reflecting the fact that the average fuel consumption of new cars actually increased over this period. The emphasis on reducing fuel consumption became positive again between 2000 and 2005, and increased further between 2005 and 2009, a time when fuel prices were increasing.

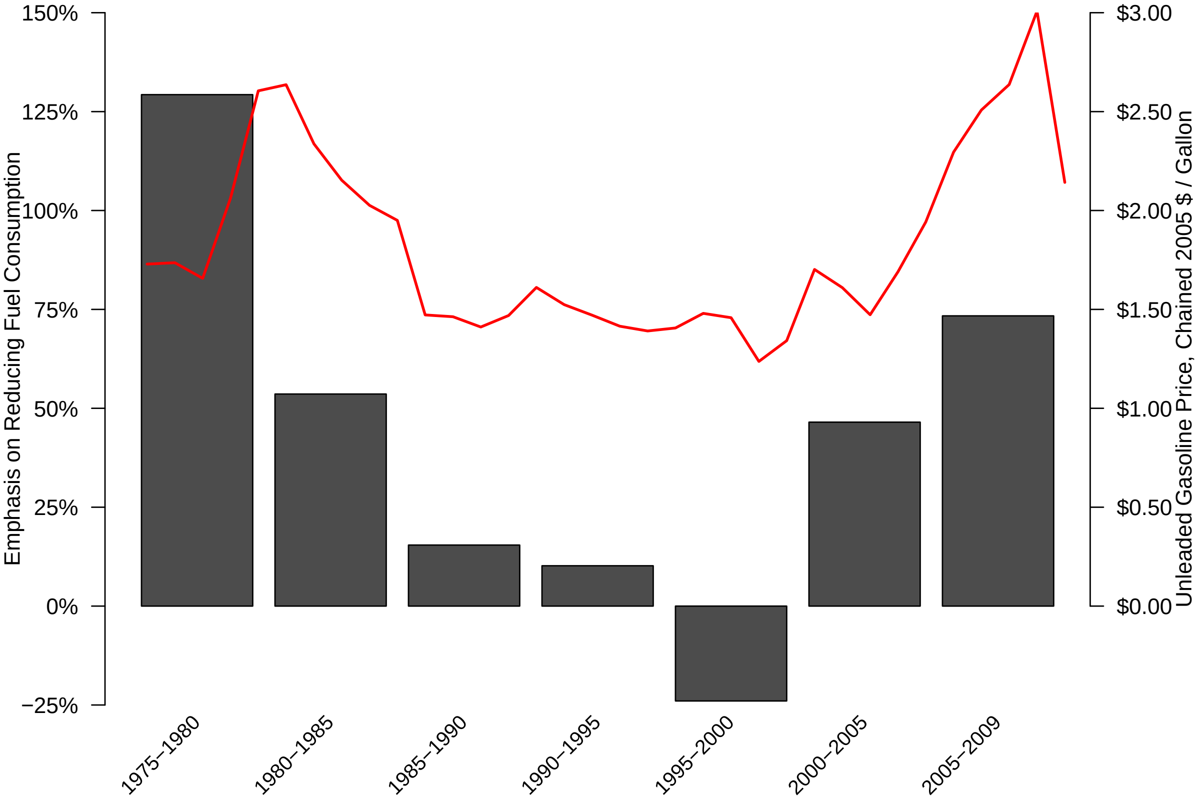

Figure 4-3: Emphasis on reducing fuel consumption among new U.S. cars, 1975–2009. The real price of unleaded gasoline is plotted on the secondary axis.

Whereas ERFC addresses the question, “How much of the potential reduction in fuel consumption has actually been realized?” a related question is “To what ends were improvements in technology applied?” To answer the latter question quantitatively, it is possible to express the changes in acceleration, size, and feature content between two years, t_1 and t_2, in terms of equivalent reductions in fuel consumption. This was done based on the tradeoff coefficients from Equation 4.4, as shown in the following equations:

$$Tech_{Size/Features} = 1 - \left(\frac{IWT_{t_1}}{IWT_{t_1} + \Delta W_{t_1,t_2}} \right)^{\beta_1} \tag{4.7}$$

$$Tech_{Z97} = 1 - \left(\frac{Z97_{t_1}}{Z97_{t_2}}\right)^{\beta_2} \tag{4.8}$$

In the above equations, $Tech_{Size/Features}$ is the fractional reduction in per-mile fuel consumption that could have been realized in lieu of the observed change in weight from greater size or feature content, holding technology constant. IWT_{t_1} is the average inertia weight of new cars in year t_1, and the $\Delta W_{t_1,t_2}$ is the change in weight between t_1 and t_2 that is attributed to either changes in the size mix or changes in feature content, as summarized in Figure 3-6. Similarly, $Tech_{Z97}$ is the fractional reduction in fuel consumption that could have been achieved in lieu of changes in acceleration performance.

Let us consider a concrete example applying Equation 4.7 to the weight of new features. From Figure 3-6, it is evident that the weight of features in the average new car increased by an estimated 135 kg between 1990 and 2009. Starting from a baseline inertia weight of 1,443 kg in 1990, this represents a 9.4% increase in weight. Applying the estimated value of $\beta_1 = 0.686$ from Model 5 in Table 4.1, we find that with the same technology needed to maintain constant per-mile fuel consumption while increasing weight by 9.4%, feature weight could have been maintained and fuel consumption reduced by 6.0%. Thus, the technology required to offset the new features added to cars from 1990–2009 is taken to be the equivalent of a 6.0% reduction in fuel consumption.

Figure 4-4 summarizes the equivalent fuel consumption reductions that were needed to offset chances in acceleration, feature content, and size changes in the average new U.S. car from 1975–1990 and from 1990–2009. Between 1975 and 1990, the average per-mile fuel consumption of new cars decreased by 43%. Over the same period, the average acceleration time decreased by 30%, which "consumed" enough technology to have reduced fuel consumption by 15%. Greater feature content and size had relatively minor effects in this period. Comparing 1975–1990 with 1990–2009, the most striking difference is the large decrease in the actual fuel consumption change between the two periods. Average fuel consumption changed much less over the second period than over the first, but nevertheless still constituted the largest "sink" for technology changes over the second period. In the second period, slightly less technology was dedicated to offsetting faster acceleration times, while offsetting the weight impacts of new features consumed considerably more technology than in the first period (though this was still a smaller technological burden than faster acceleration times and fuel consumption reductions). The weight effects of increased size consumed very little technology in either period, reflecting the fact that net size shifts were relatively small over these periods.[7]

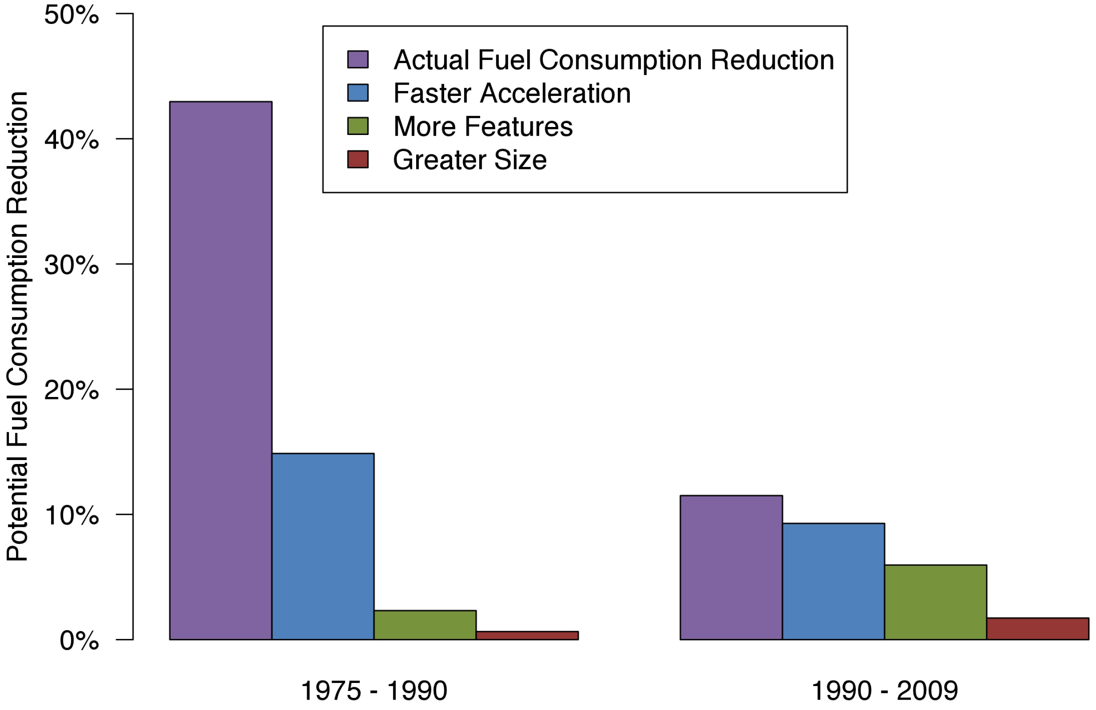

Figure 4-4: Applications of technological improvements to fuel consumption reductions and to offsetting other attribute changes in new cars over two periods.

---

[7] Throughout this chapter, but especially right here, the reader should bear in mind that the scope of the analysis is limited to new cars. Therefore, the size shift embodied in the transition from cars to light trucks is not reflected in the results reported here.

An alternative view of the technology sinks is provided in Figure 4-5. In this figure, the lower edge of the stacked areas represents the potential fuel consumption reduction that could have been achieved if size, acceleration performance, and feature content had remained unchanged at their 1975 levels. Each wedge represents the potential fuel consumption reduction that could have been achieved if a certain attribute had remained at its 1975 level. For example, the red wedge shows that at its peak, offsetting the fuel consumption effects of greater size consumed enough technology to have reduced fuel consumption by about 5% or less. The figure illustrates how offsetting the fuel consumption penalties of faster acceleration has consumed a large and continually growing amount of new efficiency technologies.

4.4.4 Comparison with Other Published Results

In this section, some of the key results reported earlier in the chapter are compared with analogous estimates previously reported by other authors. First, estimates of the tradeoff parameters between fuel consumption, acceleration performance, and weight are considered. Next, the estimates of the potential per-mile fuel consumption reduction since 1975 are compared with other authors' estimates of this quantity.

Weight, Power, and Acceleration

The model specifications that were reported in Table 4.1 indicated that holding acceleration constant, a 1% increase in weight is associated with a 0.59–0.73% increase in fuel consumption. Based on the preferred model specification (Model 5), a 1% increase in weight is expected to cause a 0.69% increase in fuel consumption. These results are consistent with literature, empirical, and simulation results presented by Cheah (2010). Cheah reviewed literature estimates and found estimates ranging from a 2–8% increase in fuel consumption for a 10% increase in weight. Her empirical analysis found that for a weight increase of 10%, fuel consumption of cars increases by about 5.6%, though she did not simultaneously control for other vehicle attributes. Finally, Cheah reported a set of vehicle simulation exercises, which yielded a 6.9% increase

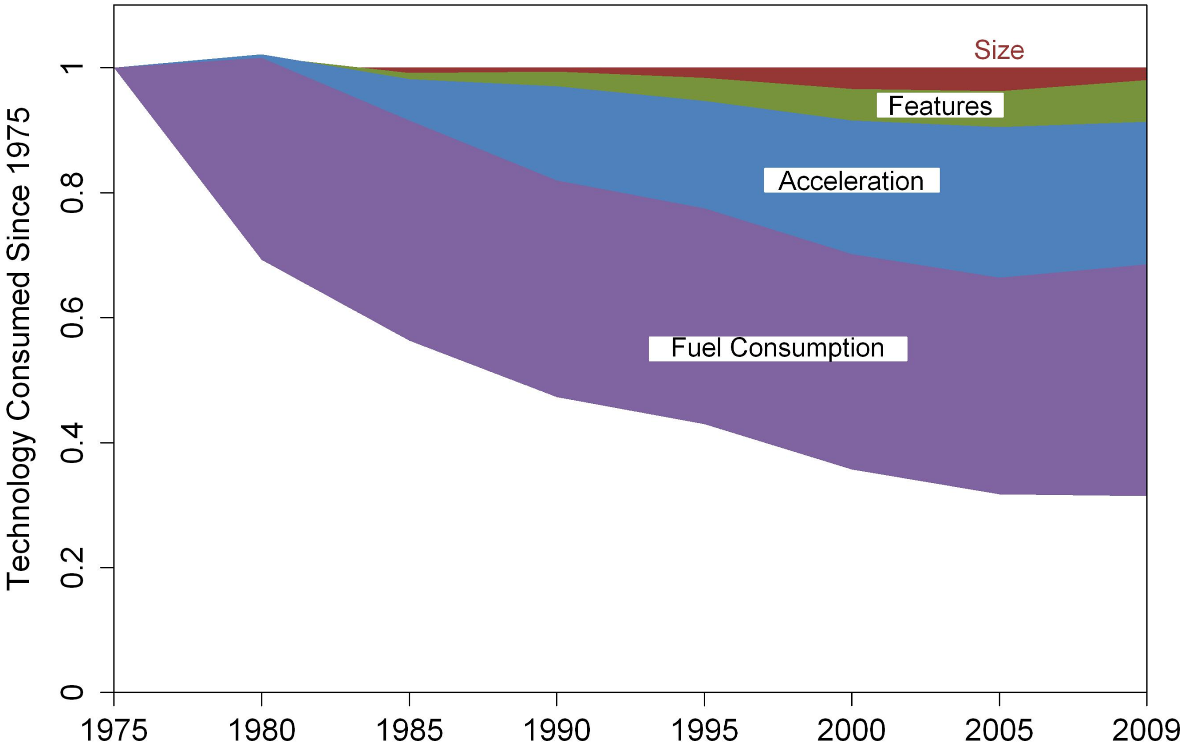

Figure 4-5: Cumulative applications of technological improvements toward major attribute changes in new cars in five-year intervals. Each wedge represents the potential fuel consumption reduction that was dedicated to offsetting the fuel consumption penalties of other attribute changes.

in fuel consumption for a 10% increase in weight, holding acceleration performance constant.

Knittel (2011) estimated that holding power constant, a 1% increase in weight would cause a 0.4% decrease in fuel economy. As discussed in Section 4.4.1, it is not surprising that the sensitivity of fuel consumption to weight is higher when holding acceleration constant (as in the present work) than when holding power constant (as in Knittel's work). Maintaining acceleration performance while increasing weight requires a commensurate increase in power. Thus, the overall effect of a weight increase on fuel consumption includes both the direct effect of greater weight, and the additional effect of increasing power to maintain acceleration performance.

Several papers dating from the early 1990s addressed the tradeoff between weight and fuel consumption. Among these, typical effects of a 10% reduction in weight were a 3% increase in fuel economy at constant power, or a 6.6% increase in fuel economy at constant acceleration performance. Similarly, they used a value of a 0.44% increase in fuel consumption for a 1% decrease in the 0–97 km/h acceleration time (OTA, 1991; DeCicco and Ross, 1993; Greene and Fan, 1994).

More recently, a number of authors have used vehicle simulations to explore the tradeoffs between fuel consumption and power or acceleration performance. Figure 4-6 illustrates the results of several such exercises for midsize U.S. cars, along with the tradeoff calculated in this chapter. The tradeoff identified in this chapter is very similar to that reported by Whitefoot et al. (2011). Compared with the results of Cheah et al. (2009), the present work and the findings of Whitefoot et al. imply a smaller fuel consumption penalty for decreasing acceleration time. The discrepancy between the results of Cheah et al. and the others may be a result of the small number of vehicle simulations carried out by the former. The tradeoff estimated by Shiau et al. falls between the current results and those of Whitefoot et al. on the one hand, and those of Cheah et al. on the other.

Powertrain Technologies

The results that were presented in Table 4.1 also contained estimates of the fuel consumption effects of various powertrain technologies. Holding acceleration and weight constant, a manual transmission was estimated to deliver a 12–14% reduction in fuel consumption in the base year (1975), though this advantage has been declining by about 0.3% per year. These results are similar to those reported by Knittel, though slightly larger. This most likely reflects the fact that manual transmissions offer better acceleration performance than automatics, so the fuel consumption benefit of a manual is greater when controlling for acceleration performance than when controlling for power. Similarly, all-wheel drive or 4-wheel drive was estimated to incur a 8–11% fuel consumption penalty in the base year (in addition to the associated weight

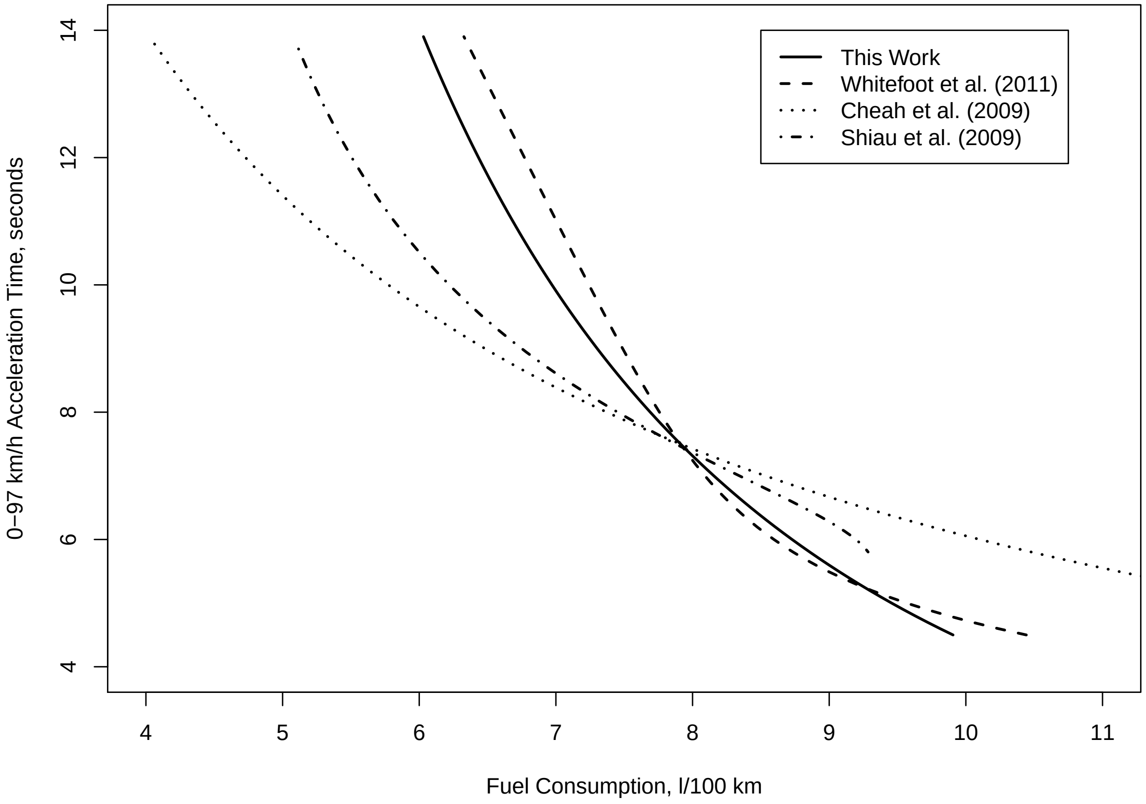

Figure 4-6: Tradeoff between 0–97 km/h acceleration time and combined city/highway fuel consumption. Curves represent tradeoffs as estimated in this work and by other authors (Whitefoot et al., 2011; Cheah et al., 2009; Shiau et al., 2009).

penalty), though this has been declining by an estimated 0.2% per year. The shrinking fuel consumption penalties associated with automatic transmissions and all-wheel drive can be interpreted as an indication that automobile manufacturers are making technological improvements in these particular subsystems.

Naturally aspirated diesels delivered an estimated 18% fuel consumption reduction, and turbodiesels approximately a 25% reduction. This is within the range of 20–30% reported by EPA (2013b), and is similar to the 24% fuel consumption benefit of dieselization reported by Bandivadekar et al. (2008).[8] In contrast, the fuel con-

---

[8] Bandivadekar et al. (2008) reported that current diesel cars offer about a 16% reduction in fuel consumption on an energy-equivalent basis, relative to a naturally aspirated gasoline car. Adjusting for the 10% greater energy content of diesel, this is equivalent to a 24% reduction in fuel consumption

sumption benefits estimated here for boosted gasoline engines (3–5%) are lower than the 10% calculated by Bandivadekar et al. (2008) for current turbocharged engines. This is most likely because boosted gasoline engines tend to have high specific power. Being in the top quintile for engine specific power is estimated to reduce fuel consumption by a further 2–4% relative to the middle quintile, so if boosted gasoline vehicles are all in the top quintile, then a more accurate estimate of their fuel consumption benefit would be 5–9% compared with the average new car. The estimated 30% fuel consumption benefit of hybridization was similar to the results of Bandivadekar et al. (2008), and within the 20–40% range reported by EPA (2013b).

Technological Improvements

In this chapter, estimates have been developed for the rate of technological improvement in new U.S. cars since 1975, expressed as the potential reductions in per-mile fuel consumption that would have been achieved if not for changes in other vehicle attributes, namely acceleration, features, and functionality. Other investigators have addressed this question in the past, using different time periods, different methodologies, and controlling for different vehicle attributes. Table 4.2 summarizes some of these studies and compares their estimates of technology change with the results reported in this chapter.

Each row of Table 4.2 contains a reference to a prior estimate of annual technology improvement for new U.S. cars, expressed as the potential reduction in per-mile fuel consumption if other attributes had remained constant. Also listed are the the particular attributes that were controlled for in each analysis, and the years that were considered. The last column contains the potential reduction in fuel consumption over the period in question as estimated in this chapter. The estimates from this chapter are intended to reflect the potential improvements holding acceleration performance, features, and functionality (including size and comfort) constant.

Most of the estimates in Table 4.2 are lower than the estimates developed in this chapter, reflecting the broader scope of technology improvements encompassed by

---

on a volumetric basis. The EPA fuel economy numbers used in this work are reported on a volumetric

Table 4.2: Annual rates of technology improvement: comparisons between results of this chapter and literature estimates. Each literature study provided estimates of how much fuel consumption would have been reduced if other vehicle attributes had remained constant over time.

Reference	Controlled For	Years	Annual Technology Improvement	
			Cited Work	This Work
Greene and Fan (1994)	hp/wt	1975–1993	3.6%	4.6%
Greene and Fan (1994)	hp/wt, wt	1975–1993	2.8%	4.6%
An and DeCicco (2007)	hp/wt, size	1977–2005	3.2%	3.3%
Knittel (2011)	hp, wt	1980–2006	1.9%	3.0%
Lutsey and Sperling (2005)	size, wt, accel	1987–2004	0.7%	2.3%
EPA (2013b)	size	1977–2009	1.8%	3.3%
EPA (2013b)	wt	1975–2009	1.4%	3.4%
EPA (2013b)	size, wt	1977–2009	1.4%	3.3%

the present work. The technology change estimates reported by An and DeCicco (2007) come closest to the estimates developed here, but as discussed in Section 4.1 there are important disadvantages to An and DeCicco's assumption of a 1:1:1 tradeoff between fuel economy, interior volume, and power/weight ratio. When considering different time periods, An and DeCicco's estimates yield different results than the present work. For example, their method yields 4.2% per year from 1977–1990 and 2.3% per year from 1990–2005. In contrast, the current work indicates a sharper difference in the rate of technology change between these two periods: 4.9% per year from 1977–1990 versus 1.9% per year from 1990–2005.

4.5 Conclusions

In order to assess the potential for fuel consumption reduction, it was necessary to quantify the tradeoffs between vehicle fuel consumption, weight, and acceleration performance. Empirical analysis of cars offered for sale in the U.S. since 1975 yielded an estimate of a 0.69% reduction in fuel consumption for a 1% reduction in inertia weight, which is consistent with values reported in the literature. The effect of a 1% increase in acceleration time was sensitive to model specification, but in the preferred model specification was estimated to cause a 0.44% decrease in fuel consumption, holding all else equal.

This chapter developed a broader view of technology improvements than has been reported in previous studies of automotive technology improvement. Improvements in the fuel efficiency technology of new cars in the U.S. since 1975 have been impressive. The work reported in ths chapter has shown that between 1975 and 2009, per-mile fuel consumption could have been reduced by approximately 70%, or 3.4% per year, if not for reductions in acceleration time and the introduction of new features and functionality to vehicles. However, this progress has not been uniform: improvements averaged 5% per year from 1975–1990, but only 2.1% per year from 1990–2009. These estimates of potential fuel consumption reductions are greater than estimates previously reported in the literature. This is because the present work takes a broader view of technology improvement than have previous investigations, and captures additional sources of improvement. These include improvements in acceleration performance for a given level of weight and power, and weight-saving technologies that enable more features and functionality to be added to a vehicle without increasing weight.

The ends to which technological improvements have been applied has varied over time. In the late 1970s, all of the improvements in car efficiency technology (and then some) were realized as reductions in per-mile fuel consumption. Since that time, the emphasis on reducing fuel consumption declined, and was even negative for a few years in the 1990s as the average fuel consumption of new cars actually increased. In recent years, it has again rebounded, and between 2005 and 2009 about 75% of the potential fuel consumption reduction was realized.

In light of the findings in this chapter, the fuel economy standards recently finalized for 2025 appear to be rather ambitious. Even if features, functionality, and acceleration remain at today's levels, cars will need to sustain average annual improvements of 4.3% per year for 14 years in order to comply with the 2025 standards. This is much higher than has been observed in recent years, though it is within the range of the improvements that were achieved between 1975 and 1990. If automakers hope to further reduce acceleration times, or to increase features or functionality in any way that adds weight to their vehicles, they will need to improve their technology even faster to offset these changes while still meeting the standards. Continued reductions in acceleration times of 10–15%, consistent with the trends identified in Chapter 2, would require an additional 0.3–0.5 percentage point increase in the rate of technology improvement. Continued increases in feature weight of 7 kg/year, consistent with the trend since 1980, would require a further 0.4 percentage point increase in the rate of technology improvement. All told, meeting the 2025 CAFE standard while continuing historical trends in acceleration performance and feature content will require technology improvement of at least 5% per year, slightly exceeding the rates observed between 1975 and 1990.

Chapter 5

Do Automotive Fuel Economy Standards Increase Rates of Technology Change?

In Chapter 4, it was shown that recent rates of improvement in fuel efficiency technology will not be adequate to meet 2025 Corporate Average Fuel Economy (CAFE) standards for new cars, unless sacrifices are made in other vehicle attributes. However, a compelling argument can be made that the presence of CAFE tighter standards may itself cause firms to develop and deploy fuel efficiency technologies more quickly. Indeed, CAFE standards are popularly presented as technology-forcing regulations, which will have precisely this effect. It is not clear, however, that such regulations actually spur more advanced technology; fuel economy gains may also come through sacrificing other vehicle attributes, such as acceleration performance, size, or features. In this chapter I test whether binding CAFE regulations increased the rate of technology deployment within a firm's fleet of automobiles between 1978 and 2008. I build on recent applications of a product characteristics framework to quantify technological change in automobiles. I use fleet and firm-level regulatory compliance data to identify changes in the rate of technology improvements when a manufacturer's fleet was more tightly constrained by a CAFE standard. In a variety of panel regression specifications, I found little to no significant change in the the rate of technology improvement when fleets were more tightly constrained by a CAFE standard. This was the case for both technology improvement among the menu of cars offered for sale, and the sales-weighted mix of all cars sold. The failure to find a significant effect of the standards on technology change does not preclude the possibility of such an effect, and I discuss several limitations to my results in this chapter.

5.1 Introduction

The development and diffusion of technologies that enhance energy efficiency is a topic of recurring public policy interest, as such technologies offer the promise of reducing resource consumption and externality generation without sacrificing the utility realized by consumers. As Jaffe and Stavins (1994) have argued, understanding the effectiveness of different policies for stimulating the diffusion of energy-saving technologies is essential to sound policymaking in this area. To better understand these issues, those authors identified "two inextricably linked questions: What factors influence the rate of adoption of energy-conserving technologies; and what types of public policy can accelerate their diffusion?" In this chapter, I present an empirical analysis of the latter type: testing whether binding automotive fuel economy standards increased the rate of technology change in U.S. cars between 1978 and 2008.

Corporate Average Fuel Economy ("CAFE") standards are the primary policy tool used to curtail petroleum consumption and greenhouse gas emissions from light-duty vehicles (i.e. cars and light trucks) in the U.S. The production and use of fuels for these vehicles accounted for nearly one half of all petroleum consumption and one quarter of greenhouse gas emissions in the U.S. in 2010. CAFE standards require that each fleet of light-duty vehicles sold in the U.S. by each manufacturer in each year meet a minimum average level of fuel economy (sales-weighted, harmonically averaged). Firms have been permitted to bank credits earned through overcompliance, and to borrow credits against promised future overcompliance, for up to three years.[1]

---

[1] Recently, some additional flexibility mechanisms have been introduced to the CAFE program, but were not in effect during the years covered in this analysis.

CAFE standards are commonly characterized in both policy debates and in the literature as "technology-forcing" regulations that accelerate the development and deployment of efficiency-improving technologies; see, for example, Kleit (2004), and NESCAUM (2008). Numerous studies have estimated the effects of CAFE using economic models that are premised on an assumption that the regulations are technology-forcing. To provide just two examples, Fischer et al. (2007) and Kleit (2004) both permit the level of technology adoption to vary in response to prices or regulations. Despite the prevalence — and acknowledged plausibility — of this premise, empirical assessments of regulation-induced technological change in automobiles are scarce.

Past work has affirmed the role of CAFE standards in stimulating fuel economy increases. Greene (1990) concluded that the standards were more important than fuel prices in determining fuel economy levels over the first twelve years of the CAFE program. However, this is not the same as saying that the standards spurred more rapid changes in technology. Fuel economy gains may also have come at the expense of other vehicle attributes, such as acceleration performance (Knittel, 2011). Or, in the parlance of Newell et al. (1999), it is possible that CAFE standards drove changes in the direction, but not the rate, of technological progress in automobiles.

To determine whether standards have affected the rate of technology change, it is necessary to shift away from a focus on fuel economy (which is but one of many vehicle attributes affected by technological change) and toward a focus on some measure of technological change *per se*. One such approach relies on patent counts as a measure of innovation (Popp, 2002). Crabb and Johnson (2010) recently applied this method to energy-saving innovations in U.S. automobiles, concluding that automotive energy efficiency patents are not responsive to CAFE standards. Attempting to square this conclusion with the literature that has found an effect of CAFE, the authors speculate that "Perhaps it is price that drives innovation, as we have shown, and regulation encourages only the final step of adoption by automobile manufacturers." (Crabb and Johnson, 2010) This points to another approach to measuring technology change: one that uses a product characteristics space to measure the diffusion of new technologies.

Newell et al. (1999) measured the rate and direction of technological change in air conditioners and water heaters using the product characteristics approach. They included energy prices and policy stringency as regressors, and concluded that these factors influenced the direction, but not the rate, of technology change in air conditioners. They found no significant effects on either the rate or direction of technological change in gas water heaters. Knittel (2011) recently employed a similar product characteristics approach to estimate the average fleet-wide technology improvements in new automobiles, finding that"Technological progress was most rapid during the early 1980s, a period where CAFE standards were rapidly increasing and gasoline prices were high." Although this suggests that faster rates of technology change are associated with faster rates of standard-tightening (and also higher gas prices), no attempt was made to estimate the effect of CAFE on the rate of technological change, nor to distinguish it from the effect of gasoline prices (Knittel, 2011).

In this chapter, I build on past investigations by extending the product characteristics framework for automobiles to estimate technological changes at the level of individual manufacturers' car fleets. I then use these fleet-specific estimates to test whether being constrained by a CAFE standard has been associated with higher rates of technology change in the constrained fleet. In so doing, I exploit variation in the gap between actual fuel economy and applicable standards, both between firms and within the same firm over multiple years, to identify the magnitude of the technology-forcing effect of a binding CAFE standard.

The chapter proceeds as follows: in the next section, I discuss my empirical strategy and how it is informed by the structure of the CAFE program; in Section III I discuss the data sources used; in Section IV I present results; and in the final section, I offer some conclusions.

5.2 Empirical strategy

Several features of the CAFE program are relevant to my empirical strategy. First, CAFE standards are applied separately to each manufacturer's fleets of domestic passenger cars, imported passenger cars, and light trucks, each of which must independently meet an applicable standard.[2] Light trucks are subject to generally looser standards, and have in the past been subdivided into numerous subcategories for compliance purposes. The time-varying classification of light trucks makes it impractical to match a group of trucks with a certain rate of technology change to the relevant CAFE fleet, and so the present analysis focuses only on cars. As noted earlier, new provisions have been introduced to allow credit trading between fleets, but these were not in place for the historical period considered in this work.

A second important feature of the CAFE program is its penalty structure. Firms that fail to meet a standard must pay a fine proportional to the number of vehicles in the noncompliant fleet and to the number of miles per gallon (mpg) by which that fleet missed the standard. Although there are penalties for failing to meet the standard, there is no bonus for exceeding it (doing so can generate credits to offset future noncompliance, but until recently these credits could not be sold and had to be used within 3 years). As such, we would expect that the shadow value of the CAFE standard, and therefore firms' responses to the standard, would be different when firms are above the standard compared with when they are below it.

Finally, it has been documented by Jacobsen (2012) that different firms behave differently in the face of the CAFE standard. One set of firms (the Detroit Three — GM, Ford, and Chrysler) appears to treat it as a binding constraint. Another set of firms (mainly European) routinely violates the standard and pays the penalty, and thus is largely unaffected by the standard (since the penalty rate is extremely low relative to the cost of technical solutions). A final set of firms (mainly Asian) routinely

---

[2] Many firms produce both a domestic car fleet and an import car fleet, each of which must meet the applicable standard.

exceeds the standards, and so is also unaffected by them. Thus, the shadow value of the standard is very different for these different groups of firms, and we may expect a different response to the standard for different classes of firms. We are especially interested in the response of those firms that treat the standard as a binding constraint, as we may reasonably expect little or no response from the other sets of firms.

5.2.1 Fleet-level rates of technological change

Knittel (2011) applied a product-characteristics framework in order to estimate yearly averages for technological progress in automobiles. Using an adaptation of this methodology, I first generate estimates of the rate of technology change for each fleet in each year. Whereas Knittel estimated fixed effects for each year and interpreted these as a measure of cumulative technological change since his base year, I estimated fixed effects for each combination of firm-fleet-year (e.g. GM's import car fleet in 2002). At the same time, I relaxed Knittel's assumption of constant coefficients on weight and power across all firms, allowing for firm-specific coefficients on these variables. Thus, I modeled the fuel economy of model i in fleet j from firm k in year t as:

$$\ln mpg_{ijkt} = T_{jkt} + \beta_{1k} \ln W_{ijkt} + \beta_{2k} \ln P_{ijkt} + \mathbf{B}\mathbf{X}_{ijkt} + \epsilon_{ijkt} \tag{5.1}$$

Where T is a set of fixed effects for the fleet-firm-year, W is car weight, P is engine peak power, Dsl is a dummy variable indicating whether a car has a diesel engine, Man is a dummy variable indicating whether it has a manual transmission, and ϵ is an i.i.d. random error term.[3] I took the first differences in the fixed effects T_{jkt} in order to obtain my dependent variable: the year-over-year change in the technology deployed within each fleet by each firm, denoted ΔT_{jkt}.

---

[3] This closely resembles Knittel's model specification #3, but with the year fixed effects replaced by firm-fleet-year fixed effects.

$$\Delta T_{jkt} = T_{jk,t+1} - T_{jkt} \tag{5.2}$$

When Equation 5.1 is estimated using an unweighted regression (ordinary least squares), the fixed effects can be interpreted as a measure of the average technological sophistication of the cars offered for sale in a given fleet/firm in a given year. Alternatively, Equation 5.1 can be estimated using weighted least squares, where the weights are equal to the sales of each observed car. In the latter case, the fixed effects may be better interpreted as the average technological sophistication of the cars sold by a given fleet/firm in a given year. I estimated Equation 5.1 using both unweighted and sales-weighted least squares, and estimated the effects of CAFE on both of these measures of technological change.

Over the years and fleets considered in this work, the mean value of the (unweighted) annual change in technology potential was 0.017, with a standard deviation of 0.032. This means that technological improvements could have increased fuel economy of cars offered for sale by an average of 1.7% annually, conditional on power, weight, transmission types, and diesel share remaining unchanged. On a sales-weighted basis, this value was 0.015, with a standard deviation of 0.040. This indicates that the technological improvements, averaged across the mix of cars actually sold, could have increased average fuel economy by an average of 1.5% per year, conditional on power, weight, transmission types, and diesel share remaining unchanged.

5.2.2 Effects of CAFE on technological change

The CAFE penalty system creates an incentive structure in which penalties are proportional to the amount by which a fleet falls short of the standard, but fleets that exceed the standard are neither penalized nor rewarded. Thus, the shadow value of the constraint can be expected to differ when a firm is above the standard compared with when it is below the standard. I define an independent variable which I call *shortfall* as the gap between the current year's actual fuel economy and the next year's required fuel economy:

$$S_{jkt} = (Std_{jk,t+1} - MPG_{jkt}) \tag{5.3}$$

The shortfall therefore captures the stringency of a CAFE constraint, in that it represents the amount by which a fleet would have to increase its fuel economy over the coming year to be compliant with the next year's standard. Since there is a penalty for missing the standard but no reward for exceeding it, I also define a binary treatment variable indicating whether or not the fleet is CAFE-constrained (i.e. whether it has to improve its fuel economy to meet the new standard in the coming year):

$$D_{jkt} = 1\{S_{jkt} > 0\} \tag{5.4}$$

I then flexibly model the effect of the shortfall and the CAFE-constrained indicator on the rate of technology change using a set of panel regression models with fixed effects for fleet, the general form of which is:

$$\Delta T_{jkt} = \gamma P_t + \alpha_0 D_{jkt} + \alpha_1 S_{jkt} + \alpha_2 S_{jkt} D_{jkt} + \mu_{0jk} + \mu_{1jk} t + \mu_{2jk} t^2 + \mathbf{B}\mathbf{X}_{jkt} + \epsilon_{jkt} \tag{5.5}$$

In the above equation, P_t is the price of gasoline in year t;[4] μ_{0jk} is a vector of fixed effects for each fleet; μ_{1jk} and μ_{2jk} allow for fleet-specific time trends; S_{jkt} and D_{jkt} are as defined above; and X_{jkt} is a vector of covariates. I also investigated a specification that included interactions of S_{jkt} and D_{jkt} with a dummy variable for the Detroit Three, allowing for a different response to the standard for this subset of companies.[5]

---

[4] Lagged gasoline prices were also investigated, but returned an inferior fit.

[5] A model specification that included interaction terms between gasoline prices and shortfall was also investigated, but it did not indicate any significant interaction effects.

In some specifications of the model, I also included as covariates average weight, average size (interior volume), average engine peak power, average fuel consumption (gallons per mile), fraction of small cars (mini-compacts, subcompacts, and two-seaters), fraction of cars with 4-wheel drive, fraction of wagons, and fraction of convertibles. The rationale for including these covariates is that the composition of a firm's fleet may plausibly affect the ease with which new technologies can be adopted.

As discussed later, in the section on Identification, the specification in Equation 5.5 may be prone to spurious correlations between increases in CAFE standards and faster technological improvements. These concerns are mitigated by the use of a specification that employs fixed effects for year, as shown below:

$$\Delta T_{jkt} = \delta_t + \alpha_0 D_{jkt} + \alpha_1 S_{jkt} + \alpha_2 S_{jkt} D_{jkt} + \mu_{0jk} + \mu_{1jk} t + \mu_{2jk} t^2 + \mathbf{B}\mathbf{X}_{jkt} + \epsilon_{jkt} \quad (5.6)$$

In the above equation, δ_t is a vector of fixed effects for year, and other terms are as defined previously. Since fuel prices were observed as annual averages, they were omitted when year fixed effects were included.

5.2.3 Standard errors

The number of fleets in the data set was 29, meaning that conventional methods of estimating clustered standard errors were inappropriate. I therefore used a double bootstrapping procedure to estimate the standard errors. To do this, I first resampled the fleets, with replacement. Next, from the data set created by resampling the fleets, I resampled blocks of years, with replacement, as in block bootstrapping. I used non-overlapping blocks with length 3 years. Finally, I estimated the parameters of interest using the resampled data, repeating the above steps 2000 times to simulate a distribution of parameter values.

5.2.4 Identification

It is worth discussing briefly the ability of the panel regression specification to correctly identify the effect of the binding CAFE standard on the rate of technological change. For the panel regression to identify the causal effects of the shortfall S and the state of being constrained by the standard D, we need the error terms to be independent conditional on the fixed effects and the covariates, i.e. $E[\epsilon_{jkt}|P_t, X_{jkt}, S_{jkt}, D_{jkt}, t] = 0$ for all t in the case of Equation 5.5, or $E[\epsilon_{jkt}|\delta_t, X_{jkt}, S_{jkt}, D_{jkt}, t] = 0$ for all t in the case of Equation 5.6. In other words, we need there to be no unobserved confounders that are varying with time.

One possible reason that the panel regression may not correctly identify the causal effects is if the standards were set at levels that the firms deemed feasible. In this case, there could be an omitted variable, "firm's technology plan," that is correlated with increases in observed technology. If firms had influenced the level of the standard, then this technology plan variable could also be correlated with the next year's CAFE standard. An important point here is that the confounding effect of the technology plan variable would have to be varying over time. Nevertheless, it is possible that such an effect could account for some of the correlation between higher standards and faster rates of technology improvement.

The fixed effect specification (Equation 5.6) should be more robust in such a situation. To the extent that the the stringency of regulations is shaped by the capabilities and plans of the industry as a whole, this confounding should be soaked up in the fixed effects. However, CAFE standards have been set, at least some of the time, according to a "least capable manufacturer" heuristic, in which the standards are set so as not to be overly burdensome on any single firm (NHTSA, 2006). This type of process could indeed lead to a situation in which the size of a particular firm's fuel economy shortfall is correlated with its future technology improvements, without causing those improvements.

A further approach to dealing with this comes from the fact that after 1989, CAFE standards for cars did not change for the remainder of the years in this analysis. The standard for cars was constant at 27.5 mpg, which is essentially an arbitrary value. Since we know that technology continued to improve during this period — as evidenced by the continued changes in the year fixed effects from the fuel economy regression — we can conclude that the level of the CAFE standard in these years was not determined by the technology plans of the manufacturers. So, if we re-estimate our model(s) using only the data from 1989 onwards, we can be more confident that the results are not an artifact of CAFE standards being set based on manufacturer technology plans.

5.3 Data

I used two principal data sources in this work. Data on the annual fuel economy performance and applicable CAFE standards came from CAFE compliance reports held by the National Highway Traffic Safety Administration (NHTSA), which administers the CAFE program. These were used to calculate whether a fleet was above or below the next year's standard, and by how much. Another database maintained by NHTSA provided model-level data on car attributes including fuel economy, weight, engine type and power, transmission type, drive type, and body style. I used these data to estimate the rates of technological progress in each year (the outcome of interest) as well as various covariates. I used data from the years 1978–2008, and omitted only manufacturers of limited-volume, specialty vehicles.

5.4 Results

Table 5.1 summarizes the results of the regressions for technological change in cars offered for sale, based on Equation 5.5. These regressions include gasoline prices end exclude year fixed effects. The bootstrapped standard errors are reported in parentheses. If the CAFE standard had a forcing effect on technology, we would expect the coefficients on the shortfall variable, S, and on the shortfall–CAFE-constrained interaction term, $S * D$, to be positive and significant. From columns 1–4 in Table 5.1 it is clear that although the coefficients on the binary (constrained/unconstrained) and continuous treatment intensity (shortfall) variables mostly have the expected sign, none of them are statistically significant. This remained true regardless of the inclusion of firm-specific time trends and covariates, and of the restriction of the data to only those years when standards were arbitrarily fixed at 27.5 MPG.

As noted earlier, we would expect the effects of a CAFE standard to be larger for those firms (the Detroit Three) that treat the standards as a binding constraint. It is possible that the estimated effects in columns 1–4 were dampened by the inclusion of firms that tend to ignore their CAFE obligations. Therefore, I re-estimated the models in columns 1–4 while allowing for a different CAFE response for the Detroit firms. These results are shown in columns 5–8, while column 9 shows the results when only the Detroit firms were included in the data set. The coefficient on shortfall for Detroit firms, when they are constrained by the CAFE standard, was positive and significant in model specification 7. This is in line with predictions of technology-forcing, but the result was not robust to different model specifications and is not compelling on its own.

The coefficients for gasoline price in Table 5.1 are all positive, and in a few cases are marginally significant statistically. However, the estimates are not robust to the different specifications.

Automobile manufacturers generally make product planning decisions several years ahead. Therefore, it is plausible that being constrained by a CAFE standard in one year would not be sufficient to alter the average technological sophistication of the mix of cars offered in the next year. One approach to dealing with this is to recognize that being constrained by the standard might alter the mix of products actually sold, an effect which would show up in the *sales-weighted* measure of technological sophistication. Table 5.2 shows the same set of models as before, but with the sales-weighted technological progress as the dependent variable. The results are somewhat different than in Table 5.1. The estimated coefficients on the shortfall variable are significant in some specifications, though the estimates are not stable across these different specifications and become insignificant when covariates are included.

In contrast to the results in Table 5.1, the coefficients on gasoline price are significant in most specifications, though their specific values are sensitive to the specification. It is also apparent that the coefficient estimates are larger in magnitude than the corresponding estimates in Table 5.1. The same is true of the coefficient estimates for the shortfall variable. This is consistent with the view that the mix of vehicles sold is more sensitive to standards and fuel prices than is the mix of vehicles offered.

The results of the panel regressions that employed fixed effects for year are shown in Table 5.3 (for the mix of vehicles offered for sale) and Table 5.4 (for the mix of vehicles actually sold). As with the regressions that included gasoline price, none of the coefficients on the CAFE variables had a significant effect on the rate of technology change in cars offered for sale. Also, as before, the CAFE shortfall variable had a larger effect on the mix of vehicles sold than on the mix of vehicles offered for sale. In some specifications this effect was statistically significant, though this significance vanished when covariates were included.

Since it is plausible that the mix of product offered would not respond to policy signals over a single year, another approach to accommodating the dynamics of the product cycle is to estimate rates of technological change and stringency of the CAFE constraint over periods longer than one year. Table 5.5 summarizes the results of several regressions like those already presented. However, these results were based on technological change and CAFE shortfall measured over three-year intervals instead of one-year intervals. As with the earlier results, no significant effect of CAFE on the rate of technological change is evident.

Table 5.1: Results for year over year technological change in cars offered for sale (unweighted), without year fixed effects.

	1	2	3	4	5	6	7	8	9
Gas Price	0.0173	0.0276 +	0.0328 +	0.0110	0.0158	0.0278	0.0338 +	0.0097	0.0106
	(0.0108)	(0.0165)	(0.0168)	(0.0838)	(0.0102)	(0.0177)	(0.018)	(0.0369)	0.0360
CAFE Constrained	0.0016	0.0008	0.0051	0.0047	-0.0082	0.0026	0.0030	-0.0069	-
	(0.0084)	(0.0103)	(0.0109)	(0.0139)	(0.0141)	(0.0191)	(0.0406)	(0.0165)	-
CAFE Shortfall	0.0017	0.0025	0.0010	-0.0005	0.0008	0.0036	0.0012	0.0002	-
	(0.0015)	(0.0025)	(0.0033)	(0.0080)	(0.0013)	(0.0030)	(0.0035)	(0.0093)	-
Constrained * Shortfall	0.0173	0.0053	0.0010	-0.0014	0.0028	0.0045	-0.0011	-0.0043	-
	(0.0035)	(0.0057)	(0.0064)	(0.0150)	(0.0054)	(0.0143)	(0.0113)	(0.0234)	-
Detroit * CAFE Constrained	-	-	-	-	0.0087	0.0008	0.0007	0.0126	0.0064
	-	-	-	-	(0.0179)	(0.0177)	(0.018)	(0.0569)	(0.0421)
Detroit * CAFE Shortfall	-	-	-	-	0.0027	-0.0033	-0.0020	-0.0033	0.0032
	-	-	-	-	(0.0035)	(0.0191)	(0.0406)	(0.0124)	(0.0086)
Detroit * Constrained * Shortfall	-	-	-	-	0.0001	0.0030	0.0083 *	0.012	0.0044
	-	-	-	-	(0.0143)	(0.003)	(0.0035)	(0.2084)	(0.0322)
Firms Included	All	All	All	All	All	All	All	All	Detroit Only
Years Included	All	All	All	1989 on	All	All	All	1989 on	All
Firm-Specific Time Trends	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Covariates Included	No	No	Yes	Yes	No	No	Yes	Yes	Yes

+ Significant at the 0.1 level * Significant at the 0.05 level

Table 5.2: Results for year over year technological change in mix of cars sold (sales-weighted), without year fixed effects.

	1	2	3	4	5	6	7	8	9
Gas Price	0.0253 *	0.0366 +	0.0453 *	0.0140	0.0238 *	0.0367 +	0.0469 *	0.0127	0.0414
	(0.0123)	(0.0189)	(0.0209)	(0.0385)	(0.0114)	(0.0201)	(0.0212)	(0.0389)	0.0371
CAFE Constrained	-0.0010	-0.0041	0.0029	0.0080	-0.0137	-0.0044	-0.0028	-0.0012	-
	(0.011)	(0.0131)	(0.0124)	(0.0291)	(0.0156)	(0.0186)	(0.0348)	(0.0147)	-
CAFE Shortfall	0.0032 +	0.0070 *	0.0051	0.0040	0.0027	0.0086 *	0.0054	0.0051	-
	(0.0019)	(0.0034)	(0.0048)	(0.0103)	(0.0019)	(0.0039)	(0.0052)	(0.0105)	-
Constrained * Shortfall	0.0020	0.0035	-0.0031	-0.0020	0.0029	0.0020	-0.0069	-0.0059	-
	(0.0049)	(0.0074)	(0.0073)	(0.0383)	(0.0074)	(0.0085)	(0.0105)	(0.0198)	-
Detroit * CAFE Constrained	-	-	-	-	0.0129	0.0048	0.0034	0.0082	-0.0006
	-	-	-	-	(0.0314)	(0.0201)	(0.0052)	(0.7689)	(0.021)
Detroit * CAFE Shortfall	-	-	-	-	0.0015	-0.0052	-0.0033	-0.0049	0.0026
	-	-	-	-	(0.0033)	(0.0186)	(0.0212)	(0.0084)	(0.0088)
Detroit * Constrained * Shortfall	-	-	-	-	0.0032	0.0061	0.0154	0.0167	0.0072
	-	-	-	-	(0.0222)	(0.0039)	(0.0348)	(0.9591)	(0.0207)
Firms Included	All	All	All	All	All	All	All	All	Detroit Only
Years Included	All	All	All	1989 on	All	All	All	1989 on	All
Firm-Specific Time Trends	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Covariates Included	No	No	Yes	Yes	No	No	Yes	Yes	Yes

+ Significant at the 0.1 level * Significant at the 0.05 level

Table 5.3: Results for year over year technological change in cars offered for sale (unweighted), including year fixed effects.

	1	2	3	4	5	6	7	8	9
CAFE Constrained	0.0035	0.0046	0.0068	0.0058	-0.0012	0.0081	0.0059	-0.0091	-
	(0.0086)	(0.0113)	(0.0116)	(0.0167)	(0.0142)	(0.0233)	(0.0286)	(0.0235)	-
CAFE Shortfall	0.0019	0.0022	0.0030	-0.0001	0.0012	0.0038	0.0034	0.0009	-
	(0.0015)	(0.0023)	(0.0049)	(0.0085)	(0.0013)	(0.0027)	(0.0053)	(0.0092)	-
Constrained * Shortfall	0.0008	0.0058	0.0042	-0.0024	0.0012	0.0042	0.0028	-0.0054	-
	(0.0037)	(0.0052)	(0.0072)	(0.0283)	(0.0089)	(0.0233)	(0.0123)	(0.0250)	-
Detroit * CAFE Constrained	-	-	-	-	0.0008	-0.0017	0.002	0.0196	0.0103
	-	-	-	-	(0.0202)	(0.0027)	(0.0056)	(0.0570)	(0.0555)
Detroit * CAFE Shortfall	-	-	-	-	0.0022	-0.0046	-0.0025	-0.0046	0.0118
	-	-	-	-	(0.0038)	(0.0233)	(0.0286)	(0.0119)	(0.0291)
Detroit * Constrained * Shortfall	-	-	-	-	0.0019	0.0065*	0.0046	0.0109	0.0115
	-	-	-	-	(0.0168)	(0.0027)	(0.0053)	(0.0684)	(0.0679)
Firms Included	All	All	All	All	All	All	All	All	Detroit Only
Years Included	All	All	All	1989 on	All	All	All	1989 on	All
Firm-Specific Time Trends	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Covariates Included	No	No	Yes	Yes	No	No	Yes	Yes	Yes

+ Significant at the 0.1 level * Significant at the 0.05 level

Table 5.4: Results for year over year technological change in mix of cars sold (sales-weighted), including year fixed effects.

	1	2	3	4	5	6	7	8	9
CAFE Constrained	0.0001	-0.0003	0.0036	0.0078	-0.0058	0.0021	-0.0004	-0.0036	-
	(0.0105)	(0.0132)	(0.0126)	(0.1567)	(0.0165)	(0.0195)	(0.0372)	(0.0176)	-
CAFE Shortfall	0.0036*	0.0066*	0.0067	0.0042	0.0032 +	0.0088 *	0.0072	0.0055	-
	(0.0018)	(0.0033)	(0.0068)	(0.0161)	(0.0018)	(0.0037)	(0.0068)	(0.0102)	-
Constrained * Shortfall	0.0008	0.0041	0.0000	-0.0027	0.0011	0.0023	-0.0027	-0.0057	-
	(0.0049)	(0.0067)	(0.0078)	(0.4362)	(0.0093)	(0.0146)	(0.0125)	(0.0202)	-
Detroit * CAFE Constrained	-	-	-	-	0.0042	0.0024	0.0061	0.0154	0.0063
	-	-	-	-	(0.1097)	(0.0057)	(0.0372)	(0.0875)	(0.035)3
Detroit * CAFE Shortfall	-	-	-	-	0.001	-0.0065	-0.004	-0.0054	0.0118
	-	-	-	-	(0.0032)	(0.0195)	(0.0068)	(0.0100)	(0.0201)
Detroit * Constrained * Shortfall	-	-	-	-	0.0031	0.0075 *	0.0091	0.0116	0.0117
	-	-	-	-	(0.0738)	(0.0037)	(0.0125)	(0.0777)	(0.0270)
Firms Included	All	All	All	All	All	All	All	All	Detroit Only
Years Included	All	All	All	1989 on	All	All	All	1989 on	All
Firm-Specific Time Trends	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Covariates Included	No	No	Yes	Yes	No	No	Yes	Yes	Yes

+ Significant at the 0.1 level * Significant at the 0.05 level

Table 5.5: Results for technological change in cars offered for sale (unweighted), over three-year intervals.

	1	2	5	6
CAFE Constrained	0.0017	-0.0123	-0.0085	-0.0421
	(0.0299)	(0.0675)	(0.0524)	(0.0503)
CAFE Shortfall	0.0021	0.0053	0.0017	0.0045
	(0.0034)	(0.0167)	(0.0040)	(0.0503)
Constrained * Shortfall	0.0050	0.0110	0.0126	0.0187
	(0.0112)	(0.0675)	(0.0207)	(0.0503)
Detroit * CAFE Constrained	-	-	0.0228	0.0493
	-	-	(0.0816)	(0.0503)
Detroit * CAFE Shortfall	-	-	0.0072	0.0035
	-	-	(0.0121)	(0.0503)
Detroit * Constrained * Shortfall	-	-	-0.0311	-0.0303
	-	-	(0.0429)	(0.0503)
Firms Included	All	All	All	All
Years Included	All	All	All	All
Firm-Specific Time Trends	No	Yes	No	Yes
Covariates Included	No	No	No	No

+ Significant at the 0.1 level * Significant at the 0.05 level

5.5 Conclusions

In this work, I found little to no evidence that either the state of being constrained by a CAFE standard in a given year, or the stringency of that constraint, has increased the historic rates of technological change by automobile manufacturers in the U.S. since 1978. In the few cases where significance was found, the results were not robust to alternative model specifications. This held true when the rate of technology change was defined in terms of (1) the year to year change in the technology of cars offered for sale, (2) the year to year change in the sales-weighted average technology level, and (3) the change in technology of cars offered for sale over three-year intervals. However, these results must not be miscontrued as definitively ruling out an effect of CAFE on technology change, and several caveats to this work are discussed below.

This work does suggest that the price of gasoline may have increased the rate of technology change among the mix of vehicles offered for sale. When considering the sales-weighted mix of vehicles actually sold, the evidence is more convincing, but still mixed. Gasoline prices were associated with significant increases in the sales-weighted rate of technological improvement, but the estimates of this effect ranged between 1.3% and 4.7%, depending on the model specification, and were not significant in all cases.

Although little prior empirical work has been reported on the effects of CAFE standards on rates of technology change in automobiles, the negative results reported here do run counter to a good deal of well-grounded microeconomic modeling. This modeling work (for example, Kleit (2004), Fischer et al. (2007), Shiau et al. (2009), as well as the work reported in Chapter 6 of this document) generally predicts that when a firm has the ability to change its adoption of technology, the most profitable response to a binding CAFE standard will involve at least some increase in the adoption of technology.

The gap between the theoretical prediction of faster technology deployment in response to standards and the failure to find such an effect here warrants some discussion. One possible explanation for this discrepancy is that the cost of accelerating technology deployment is higher than commonly thought. It may be that although new technologies can be added at modest cost over the long run, manufacturers are constrained by their R&D pipelines, knowledge bases, and supply chains over the shorter term. As Zoepf and Heywood (2012) has shown, there are limits to how quickly new technologies have been deployed across the vehicle fleet. However, such an explanation begs the question as to why a technology response is detected for gasoline prices but not for CAFE constraints. One possibility is that firms respond differently to the increase in demand for fuel economy resulting from higher fuel prices than to the shadow cost of the fuel economy standard. However, it is not clear why this should be the case.

A second possible explanation for the discrepancy between theory and observation is that although the results were not statistically significant, the signs of the coefficients were in most cases consistent with a small positive effect of CAFE on the pace of technology change. It is possible that a larger sample could reveal a significant effect of CAFE, but we are already using all of the available years and major manufacturers for which data are available.

Thirdly, the failure of the model to show a significant effect of the standards on technology change certainly does not preclude the possibility of such an effect, especially if such an effect operates on multi-year timescales. During the period covered in this analysis, manufacturers had the ability to bank and borrow credits if they over- or under-complied with CAFE in a given year. Moreover, firms cannot introduce new products overnight, and typically redesign vehicle models on a 4–7 year cycle that includes several years of design leadtime before product launch. These factors would be expected to dampen the firm's response to a CAFE shortfall in any single year.

Finally, conclusions based on historical CAFE standards (which were last increased in the 1980s) may not be applicable for predicting responses to future CAFE standards. This work has by necessity been entirely retrospective, and there are several reasons that future responses to CAFE standards may differ from obsreved past responses. First, the CAFE standards considered in this work were uniform, with each car fleet was required to meet the same standard. Adjusting prices and production to sell more small cars was a viable strategy for meeting the standards. Going forward, standards are size-based, with manufacturers of smaller cars being held to tighter fuel economy standards. To a first approximation, this is expected to remove downsizing as a viable compliance strategy. Without the option of downsizing, firms may have to rely more heavily on other compliance strategies, including the deployment of more advanced technologies. Second, modern vehicles are more durable than vehicles in the past, and there are more full-line manufacturers today than when the Big Three ruled the U.S. market in the 1970s. These changes may increase the competitive pressures on manufacturers to deliver attributes that consumers desire, in order to avoid losing customers to competitors or to the used car market.

Notwithstanding the above caveats, it is still interesting that an effect of binding CAFE standards on the historic rate of technology change was not found. Taken with the results of Greene (1990), which found that CAFE significantly influenced fuel economy levels, the results reported here suggest a pattern for cars similar to that reported by Newell et al. (1999) for air conditioners: namely, that standards affect the direction, but not the rate, of technological change. Opportunities remain to apply the product characteristics framework to automobiles in order to test this hypothesis more thoroughly.

Chapter 6

A Simple Model of Manufacturer Responses to Nested Fuel Consumption Standards

This chapter develops a framework for modeling automotive firms' technology adoption and product characteristics decisions in an environment of multiple "nested" regulatory constraints and non-negligible costs to firms of diversifying their product portfolios.

A simple model of an automobile market is constrained by a binding minimum fuel economy standard. Firms compete with one another, seeking to maximize their profits. Variable costs of vehicle manufacturing are assumed to depend on the level of efficiency technology used on the vehicle and the vehicle's acceleration capability, and firms' costs are assumed to be symmetric. Consumer choices of vehicles are modeled using a nested logit model. It is possible for a subset of the market to impose a more stringent fuel economy standard than is required across the broader market, creating a "nested" standard. In such a situation, manufacturers must choose whether to build to the more stringent standard everywhere, or to build specific variants for each market. Building multiple variants is assumed to increase the fixed costs of product development, manufacturing, and certification. The model is currently implemented for firms competing in a single period, within a single vehicle class. The results demonstrate why it may be sometimes — but only sometimes — rational for firms to design specific variants for a nested market, but to also prefer that the larger market harmonize its standards with those of the nested market.

6.1 Nested Fuel Consumption Regulations in the United States

The process for establishing fuel economy standards in the United States has become considerably more intriguing in recent years. Once the strict purview of the federal Department of Transportation and its subsidiary, the National Highway Traffic Safety Administration (NHTSA), today fuel economy policy is also being influenced by the U.S. Environmental Protection Agency (EPA) and the California Air Resources Board (CARB). Currently, by agreement, standards promulgated by all three of these agencies are consistent with one another. However, it remains a possibility that CARB could once again break from the federal government and attempt to establish its own standards, which would create a more strictly regulated market "nested" in the larger federal market. Were such a situation to occur, the implications for consumers, automobile manufacturers, and the environment could be substantial. These outcomes would depend on the structures and relative stringency of the state and federal regulatory regimes, the relative size of the nested state market, and the costs to automobile manufacturers of bringing to market special variants[1] of their vehicles which are intended for the nested market only. In this chapter, I present a framework for modeling the response of automobile manufacturers and consumers to nested state and federal fuel economy standards. I implement this framework using illustrative parameter values for ten firms competing in a single vehicle class. The results demonstrate why sometimes — but only sometimes — it may be rational for firms to design specific variants for a nested market, while preferring that the larger market adopt the more stringent standards of the nested market.

---

[1] Throughout this chapter, I will refer to firms producing one or more *variants* of a vehicle model. I use this to refer to different versions of the same basic model, which differ in their fuel consumption, acceleration performance, or application of technology. For example, variants of the Toyota Camry include those with the 4-cylinder, 6-cylinder, and hybrid electric powertrains.

6.1.1 The Road to Nested Fuel Consumption Regulations

Automotive fuel economy standards have been a part of U.S. energy policy since 1978, after being mandated by the Energy Policy and Conservation Act of 1975 (EPCA). The Act assigned the authority to set CAFE standards to the Secretary of Transportation, who has delegated it to NHTSA. Although EPCA specified a minimum level of 27.5 mpg as the fuel economy standard for cars, NHTSA had the legal responsibility to set higher standards if such standards were feasible. However, standards for cars were not increased above 27.5 mpg, and legislative attempts to raise them were unsuccessful, leaving them unchanged throughout the 1990s. Figure 6-1 shows the historic levels of the CAFE standard for cars, along with the actual average fuel consumption of new cars.

In 2002, the state government of California passed Assembly Bill No. 1493, which required CARB to promulgate standards for the maximum permissible levels of per-mile greenhouse gas (GHG) emissions of vehicles sold in the state. Under §209 of the federal Clean Air Act (42 USC §7543), California has the right to set its own emissions standards for motor vehicles, as long as those standards are at least as stringent as any comparable federal standards, are not set arbitrarily, and are necessary to meet "compelling and extraordinary conditions" in the state. In addition, such standards must be consistent with the regulatory authority laid out in §202 of the Clean Air Act (42 USC §7521), which includes a requirement that the emissions in question "endanger public health or welfare." Under §177 of the Clean Air Act (42 USC §7507), other states can opt to follow California's standards instead of the corresponding Federal standards. By 2009, 13 states and the District of Columbia had opted into California's "Pavley standards," (named after the sponsor of A.B. 1493), meaning that approximately 40% of new vehicles sold in the U.S. would be subject to these tighter standards (EPA, 2009b).

Even as California and a number of other states were pushing ahead with automotive GHG standards, automobile manufacturers and dealers sued to block those regulations, arguing that they were pre-empted by federal CAFE law. Specifically, 49

Figure 6-1: Corporate Average Fuel Economy standards (expressed as fuel consumption) for new cars, 1978–2025, and average fuel consumption of new U.S. cars, 1975–2011.

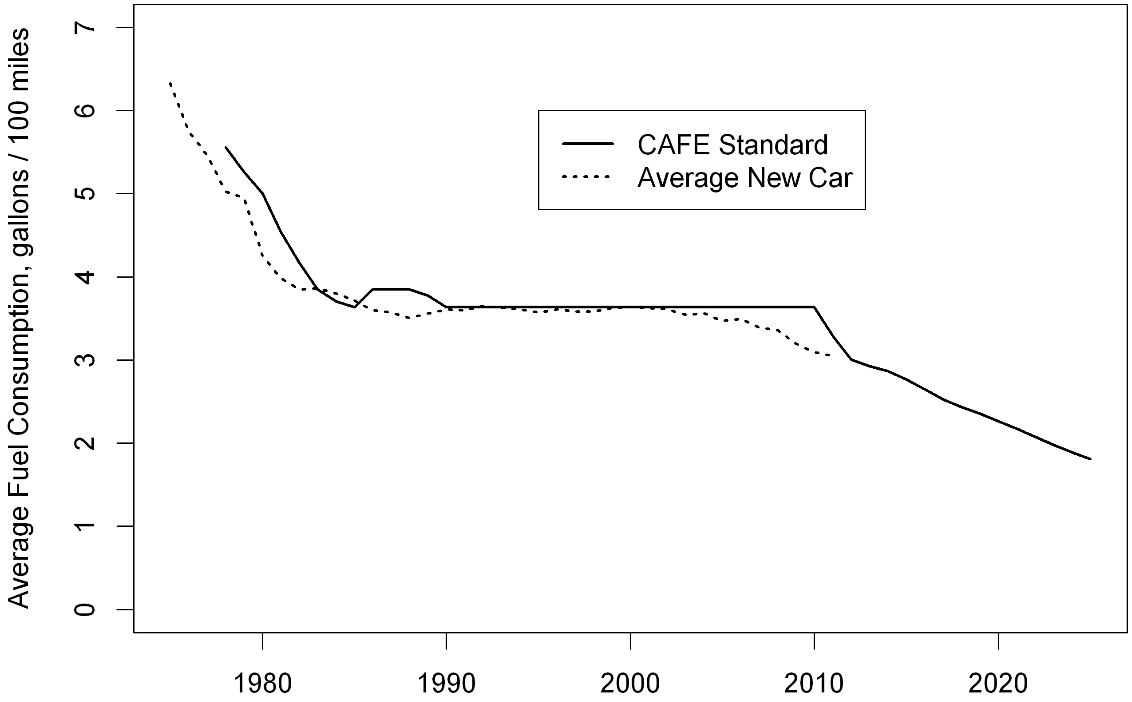

USC §32919 states that "When an average fuel economy standard prescribed under this chapter is in effect, a State or a political subdivision of a State may not adopt or enforce a law or regulation related to fuel economy standards or average fuel economy standards for automobiles covered by an average fuel economy standard under this chapter." Since a large majority of automotive GHG emissions come directly from the combustion of fuel, the argument went, regulating GHG emissions was tantamount to regulating fuel economy and was therefore pre-empted.

The pre-emption issue was still under dispute in 2009, when the Obama administration took office. The new administration granted California the necessary waiver under §209 (EPA, 2009a), while simultaneously brokering a deal between various stakeholders that essentially tightened the federal fuel economy rules to be as stringent as the California GHG standards. The terms of this deal were laid out in commitment letters from major automobile manufacturers, their trade associations, and CARB, as documented on the EPA's website (EPA, 2013c). Key commitments from the auto industry included the following:

- A stay of all pending litigation challenging California's regulations, including challenges claiming pre-emption of state GHG regulations under EPCA;
- A commitment not to contest EPA's granting of a waiver to California under §209 of the Clean Air Act;
- A commitment not to contest CAFE or GHG standards promulgated by NHTSA or EPA, respectively, if those standards were substantially similar to those sketched out in the May, 2009 Notice of Intent to propose rules;
- A commitment not to renew or initiate such litigation in relation to vehicles from model years 2009-2016.

In return, the auto industry obtained the following commitments:

- EPA and NHTSA would propose national GHG and CAFE standards substantially similar to those outlined in the May, 2009 Notice of Intent;
- For model years 2009–2011, California would revise its automotive GHG standards such that they would be applied to the combined fleet of vehicles sold in California and other states adopting California's standards pursuant to §177 of the Clean Air Act;
- For model years 2012–2016, California would revise its standards so that any company in compliance with the federal standards would be deemed to be in compliance with the California standards;

- An assurance that federal CAFE testing procedures could be used to demonstrate compliance with California's standards.

This agreement effectively created a ceasefire, whereby disputes over California's authority to regulate GHG emissions were suspended but not actually resolved. An additional deal with broadly similar terms and covering model years 2017–2025 was subsequently announced in 2011 (EPA, 2013c). However, the 2017–2025 rulemaking included provisions for a mid-term review of the feasibility of the standards in for 2022–2025, citing the long lead time and associated uncertainty.[2] Both California and the automobile manufacturers reserved the right to contest final rules based on the outcome of this mid-term review.

Nested state and federal GHG standards remain an important topic. While it is difficult to say how the courts would have resolved the pre-emption issue, it is certainly possible that if not for this deal, we might now be living in a world where there were two sets of standards: a more lax federal standard and a more stringent standard with the subset of the states following California's standards. In addition, nested standards could return following the outcome of the mid-term review, or upon the expiry of the current deal in 2025. Goulder and Stavins (2011) have observed that "The coexistence of state and federal programs is likely to continue in the context of US climate change policy." Referring specifically to the Pavley initiative, they concluded that given uncertain federal and state rulemakings, "the leakage issue remains very much alive."

6.1.2 Emissions Leakage and Nested Regulations

A critical issue in the context of nested regulations is that of emissions "leakage." Leakage refers to an outcome that can occur when (1) a binding federal standard is in force, (2) a nested subset of the federal market ("adopting states") imposes a more stringent standard, and (3) sales of products in the adopting states are pooled with sales in non-adopting states to determine compliance with the federal standard.

---

[2] Moreover, under 49 USC §32902(b)(3)(B), NHTSA can set CAFE standards for no more than five years at a time, and thus is legally obligated to conduct a *de novo* rulemaking for the years 2022–2025.

In such circumstances, the emission reductions required in the adopting states can "leak" out by enabling increased sales of higher-emitting products in the non-adopting states. In such a situation, the costs of compliance with the tighter standards would be borne by parties in the adopting states, the costs of complying with the federal standard would be reduced for parties in the non-adopting states, and some or all of the emission reductions required in the adopting states would be lost to leakage. Moreover, the overall cost of compliance will be higher than with a uniform federal standard designed to achieve the same overall level of emission reductions (Goulder and Stavins, 2011).

Economic analysis specifically addressing the Pavley law and federal CAFE standards suggests that leakage rates could be quite high. Goulder and Stavins (2011) argue that as much as 100% of the Pavley emissions reductions could leak out of the adopting states, if all regulated firms were already constrained by the federal CAFE standard. Perhaps taken for granted by those authors, an additional necessary condition for 100% leakage is that all firms operating in the non-adopting states were also operating in the adopting states, or that credit trading between firms be permitted.

Goulder et al. (2012) simulated the interactions between the Pavley standards and the federal CAFE standards. They estimated that about 74% of the Pavley emissions reductions through 2016, and 65% of emissions reductions through 2025, could leak out to non-adopting states. They acknowledged the potential significance of credit trading between fleets but did not incorporate it into their model. Their modeling approach is summarized as follows:

- Seven firms each seek to maximize their own profits in Bertrand competition, making pricing and product attribute decisions while accounting for regulatory constraints and the effects of their decisions on costs and demand.
- Each firm offers four vehicle models: small and large cars, and small and large trucks.
- Each firm offers two variants of each model: one in adopting states and one in non-adopting states. There are eight variants in total.

- Firms improve fuel economy through the adoption of more advanced technologies, which can be either "static" or "dynamic."

- Dynamic technology changes are represented as incurring a fixed cost per model, and are assumed to apply to all variants of a model. Each firm has four dynamic technology decision variables: one for each of their models.

- Static technology changes are represented as increasing the variable cost of production, and may differ between the two variants of a model. Each firm has eight static technology decision variables: one for each of their variants.

The distinction between the static and dynamic technology variables is important, and deserves clarification. Goulder et al. (2012) listed "improved aerodynamics and certain improvements in engine design" as examples of dynamic innovations. They assume that such innovations would, for zero variable cost, benefit all variants of a model once "purchased" for the model line as a whole. This represents technologies for which the costs are heavily concentrated in the design and engineering phases of product development. This is plausible for something like aerodynamics: once the body of the vehicle is desiged, all variants benefit from the resulting improvement in efficiency. In contrast, they identified tires, low-friction lubricants, and improved transmissions as static innovations, in the sense that applying these technologies essentially increases the variable cost of vehicle production. This distinction is captured in the objective function as the authors have presented it.

However, Goulder et al. (2012) may have conflated another important concept with their distinction between static and dynamic technology change. They elsewhere distinguished between "static substitutions of car features involving known technologies (e.g. substituting smaller engines for larger ones)," and "dynamic technological progress (which improves the fuel economy associated with a given set of car features)." They noted that their goal was to "contrast moving along a given technological frontier (static substitution) with moving the technological frontier itself (dynamic technological progress)." This latter distinction is more akin to accounting for the tradeoffs between vehicle attributes such as fuel consumption and acceleration performance.

6.1.3 Gaps in Assessments of Pavley/CAFE Leakage

The existing literature addressing interactions between the Pavley standards and the federal CAFE standards has not addressed a number of factors that may affect the outcomes resulting from nested standards.

First, work reported in the literature has not accounted for the effects of different structures of regulations. Federal CAFE standards are now size-based, with each firm's fleet facing a tailored fuel economy standard based on the fleet's particular sales-weighted mix of vehicle sizes. Firms that sell more small vehicles generally need to meet a higher CAFE standard. As such, if firms complied with California's standards by selling more small cars in the adopting states, this would drive up both their nationwide average fuel economy and their federal CAFE requirement. This would be expected to diminish the leakage to non-adopting states. However, if firms already had two variants of each model in production, they could reduce prices and increase sales of the less fuel-consuming variant in the adopting states, while increasing volumes of the more fuel-thirsty variant in the non-adopting states. In this case, leakage could occur easily. Thus, actual leakage will depend on the relative stringency of the different sets of standards, as well as the differences in per-mile fuel consumption among existing variants.

A second deficiency of existing work is that it does not address the possibility that offering a different set of products in adopting and non-adopting states would incur costs for the manufacturers. In their analysis of interactions between the Pavley standards and CAFE, Goulder et al. (2012) take it as given that separate variants of each model will be offered in adopting and non-adopting states. In reality, however, firms would likely be able to save on engineering, testing/certification, and supply chain/logistics costs if they standardize each model. This could reduce leakage by creating an incentive for firms to build a single fleet of vehicles, compliant in the adopting states, and sell it in non-adopting states as well. In this way, the tighter standards in the adopting states could potentially function as a *de facto* increase in the federal standard.

The costs of "splitting" a model line into multiple variants for adopting and non-adopting states is also relevant to the political economy of regulatory competition. Goulder and Stavins (2011) noted that, "there is broad agreement that the California-led state-level tightening of greenhouse-gas-per mile standards brought about the subsequent tightening of federal CAFE standards. Automakers did not wish to face different standards at the federal and state level. Hence they were willing to support tighter federal standards so long as the state standards were removed." Therefore, rather than taking federal standards as exogenous, the analysis of the likely leakage effects of a nested standard ideally would consider the effects of such a standard on industry support for tighter standards at the federal level.

6.1.4 Political Economy of Nested Regulations

Tighter standards in a nested market can create incentives for firms to support tighter standards across a broader market as well. Such an effect is believed to be responsible for the auto industry's support of the federal fuel economy and GHG program championed by the Obama administration. As Goulder and Stavins (2011) noted, "There is, in fact, a considerable history of California air standards having precisely this effect on federal policy developments, because industry is reluctant to face different standards in different parts of the country. For example, the California-led state-level tightening of greenhouse-gas-per mile standards helped bring about the subsequent tightening of federal CAFE standards." Indeed, this would hardly be the first time that California's automotive emissions policies had driven policy at the federal level.

Vogel (1995) used the term "California Effect" to describe the process whereby stricter environmental standards in one jurisdiction motivate industry support for other jurisdictions to adopt those same standards. The eponymous example of this phenomenon is found in the repeated pattern of California tightening its automotive emission standards (historically this has been for criteria pollutants, e.g. NO_X and VOCs, rather than for GHGs). Following each such tightening, the federal government

— with the blessing of the auto industry — has essentially adopted the California standards nationwide.

Murphy (2004b) developed some generalized principles about conditions that tend to lead to the California effect, or what he calls higher-commom-denominator (HCD) outcomes. His first key variable is the type of regulation: market access regulations as opposed to process regulations. Regulations restricting access to a market (such as vehicle technology or performance standards) tend to increase industry support for HCD outcomes, since firms prefer not to make different products for each market. Jurisdictions representing larger shares of the market tend to have a greater ability to motivate broader adoption of tighter standards, thanks to their greater "market power" (Murphy, 2004b; Vogel and Kagan, 2004). In addition, the strength of this effect is influenced by the degree of market concentration on the supply side; more concentrated industries lead to more pronounced effects. The U.S. auto market can be regarded as moderately concentrated, with a four-firm concentration ratio of approximately 65% and a Herfindahl index of approximately $1,400$ (based on numbers of units sold in 2007–08). Finally, asset specificity matters. Firms are more likely to support homogeneous regulations when they have high "multinational" asset specificity, meaning that the value of a their investments is dependent upon cross-border transactions, and would lose value from different standards (Murphy, 2004b). In the context of automobile standards, firms have established production and distribution systems built around selling the same products in all 50 states, and would incur costs if forced to meet different standards in different states. This is especially true in the case of fuel consumption and GHG standards, since fuel consumption is integral to the design and capabilities of the vehicle. Reducing fuel consumption cannot be accomplished with any sort of "end-of-pipe" solution, the way catalytic converters and other aftertreatment technologies can reduce criteria pollutant emissions.

Examining the political economy of regulatory competition is notoriously hard to do quantitatively. Murphy (2004a) notes that since data are hard to obtain and it is hard to measure transaction costs and asset specificity, this sort of analysis is usually done on a comparative basis, rather than quantitatively. Often, a qualitative analysis is only practical option. The work reported in this chapter was undertaken as a first step toward the quantitative analysis of the political economy of nested state and federal GHG standards, and a better understanding the emissions leakage issue.

6.1.5 Objectives of This Work

Existing analyses of nested state and federal fuel consumption / GHG standards have given light attention to the political economy of nested standards. They have acknowledged the political economy issues in a qualitative and ex-post fashion, as their main focus has been on economic efficiency and the potential for emissions leakage while taking the regulatory environment as exogenously given. However, the political economy literature on regulatory competition suggests that we may want to treat the regulatory environment as endogenously determined. Since the automotive industry is fairly concentrated, engaged in selling regulated products across state lines, and California and other adopting states represent a sizeable fraction of the U.S. vehicle market, we would expect that the auto industry would generally prefer to deal with a single set of standards nationwide, rather than one set of standards at the federal level and a set of nested standards in the adopting states. Intuitively, however, there must be limits to this logic. If standards in the adopting states become *too* strict, then the lost revenues from building to those standards nationwide may come to exceed the costs of building separate fleets of vehicles for adopting and non-adopting states.

Presumably, lawmakers in California are interested in achieving the greatest reductions in GHG emissions at the most modest economic cost to their own state. To the extent that their policies drive GHG reductions beyond their borders, they may regard this as a positive outcome. Emissions leakage, on the other hand, will lead to economic costs in adopting states but lesser or even no emissions reductions nationwide. Since GHGs are global pollutants, heterogeneous standards and leakage of GHG emissions means economic costs are borne in adopting states, for no overall environmental benefit: a most undesirable outcome for the adopting states.

In this chapter, I develop a model of firms' responses to nested state and federal fuel consumption or GHG emissions standards. The goal of this work is to illustrate how different outcomes related to emissions leakage can depend on the relative stringency of state and federal standards, the share of the federal market that is covered by the adopting states, and the costs of developing and marketing different variants of products in adopting and non-adopting states. Using a simple numerical example, I will show three possible outcomes that may result:

1. *De facto* homogeneity: In this case, the competitive equilibrium is for firms to build all of their vehicles, in both adopting and non-adopting states, to the adopting states' standard. No leakage would occur in this case.
2. *De jure* homogeneity: In this case, the competitive equilibrium would be for firms to build separate variants of their vehicles for adopting and non-adopting states, creating emissions leakage. However, because of the costs of developing and marketing two separate variants, firms' profits are actually lower than if the stricter standards had been adopted nationwide. Thus, each firm can be made better off by selling only one variant nationwide, as long as their competitors are forced to do so as well. If firms can effectively coordinate to lobby for national adoption of the tighter standards, then emissions leakage would be mitigated.
3. Heterogeneity: The state standards are so strict that it is worth the additional cost to firms to develop and market two separate variants for the adopting and non-adopting states. In this case, emissions leakage can be expected to occur.

This suggests that additional, more detailed simulations may be justified to investigate where precisely the boundaries between these outcomes fall.

6.2 Model Description

The modeling approach developed in this section builds on a number of prior studies addressing manufacturer responses to fuel economy standards. Like Goulder et al. (2012), I model automobile producers as oligopolists in Bertrand competition, who determine the prices and attributes of their products to maximize profits, subject to regulatory constraints and expected competitor actions and consumer decisions. Following authors from the optimal design literature, I explicitly account for tradeoffs between fuel consumption and other consumer-valued attributes, the costs of advanced technology adoption, and employ a random utility model of consumer choices for vehicles (Michalek et al., 2004; Shiau et al., 2009; Whitefoot, 2011). Into this general approach, I add a term accounting for the capital costs associated with developing a product and bringing it into production, and assume that this cost depends on the number of variants of the model that are available. The general modeling framework is outlined in this section, and the functional form and parametric assumptions for a simple implementation of the model are developed in the following section.

I assume that each firm k attempts to maximize the present value of its own profits across all of the vehicles it sells:

$$\Pi_k = \sum_{i,j,r,t} p_{i,j,k,r,t} q_{i,j,k,r,t} d_{k,t} - \sum_{i,j,r,t} C^v_{i,j,k,r,t} q_{i,j,k,r,t} d_{k,t} - \sum_{j,t} C^f_{j,k,t} d_{k,t} \tag{6.1}$$

In Equation 6.1, $p_{i,j,k,r,t}$ denotes the price of variant i of model j sold by firm k in region r in year t, q denotes the quantity of that variant sold, $d_{k,t}$ is a discount factor to convert cash flows in year t into their present value for firm k, C^v is the variable cost of producing the variant, and $C^f_{j,k,t}$ is the fixed cost of bringing model j to market (including design, engineering, testing, certification, and tooling). Typically, the fixed costs of product development would be incurred some time before the launch of a redesigned product, and a particular generation would be kept in production for a number of years (Hill et al., 2007).

The quantity q of a particular variant sold depends on the total market size, Q, and the variant's market share S. Market share in turn depends on the price p and design attributes X of the vehicle (including fuel consumption, FC), and the prices $\boldsymbol{p_{r,t}}$ and design attributes $\boldsymbol{X_{r,t}}$ of all other vehicles available in the same region at the same time:

$$q_{i,j,k,r,t} = Q_{r,t} S(p_{i,j,k,r,t}, X_{i,j,k,r,t}, \boldsymbol{p_{r,t}}, \boldsymbol{X_{r,t}}) \tag{6.2}$$

Market share may be represented using a discrete choice model, such as the nested logit model, or the mixed logit model if accounting for heterogeneity in consumers' tastes for different vehicle attributes. The variable cost of manufacturing a vehicle is assumed to depend on the vehicle's attributes:

$$C^v_{i,j,k,r,t} = C^v(X_{i,j,k,r,t}) \tag{6.3}$$

The fixed cost of developing and marketing a vehicle model is assumed to depend on the characteristics of the basic model and the number of variants of model j developed by firm k in period t:

$$C^f_{j,k,t} = C^f(X_{j,k,t}, n_{j,k,t}) \tag{6.4}$$

Each manufacturer sets the prices and attributes of its own product so as to maximize its own profitability, assuming its competitors will do the same. Where a federal fuel consumption standard is in force, it acts as a constraint on the firm's optimization problem:

$$\sum_{i,j,r,t} FC_{i,j,k,r,t} q_{i,j,k,r,t} \leq \sum_{i,j,r,t} T(X_{i,j,k,r,t}) q_{i,j,k,r,t} \tag{6.5}$$

In the above inequality, FC is the fuel consumption of a variant, and is one of the attributes X. $T(X)$ is a function that defines the target fuel economy of a variant depending on its attributes. T may be size-based, as with the current U.S. CAFE standards (Figure 6-2), or it may be simply a constant. Where a nested standard or standards are in force, they impose additional constraints, as firms must satisfy the standards in each region r, where T_r is the region-specific fuel consumption target function:

$$\sum_{i,j,t} FC_{i,j,k,r,t} q_{i,j,k,r,t} \leq \sum_{i,j,t} T_r(X_{i,j,k,r,t}) q_{i,j,k,r,t} \tag{6.6}$$

As specified, the constraints in Inequalities 6.5 and 6.6 imply that firms are permitted to bank and borrow credits between periods, so that the constraint binds over all periods combined. Dropping t from the summation limits would yield a set of constraints that bind on each year individually. Under the federal CAFE law, firms have limited opportunities for banking (up to five years) and borrowing (up to three years) of credits, as well as for credit trading between firms. Bunch et al. (2011) discusses the structure of these constraints in more detail.

Figure 6-2: Corporate Average Fuel Economy standards depend on size mix of vehicles sold. Source: EPA and NHTSA (2010)

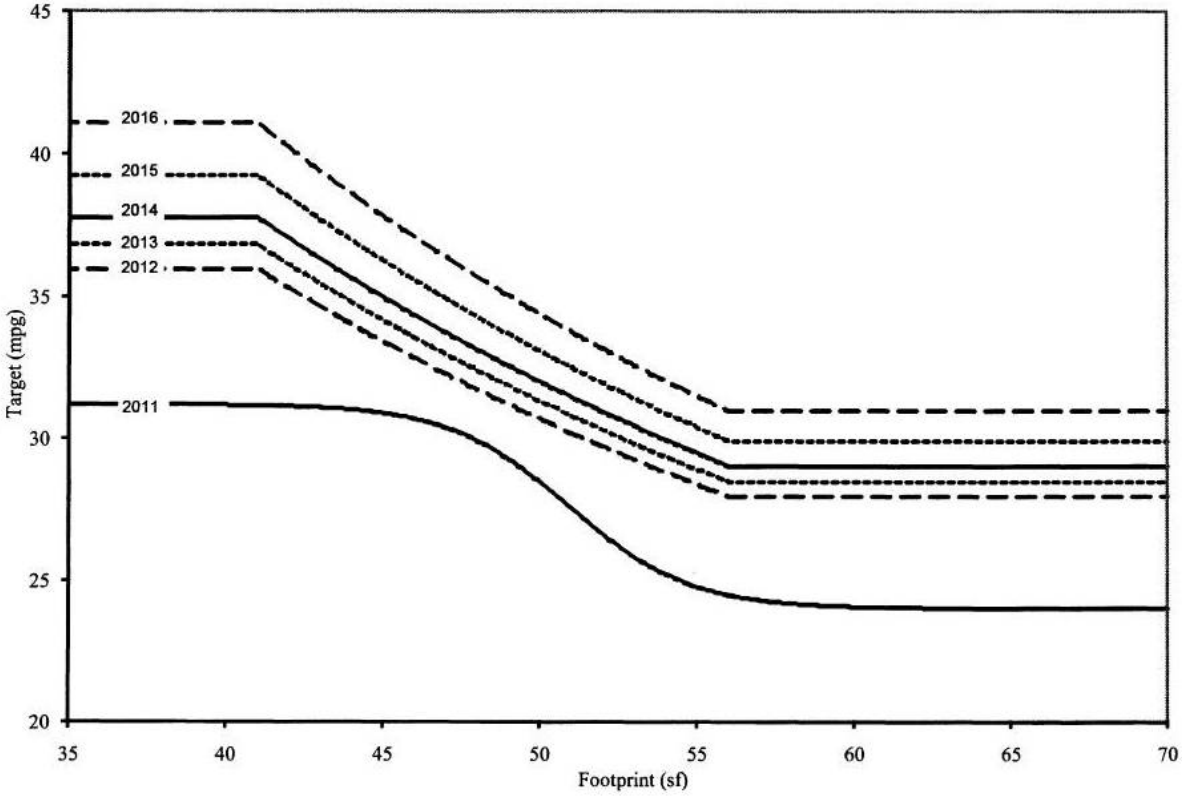

The resulting system of equations can be solved for a Nash equilibrium either by iteratively optimizing each firm's strategy until convergence (Michalek et al., 2004), or by simultaneously solving the first-order conditions(Shiau et al., 2009). The former is simple to implement while the latter is much less computationally intensive (Shiau and Michalek, 2007).

6.3 A Simple Implementation

In this section, I present a simple implementation of the above model for ten firms competing in a single vehicle class, across two regions (one nested inside the other). The objective is not to come up with definitive, quantitative answers, so much as to demonstrate the feasibility of implementing the simulation methodology and to illustrate several types of outcomes that may emerge from a market with nested fuel consumption regulations. In the following sections, I outline the assumptions employed in this modeling work.

6.3.1 Basic Structure

The single vehicle class that forms the basis of this modeling work can be thought of, roughly, as representing non-luxury midsize cars in the U.S. Between 2004 and 2008, there were on average ten firms competing in this segment, with total sales averaging 2.2 million units per year, making it one of the largest vehicle segments in the U.S. Firms are assumed to be symmetric, with identical cost structure, capabilities, and products indistinguishable apart from the variables considered in this analysis. Firms compete in a single period, bringing their products to market at the same time and keeping them in production for an equal number of years.

6.3.2 Product Life and Discount Rate

All revenues and production costs were discounted back to the point at which funds are committed for product development. The discount rate used was 7.6%, based on the average cost of capital for the automotive industry (Elter and Castedello, 2012). Each model is assumed to be in production for six years (Blonigen et al., 2013), beginning two years after the funds are committed to product development (Hill et al., 2007). To simplify calculations, a present worth factor of 4.35 is applied to annual revenues and production costs to convert the six years' worth of cash flows into a present value at a time two years before product launch.

6.3.3 Total Market Size

The total potential market is assumed to be all drivers in the U.S. Fitting a linear model to data from the U.S. Federal Highway Administration (FHWA, 2010) indicates that there are currently about 220 million licensed drivers in the U.S., and that this figure is increasing by 2.5 million per year. Most of these drivers, of course, do not buy a new midsized car in any given year. Instead, they may buy another class of new vehicle, a used car, keep their existing car, or be licensed but not own a vehicle at all. Those who do not purchase a new midsized car in any given year are considered to have chosen the “outside good,” which is a composite good treated as having constant utility.

6.3.4 Vehicle Attributes and Technology

The key design attributes considered in this simulation are per-mile fuel consumption and 0–60 mph (0–97 km/h) acceleration performance. As shown in section 4.4.3, since 1975 changes in these two variables have been the top two “sinks” for new efficiency technologies in U.S. cars.

Manufacturers are assumed to be able to trade off between acceleration and fuel consumption for a given level of technology. That is to say, they can decrease fuel consumption by increasing acceleration time, or vice versa, even without changing their use of efficiency technologies. Based on the results presented in section 4.4.1, it was assumed that every 1% increase in 0–97 km/h acceleration time would result in a 0.44% decrease in fuel consumption.

The change in efficiency technology on a vehicle was expressed in terms of equivalent fuel consumption. This is calculated as the sum of the actual change in fuel consumption and the fuel consumption equivalent of the change in acceleration. Both fuel consumption and acceleration are measured relative to baseline values of 2.8 gallons / 100 miles (35.7 mpg, unadjusted test cycle) and 8 seconds, respectively.[3] The changes in fuel consumption and acceleration performance from these baseline values

---

[3] These values were chosen based on typical values for the 2012 Toyota Camry, Honda Accord, and Ford Fusion, assuming on-road fuel economy numbers are 25% below test cycle numbers.

were used to calculate a technology metric as follows:

$$Tech_{i,j,k,r,t} = \left(\left(\frac{FC_{Base}}{FC_{i,j,k,r,t}} \right) \left(\frac{Z97_{Base}}{Z97_{i,j,k,r,t}} \right)^{0.44} - 1 \right) \cdot 100 \tag{6.7}$$

In Equation 6.7, $Tech$ is a measure of the new technology added to a car, relative to the baseline vehicle. It is expressed in terms of the percentage point increase in fuel economy that could have been achieved using the same technology if acceleration performance had been held constant at its baseline value.

6.3.5 Variable Costs of Production

Following the basic approach of Michalek et al. (2004) and Shiau et al. (2009), the variable costs of production are assumed to include a cost to produce the basic vehicle, an engine cost that depends on acceleration capabilities, and a technology cost for the advanced technologies:

$$C^{v}_{i,j,k,r,t} = C^{B}_{i,j,k,r,t} + C^{A}_{i,j,k,r,t} + C^{T}_{i,j,k,r,t} \tag{6.8}$$

Michalek et al. (2004) estimated a cost curve for the cost of an engine as depending on its power, in kW:

$$C^{E} = 670.51e^{0.0063P} \tag{6.9}$$

Michalek et al. then subtracted this estimated engine cost from the total production costs reported by Delucchi and Lipman (2001). Applying this approach and adjusting for inflation, the manufacturing cost exclusive of the engine in the work of Delucchi and Lipman was estimated to be 40% of the manufacturer's suggested retail price (MSRP). Applying this same 40% factor to the \$26,000 MSRP of a typical 2012 midsize car yields a figure of \$10,500 for the basic vehicle manufacturing cost, C^{B}.

The acceleration cost C^{A} was assumed to depend on the following cost curve, which was derived by combining the engine cost curve of Michalek et al., the acceleration model reported in Chapter 2, and specification (weight, displacement, transmission, etc.) data characteristic of current midsize cars.

$$C^{A}_{i,j,k,r,t} = 589.8e^{\frac{10.6795}{Z97_{i,j,k,r,t}}} \tag{6.10}$$

Finally, technology costs were based on a recent report by the National Petroleum Council (NPC, 2012). Costs for midsize cars were interpolated between the NPC's estimates for small and large cars, and between their high-range and low-range estimates. The resulting cost estimates were fitted to a second order polynomial of the following form:

$$C^{T}_{i,j,k,r,t} = \alpha_{1,t} Tech_{i,j,k,r,t} + \alpha_{2,t} Tech^{2}_{i,j,k,r,t} \tag{6.11}$$

As implied by Equation 6.11, reflecting the projections in the NPC report, the costs of technology are expected to fall over time. The assumed parameters for the cost curve in Equation 6.11 are summarized in Table 6.1.

Table 6.1: Assumed cost curve parameters for midsize cars, based on NPC (2012).

Year	α_1	α_2
2012	51.0880	0.25996
2015	49.9460	0.30043
2020	42.9898	0.25859
2025	38.9044	0.23401
2030	0.22257	37.0022

6.3.6 Fixed Costs of Product Redesign

The fixed costs of bringing a redesigned model to market is assumed to start at $0.9 billion for a model with a single variant, based on estimates reported by Blonigen et al. (2013) for midsize cars. Estimating the cost of an additional variant is more uncertain, because it is thought to depend on a wide range of factors, such as how different the two variants are, whether an entirely new engine needs to be developed or an existing engine can be "dropped in," and so on. If an existing engine could be used, then there would be some engineering, integration, and possibly testing & certification costs. If a new engine had to be developed, the costs could run into the hundreds of millions of dollars — a new engine plant in the 1990s was estimated to require between \$300 million and \$800 million in capital investment (Whitney et al., 1997). In the face of such uncertainty, I explore a range of costs for creating a second variant for the nested market, from \$100-500 million.

6.3.7 Models and Variants Offered in Each Region

In this implementation, it was assumed that each firm offers a single midsize car model. The firm decides whether to offer one variant across the entire federal market, or one variant in the adopting states and another in the non-adopting states. If the same variant is offered in both sets of states, the price can be set differently in each set of states. If two variants are offered, then acceleration performance, fuel consumption, technology, and price can all vary between the two regions. However, in all cases, only one variant is offered in each region. In reality, firms producing multiple variants would probably offer all variants in both regions.

6.3.8 Consumer Choice Model

Consumer choices for vehicles were assumed to be described by a nested logit model, similar to that reported by Bunch et al. (2011). An obvious difference here is that there are only two levels of choice in the current model: the choice between a new midsize car and the outside good (i.e. buy/no-buy) and the choice of a specific model of midsize car. A second key difference is that the current specification also includes a term for the utility of (log) acceleration. Finally, since firms are assumed to be symmetric and their products identical except for acceleration and fuel consumption levels, there are no alternative-specific constants for the different models of car. The observed portion of utility was assumed to be:

$$V_{i,j,k,r,t} = \beta_1 p_{i,j,k,r,t} + \beta_2 (FC_{i,j,k,r,t}) + \beta_3 \ln(Z97_{i,j,k,r,t}) \tag{6.12}$$

The coefficient on price in Equation 6.12 was assumed to be $\beta_1 = -0.00011$, which corresponds to an own-price elasticity of demand of -2.4 (within the midsize car nest).

The coefficient on fuel consumption was assumed to be $\beta_2 = -0.19$ when fuel consumption is expressed in gallons per 100 miles. This corresponds to a willingness-to-pay (WTP) of \$1,750 for a 1 gallon / 100 miles reduction in fuel consumption, which is within the range of values reviewed by Whitefoot (2011). This value is equal to the fuel costs saved over 50,000 miles of driving at a gasoline price of \$3.50 per gallon, which is a common rule of thumb of establishing consumers' WTP (Greene et al., 2005, 2009).

The coefficient on $\ln(Z97)$ was assumed to be $\beta_3 = -0.62$, which corresponds to a WTP of \$700 for a one-second reduction in 0–97 km/h acceleration time at a mean acceleration time of 8 seconds. This is consistent with the WTP values reviewed in MacKenzie (2009).

Based on the standard results for the nested logit model, the utility of buying a car relative to not buying a car was assumed to be given by:

$$V_{r,t}^{buy} = A_{buy} + \mu \ln \left(\sum_{i,j,k} e^{V_{i,j,k,r,t}} \right) \tag{6.13}$$

The average utility of buying a midsize car, relative to the outside good, was assumed to be given by $A_{buy} = -5.6$, which yields estimated midsize sales of 1.95 million in 2012. The scale parameter, μ was assumed to be 0.4, which yields a price elasticity of market demand of -1.00.

6.3.9 Fuel Consumption Constraint Structure

Fuel consumption constraints were assumed to be binding in this implementation of the model. It is well known that in the case of federal CAFE standards, some firms have historically treated the standards as binding while others have been willing to pay the modest civil penalties (\$55 per mpg per vehicle) for noncompliance (Jacobsen, 2012; Kleit, 2004). However, with EPA how enforcing a parallel set of GHG standards under the authority of the Clean Air Act, penalties for noncompliance could theoretically reach \$37,500 per vehicle, and full compliance with the new standards is anticipated (EPA and NHTSA, 2010).

In the current implementation, uniform standards were assumed in both the adopting states and federally. In the context of modeling a single vehicle class, size-based standards are essentially the same as uniform standards, since there are no other size classes to which volume can be shifted. However, if modeling the responses of manufacturers producing vehicles in multiple size classes, accounting for size-based standards would be essential.

A range of federal and adopting state standards were investigated, ranging from 2.7 gallons / 100 miles (37 mpg, not binding) down to 1.9 gallons / 100 miles (52.6 mpg).

6.4 Simulation Results

Figures 6-3 – 6-7 summarize the estimated results from the imposition of a single federal fuel consumption standard. The reductions are measured relative to an unconstrained case. A federal fuel consumption standard of 2.7 gallons / 100 miles in 2025 would not bind under the assumptions outlined above, so this case represents the modeled free-market outcome. Moving to the left in each figure, the standard becomes tighter, and prices increase, sales and profits decrease, acceleration times deteriorate, and more advanced technologies are adopted.

Each one of Figures 6-8 – 6-11 summarize the firm profitability and leakage implications of the (non-cooperative) Nash equilibria as a function of the fuel consumption standard in the adopting states and the cost to firms of developing a separate variant of their model for adopting and non-adopting states. The four figures correspond to different combinations of loose and tight federal standards and larger and smaller market shares for the adopting states.

Let us first consider Figure 6-8. The upper panel summarizes the non-cooperative leakage outcomes when there is a loose federal standard (2.5 gal/100 miles, 40 mpg) in force, and the adopting states represent 20% of the total market. Start by considering an adopting state standard of 2.5 gal / 100 miles. This is the same as the federal standard, so the national average fuel consumption ends up at this level.

Figure 6-3: Price increases are predicted as federal fuel consumption standard is tightened.

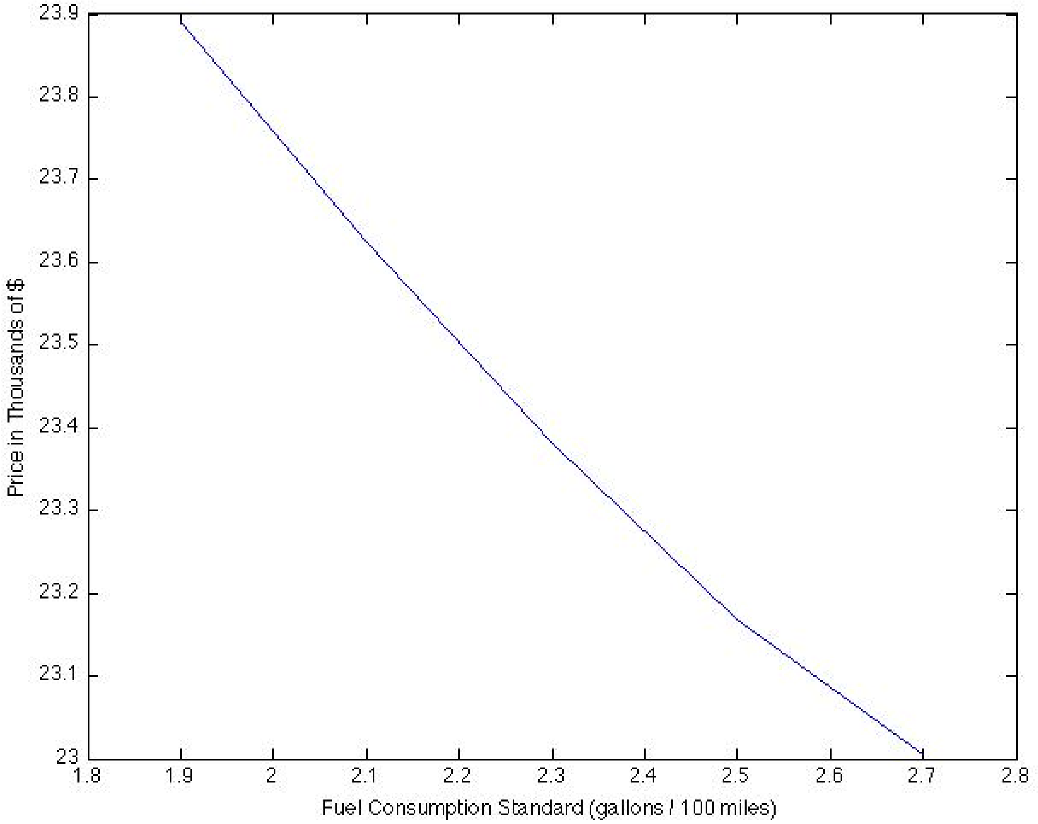

If the adopting states tighten their standard to 2.3 gal / 100 miles, however, then the outcome depends on the cost to firms of creating separate product variants. If a special variant for the adopting states can be developed for \$100 million, then the Nash equilibrium has each firm creating two variants. In this case, the national average fuel consumption stays at 2.5 gal / 100 miles, indicating that the expected fuel savings in the adopting states have leaked out to the non-adopting states. However, look now at panel 2. In the Nash equilibrium, the profits per firm are actually lower than they would be if all of the firms simply built to the tighter standard nationwide. In the absence of cooperation, however, each firm has an incentive to produce two variants. This is individually rational for each firm but collectively

Figure 6-4: Sales are predicted to fall as federal fuel consumption standard is tightened.

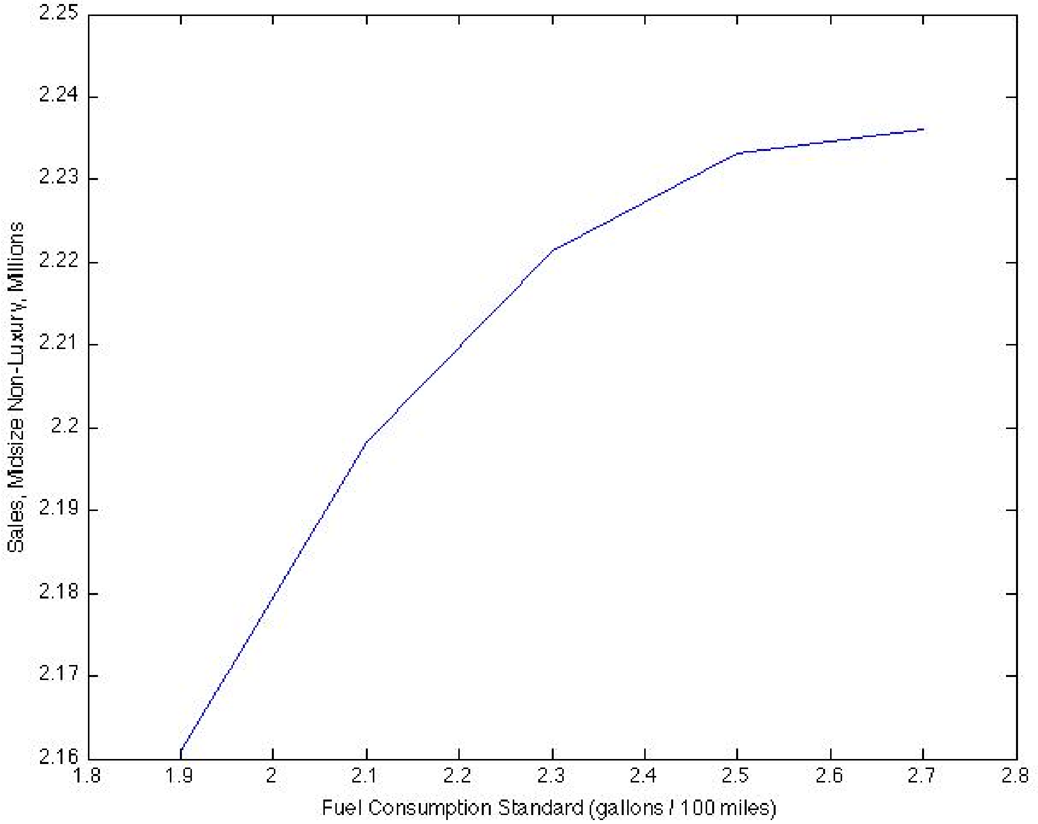

suboptimal for the industry as a whole. Thus, the automobile producers in this case would have an incentive to lobby the federal government to adopt the tighter 2.3 gallon / 100 mile standard nationwide. This standardization would remove the incentive for individual firms to create separate variants for adopting and non-adopting states, and increases profitability for all firms. This would also mitigate the predicted leakage effect, through what I earlier termed *de jure* homogeneity.

The results are different, however, if the costs of creating a separate variant are higher. If the cost of creating an additional variant is $200 million or more, then even in the absence of federal regulations, it would make sense for the each firm to build a single variant that meets the tighter, adopting states standard, and simply sell

Figure 6-5: Reduction in profits per firm under a single federal fuel consumption standard.

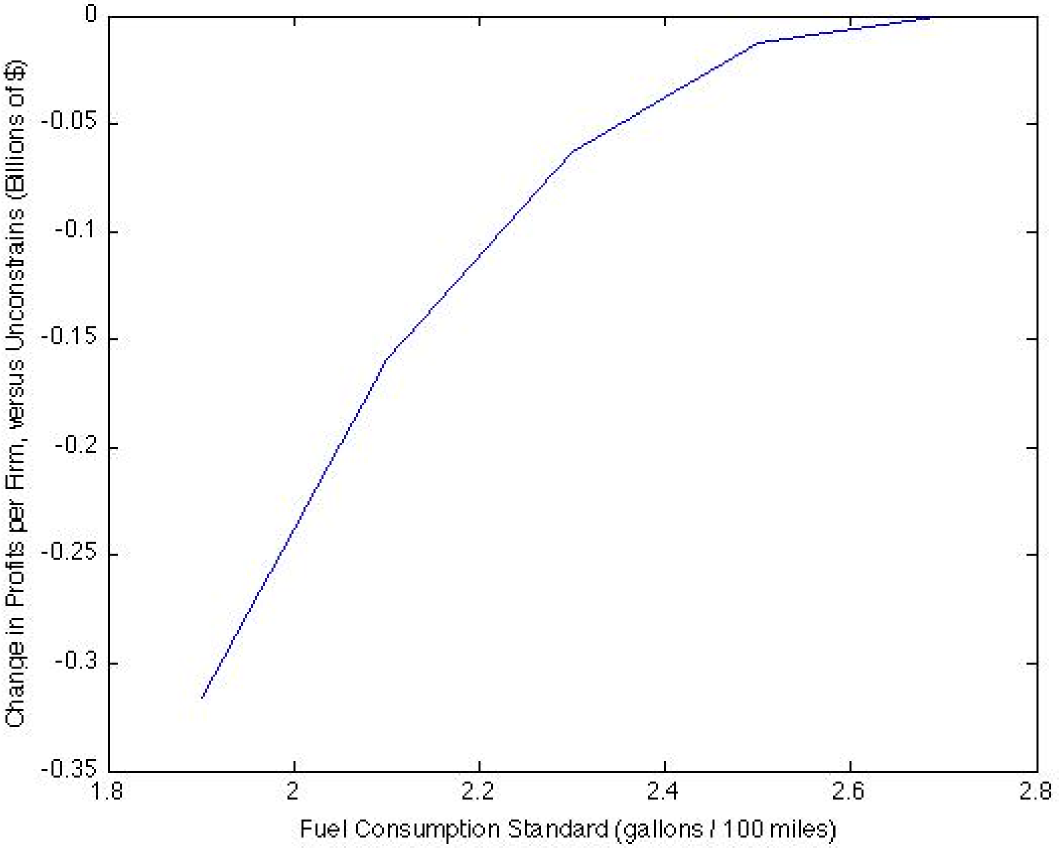

this variant in all states. In this case, the adopting state standard would be followed nationwide, no leakage would occur, and we have an example of *de facto* homogeneity.

Let us next consider the possible outcomes if the adopting states adopt a tighter standard: 2.1 gallons / 100 miles. Now, if the cost of a second variant is $100 million, then the most profitable course of action for the firms — *both individually and collectively* — is to split their models into two variants each. Here we have leakage, and no incentive for the firms to support harmonization of state and federal standards. This is the heterogeneous outcome. However, if the cost of creating two variants is higher — $200–300 million — then it is individually rational but collectively suboptimal for each firm to create two variants of its model. In this case the firms

Figure 6-6: Acceleration times deteriorate under a federal fuel consumption standard as firms trade off acceleration performance for lower fuel consumption.

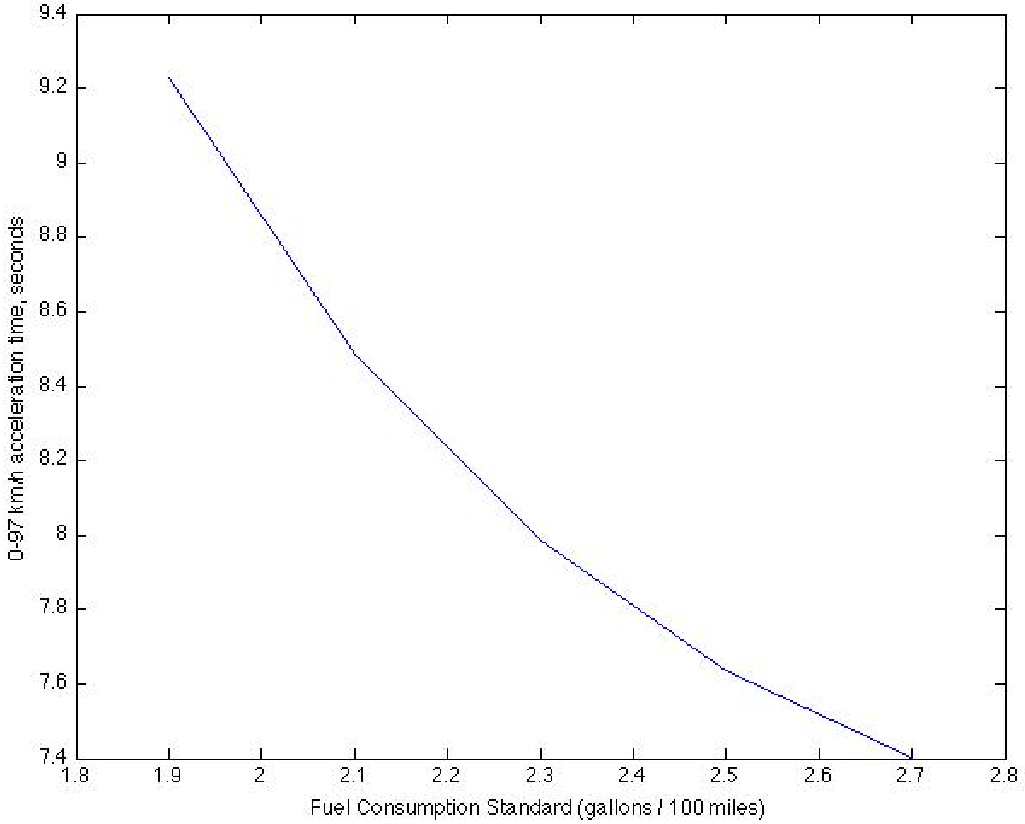

again have an incentive to support federal adoption of the tighter standard, as it is more profitable to build a single variant to the tighter standard, as long as everyone else does so as well. Once again, the incentives are aligned for the creation of *de jure* homogeneity.

Figure 6-9 shows similar results, when the adopting states are a larger share (40%) of the national market. Several differences are notable, compared with the prior case in which adopting states represent only 20% of the market. First, for an adopting state standard of 2.3 gal / 100 miles, there is no incentive for firms to create a second variant, even if they can do so for only $100 million. The smaller market share of the non-adopting states means that the benefits to building a separate variant are smaller,

Figure 6-7: A tighter federal fuel consumption standard is predicted to increase the application of advanced efficiency technologies.

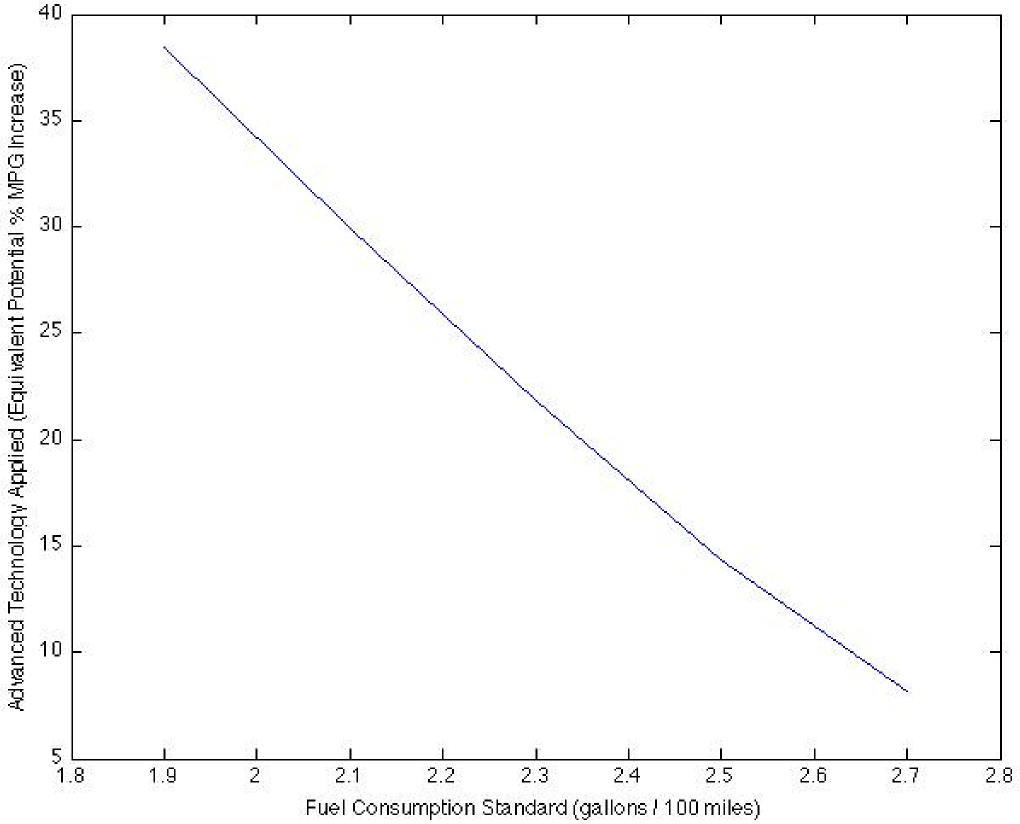

so we end up with *de facto* homogeneity regardless of the cost of developing a second variant (within the range of costs investigated). Second, even when leakage occurs, as when the adopting states adopted a standard of 2.1 gal / 100 miles, the national average fuel consumption does not return all the way to the federal standard of 2.5 gal / 100 miles, indicating that leakage is only partial. Finally, the results indicate that if the adopting states tightened their standard to 1.9 gal / 100 miles, the firms would have an incentive to lobby for national adoption of the tighter standard (*de jure* homogeneity) if the cost of developing a second variant were $200 million or more. However, if it were less than $200 million, they would be better off in the non-cooperative equilibrium, selling two variants each.

6.5 Implications

The above modeling exercise, while a vastly simplified representation of the automobile market, illustrates the potential for several outcomes that may result from nested state and federal fuel consumption regulations:

1. When the states adopting a tighter fuel consumption or GHG standard constitute a large share of the overall market, when the cost of developing multiple variants is high, or when the nested standard is not too stringent, we can end up with *de facto* homogeneity: manufacturers find it most profitable to simply build to the tighter standard and sell those products everywhere, even if they are not required to do so.
2. When the adopting states represent a smaller share of the total market or the cost of developing a second variant is relatively low, the Nash equilibrium may have firms producing multiple variants. However, this may actually be less profitable than what firms could achieve if they (and their competitors) were all forced to build to the tighter standard nationwide. This situation creates an incentive for firms to support national adoption of the tighter standards: *de jure* homogeneity.
3. When the cost of developing a second variant is not too high, and the adopting states' standard is relatively strict, it may be most profitable for firms to accept the additional costs of building multiple variants. The savings from building a single variant are not worth the loss in revenue in non-adopting states, even if competitors are forced to build to the tighter standard as well. In this case, we end up with heterogeneous standards in the adopting and non-adopting states, and emissions leakage occurs.

This work suggests that in modeling the potential for emissions leakage under nested regulations, it may be important to consider the costs to the manufacturers of developing separate variants to sell in adopting and non-adopting states. These costs create an incentive for firms to support national adoption of the tighter standards, or in some cases to build to the tighter standards everywhere even if it is not required nationally. More detailed modeling of this market should also account for the fact that there are also market incentives to offering multiple variants of each model: namely, if consumer tastes for fuel consumption, acceleration performance, and purchase price are heterogeneous, then firms should be able to capture additional market share by creating multiple variants tailored to those heterogeneities.

From the perspective of a policymaker in an adopting state, homogeneity — either *de facto* or *de jure* — is presumably the desired outcome when dealing with global pollutants. It is hard to envision any reason that a policymaker would want to impose large costs on sellers and consumers in his own state, if the intended emission reductions are simply going to leak out to non-adopting states. The first implication of the results in this chapter for such a policymaker is that she should work in coalitions with other states: the more the merrier. Second, she may be able to drive national emission reductions even if her standards are not formally adopted nationwide. Finally, she must be careful not to overplay her hand: if she advances standards that are too onerous, support for their broader adoption will evaporate, and the hard-won emissions reductions in her own state will simply leak out to the non-adopting states.

Figure 6-8: Firm profitability and leakage outcomes with a loose federal standard and a smaller number of adopting states.

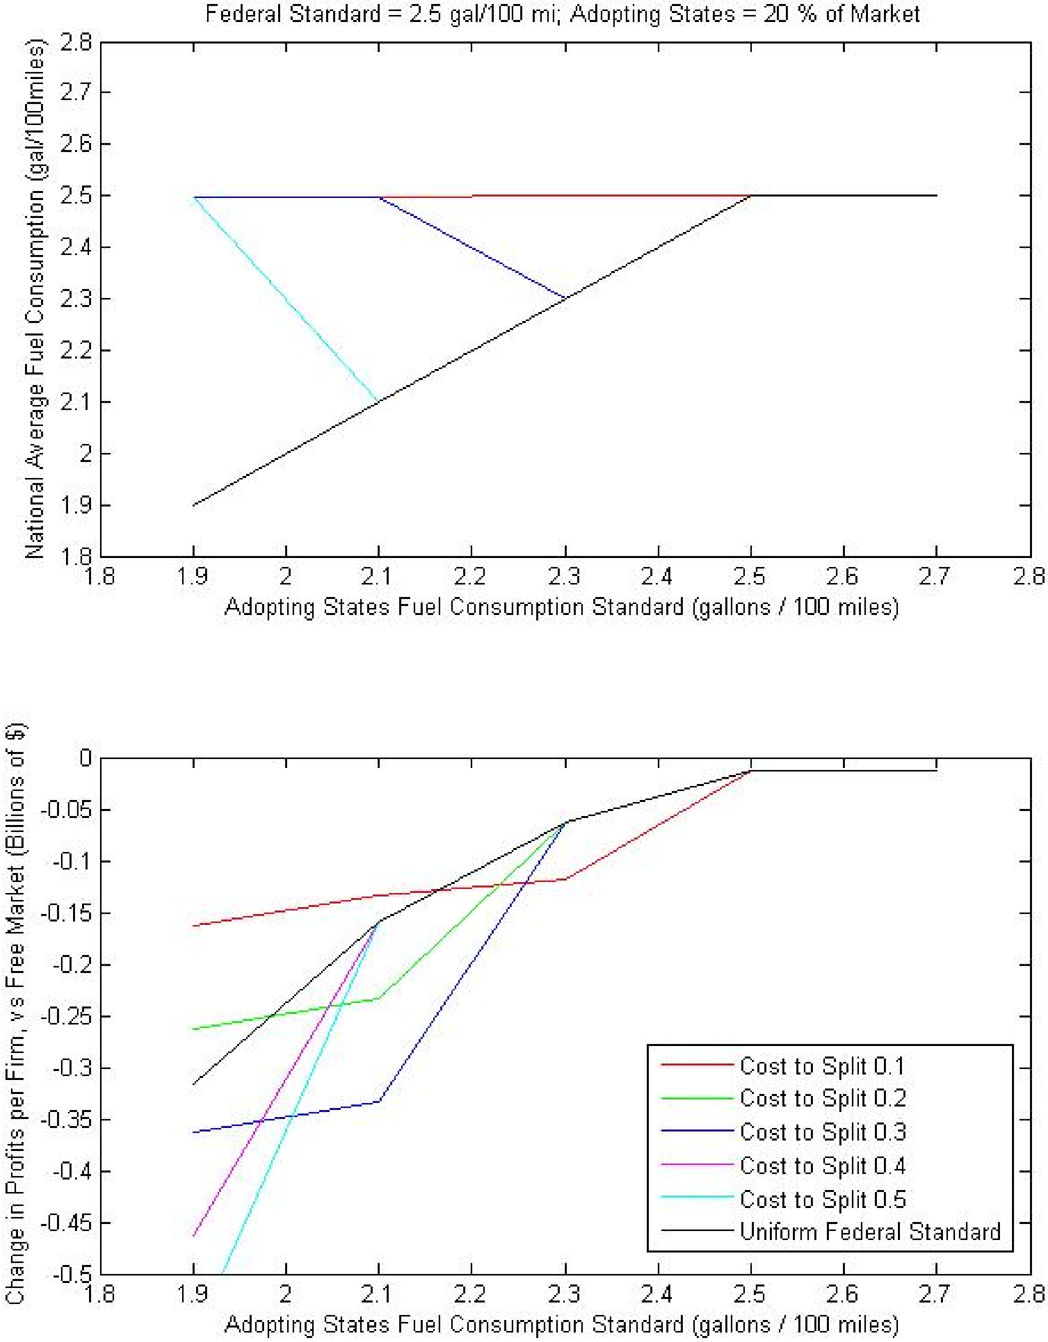

Figure 6-9: Firm profitability and leakage outcomes with a loose federal standard and a larger number of adopting states.

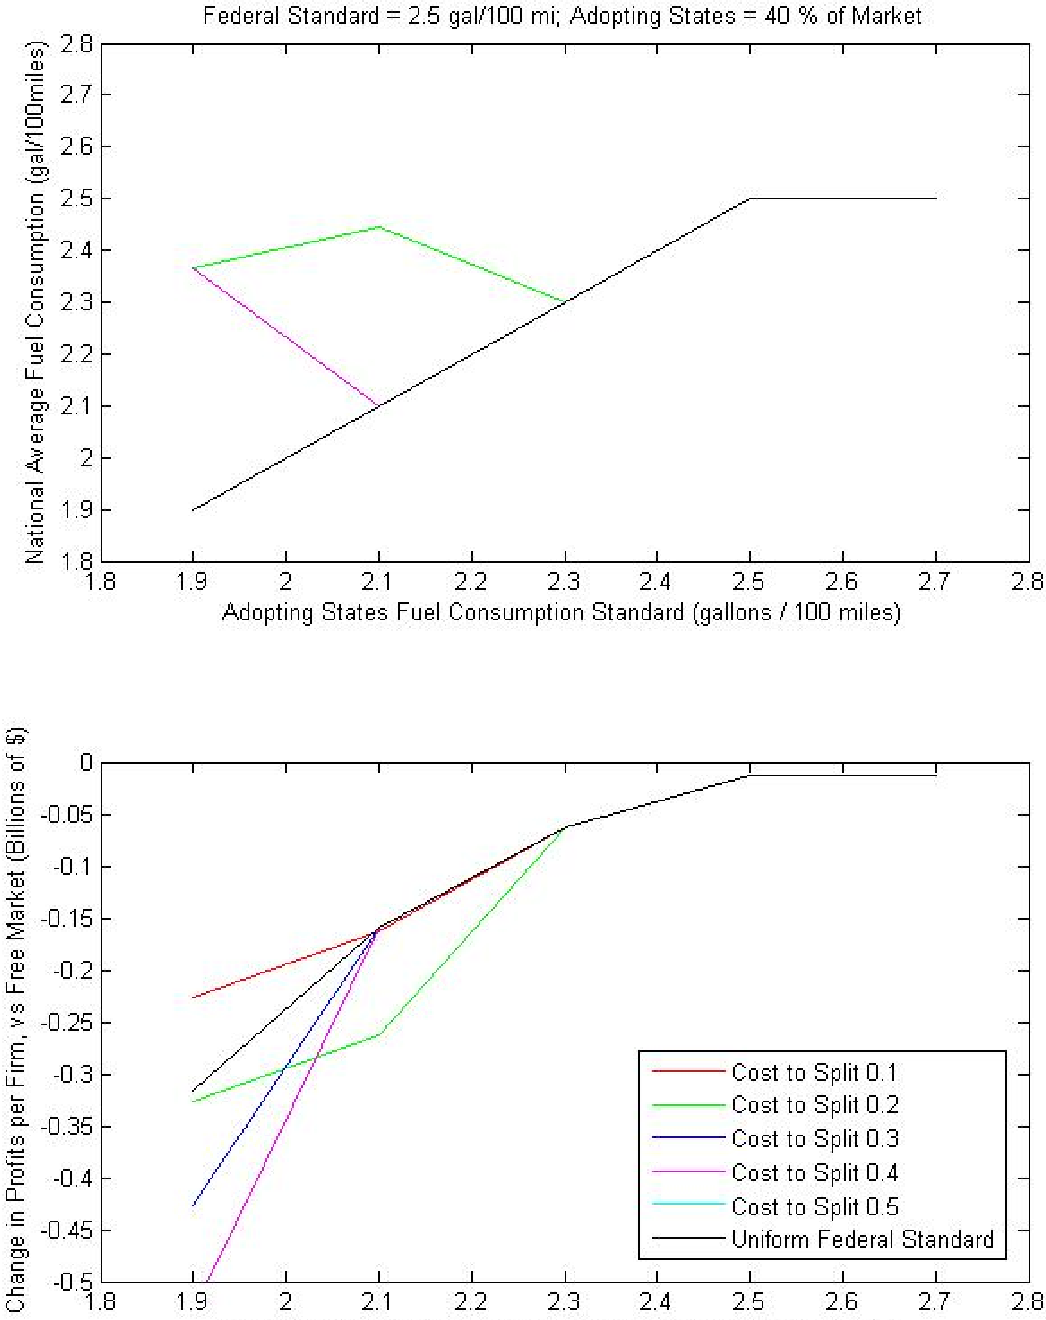

Figure 6-10: Firm profitability and leakage outcomes with a tight federal standard and a smaller number of adopting states.

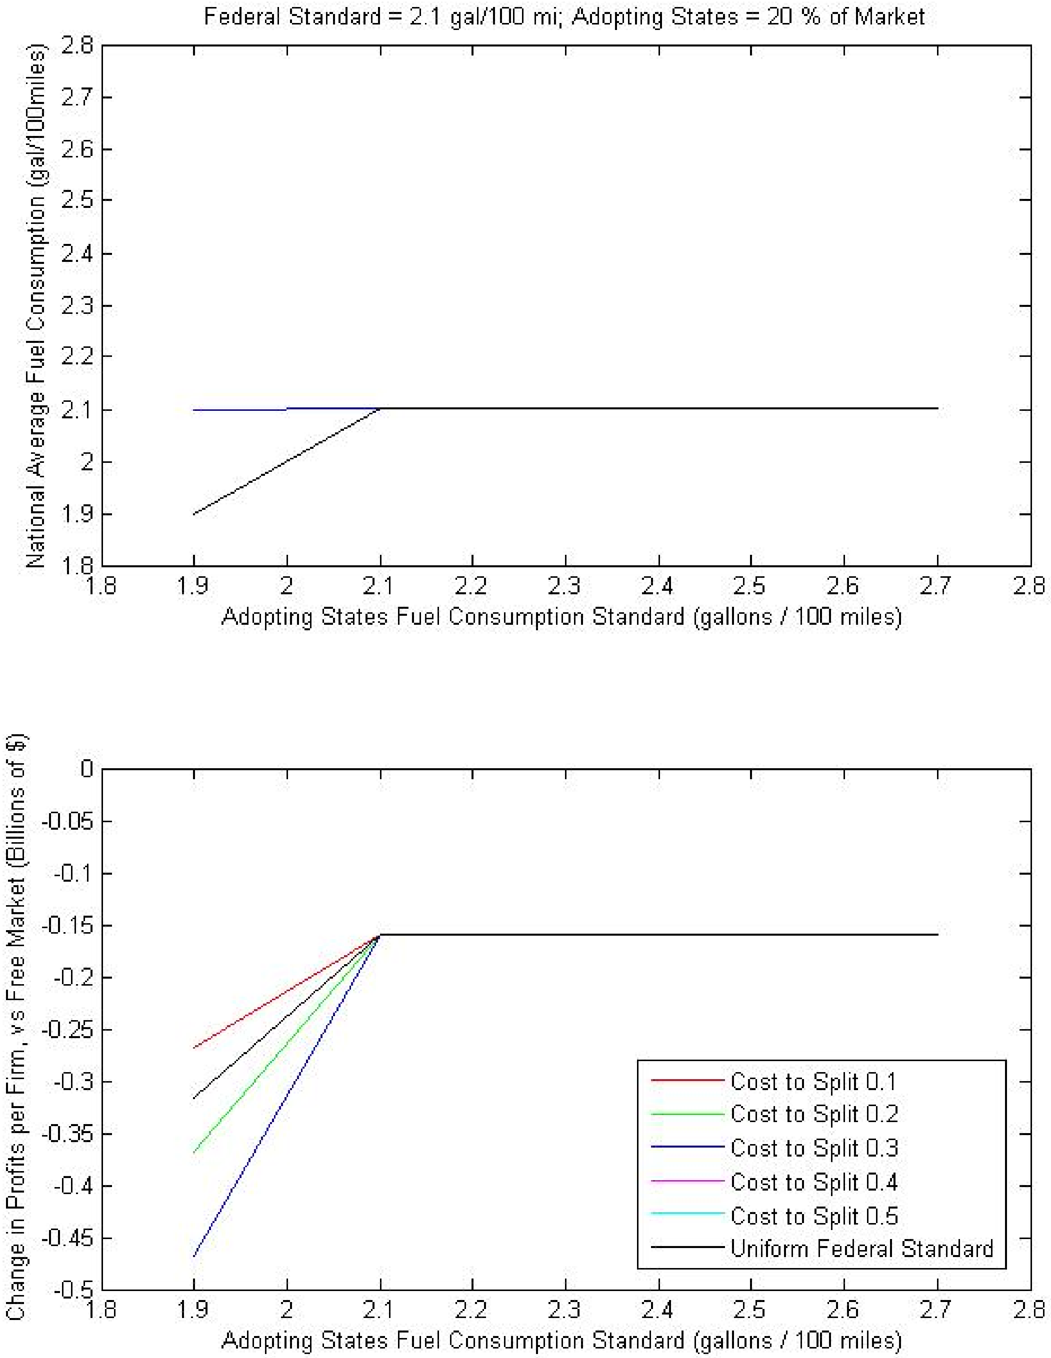

Figure 6-11: Firm profitability and leakage outcomes with a tight federal standard and a larger number of adopting states.

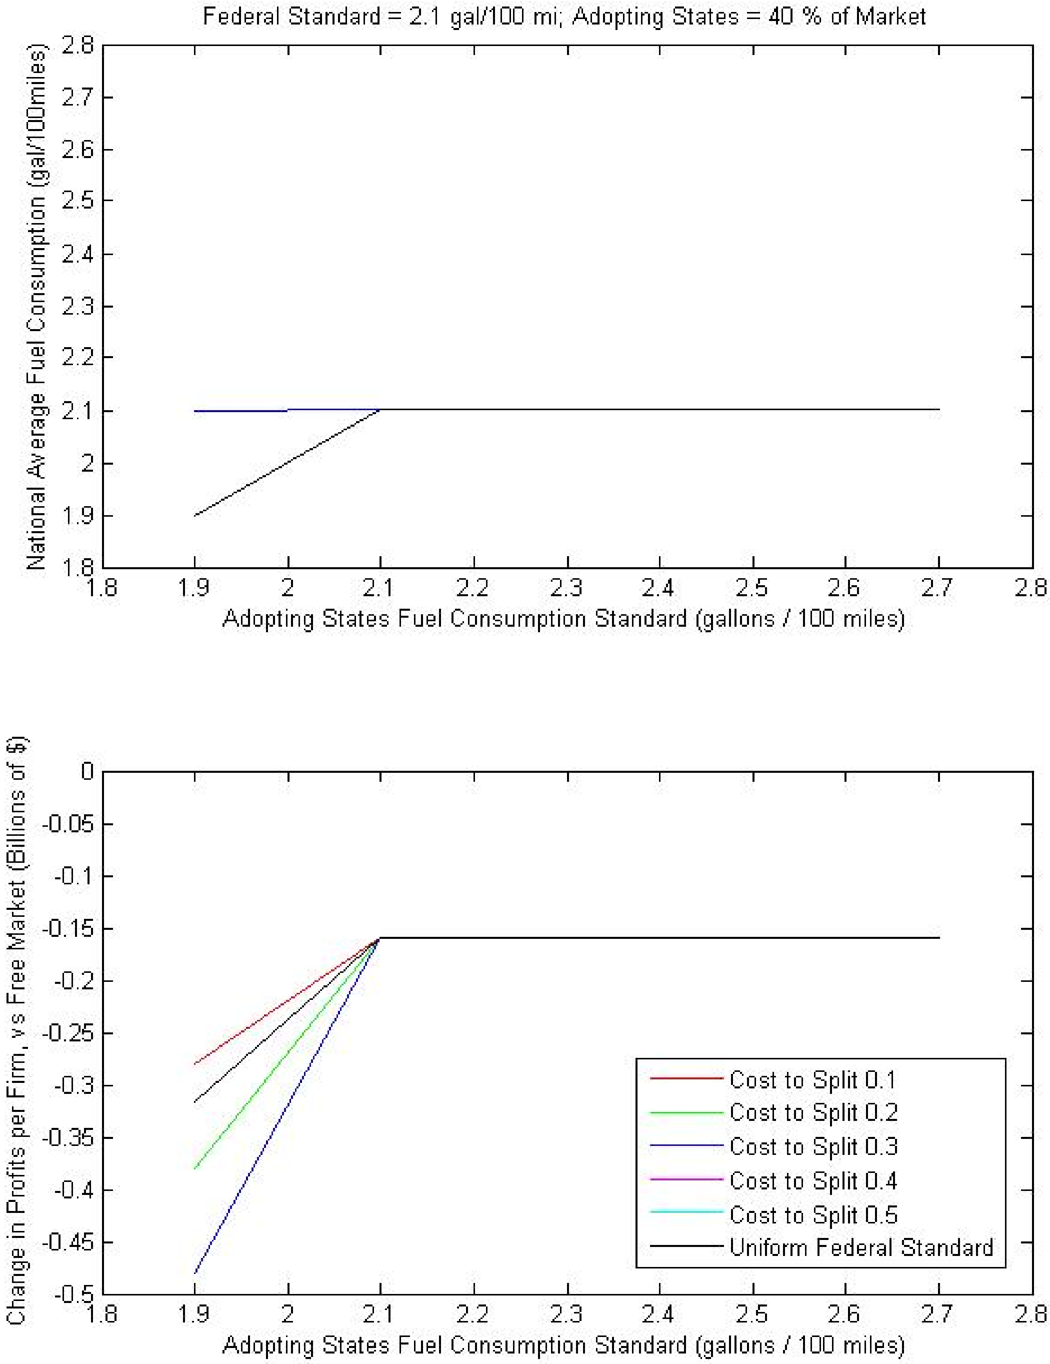

Chapter 7

Conclusions

In this dissertation, I have sought to improve our understanding of several important questions pertaining to automotive fuel efficiency technology and related policies. First, I began by exploring recent trends in automotive acceleration performance and weight, two variables intricately related with fuel consumption. Second, I developed an estimate of how much per-mile fuel consumption could have been decreased since 1975, if not for changes in the acceleration capabilities, features, and functionality of new cars. Third, I tested whether fuel efficiency technologies improved more quickly when firms are more tightly constrained by fuel economy standards, or when gasoline prices are higher. Finally, I presented a theoretical model of the incentives of firms when faced with "nested" state and federal fuel consumption regulations.

Understanding vehicle acceleration performance is critical to evaluating trends in efficiency technology, because there is a significant tradeoff between acceleration performance and fuel consumption. Apart from actual reductions in per-mile fuel consumption, the largest "sink" for new efficiency technologies since 1975 has been in offsetting the fuel economy penalties of acceleration improvements. I have presented a set of linear regression models that will enable analysts and researchers to conveniently estimate vehicle acceleration times using commonly-reported attributes such as power, weight, model year, and in some cases, selected body and powertrain characteristics. These models represent a significant improvement upon the correlations used by many analysts to date, which are based on much more limited data from the 1970s. I have shown that even after accounting for changes in weight, power, engine and transmission characteristics, today's new cars accelerate 20–30% faster than did new cars in 1975. One of the models was applied to historic vehicle sales and specification data, and the results revealed that there have been consistent reductions in 0–97 km/h acceleration times for new vehicles in the U.S. since 1982. Acceleration times have been reduced across all segments of the market, including the fastest and the slowest vehicles. A car with acceleration performance among the fastest 5% of the market in the mid-1980s would have been in the slowest 5% of the market by 2009. These changes have occurred more slowly in recent years (in both relative and absolute terms), and the acceleration times appear to be asymptoting. Extrapolating the trends suggests that reductions of a further 1 second (about 10%) in 0–97 km/h acceleration times may be plausible in the coming years.

Vehicle weight is another important characteristic, as greater weight causes increased fuel consumption. Weight *per se* is therefore is a source of disutility in a vehicle, but many other desirable vehicle attributes, such as greater size and feature content, tend to increase weight. I have explored the contributions of various weight-saving technologies, and the contributions of changes in size and feature content, to the overall changes in weight of the average new car in the U.S. since 1975. Growth in the prevalence of front-wheel drive and unibody construction, and increased use of high-strength steel, aluminum, and plastics, have all contributed substantially to reducing the weight of new cars since 1975. A shift from 6- and 8-cylinder engines to 4- and 6-cylinder engines, facilitated by markedly improved specific power in newer engines, has also contributed to weight reduction. All told, I estimate that weight-saving technologies have taken about 790 kg (about 40%) out of the average new car since 1975, while increases in size and feature content have added approximately 250 kg over the same period.

Viewing technological improvements in automobiles from a top-down, system functionality perspective — focusing on the reduction in fuel needed to deliver a certain level of size, comfort, features, and acceleration performance to the consumer — reveals vast improvements in the efficiency of new U.S. cars since 1975. New cars today burn only about half as much fuel as did a new car with the same weight and power in 1975. But this is only half of the story. Even with the same weight and power, today's cars deliver 20–30% faster acceleration times than did cars in 1975. Moreover, a car with the same safety, comfort and convenience features weighs about 40% less today than it would have in 1975. Integrating key results from the first several chapters of the dissertation, I conclude that if acceleration performance, size, and feature content had remained at their 1975 levels, the average new car today would be burning 70% less fuel per mile than in 1975. The actual reduction in per-mile fuel consumption between 1975 and 2009 was 50%.

Improvements in automotive efficiency technology have not been uniform over time. The potential reduction in per-mile fuel consumption of new U.S. cars averaged 5% annually between 1975 and 1990, and 80% of this potential was realized as actual reductions in fuel consumption. Between 1990 and 2009, however, potential fuel consumption reductions were just 2% per year, and actual fuel consumption reductions were just 34% of this potential. So, the pre-1990 and post-1990 periods are distinguished not only by the overall rate of technology improvement, but by the degree to which those improvements were dedicated to actually reducing fuel consumption. In both of these periods, modest amounts of new efficiency technologies — enough to decrease fuel consumption by roughly 1% per year — were needed in order to offset the fuel consumption penalties associated with decreasing acceleration times, increasing car size, and increasing feature content.

New federal corporate average fuel economy (CAFE) standards require sustained reductions in per-mile fuel consumption of approximately 4% each year through 2025. Meeting the standards will require either much faster technology improvements than have recently occurred, or some sacrifice in the functionality (acceleration performance, comfort, safety, or convenience features) of today's cars. Meeting these standards without sacrificing functionality (in fact, while improving other aspects of functionality) would be possible if technology could be improved at 5% per year, as occurred between 1975 and 1990.

The technological feasibility and economic practicability of meeting the 2025

CAFE standards is uncertain due to technical, policy, and market factors. First, it is unclear whether the automotive industry can return to the rates of technology change observed before 1990. Many of the new technologies adopted in the 1970s and 1980s were "low-hanging fruit" which has now been "picked." Most new cars today are already front-wheel drive, and an overwhelming majority already use unibody construction. As shown in Chapter 3, these architectural changes were important sources of weight reduction in the past, but they have saturated the market and so their potential to deliver further reductions is largely exhausted. The same holds true for technologies like lockup automatic transmissions, fuel injection, and computerized engine controls. Similarly, each additional transmission speed or valve per cylinder can be expected to offer diminishing marginal returns. Finally, while the acceleration performance for a vehicle with a given power and weight is faster than in the past, most of these improvements occurred in the 1980s, with more modest changes since 1990.

While the potential of many big-ticket technologies has already been saturated, opportunities do remain. Technologies like spaceframe construction, and advanced composite materials such as carbon-fiber reinforced polymer, offer the potential for large-scale weight reductions in the future. While these technologies present manufacturing challenges and are currently seen as suitable only for limited-volume applications, the same might have been said of front-wheel drive or unibody construction in the past. Similarly, hybrid electric vehicles are currently seen as having high potential for fuel consumption reduction, but limited potential for rapid, widespread adoption. However, we may yet be surprised and in 30 years' time look back on such technologies — in hindsight — as today's low-hanging fruit.

From a policy standpoint, a critical question is whether the stricter CAFE standards themselves might increase the rate of technology adoption. In Chapter 5 I explored this question by testing the hypotheses that fuel efficiency technologies improve more rapidly when gasoline prices are higher, and when firms are more tightly constrained by CAFE standards. Using data from the CAFE program's initiation in 1978 through 2008, I found little to no evidence of a significant effect of binding

CAFE standards on the rate of technology change in a firm's fleet of cars. However, I did find evidence that higher gasoline prices led to faster rates of innovation. These results suggest that meeting the 2025 CAFE standards without sacrificing current functionality may be more feasible if fuel prices remain at $3.50-$4.00 per gallon or above, but could be more difficult if fuel prices fall back below $2.00 per gallon, where they were between 1986 and 2003 (in 2012 dollars).

From a market standpoint, consumer preferences for different vehicle attributes are critical to the economic practicability of the CAFE standards. If consumers continue to demand attributes that detract from fuel consumption improvements, such as faster acceleration and weight-adding features, then meeting fuel economy standards will be less economically practical. Consumers will not necessarily continue to demand these changes, however. As shown in Chapter 2, U.S. consumers' thirst for acceleration performance appears to have been slowing in recent years, and may be on track to level off over the next 1–2 decades. Such an outcome could lead to "virtual performance" (DeCicco, 2010) becoming much more important as a key measure of utility. Virtual performance refers to a philosophy of shifting design efforts to vehicle characteristics that do not trade off against fuel consumption, such as richer connectivity and media capabilities. If such features — which are basically software — become the main profit center for new automobiles, then improvements in acceleration performance, and even weight-increasing hardware features, may become less critical to the profitability of vehicles, making way for greater fuel consumption reductions.

Finally, I have presented an alternative perspective on the establishment of fuel economy and GHG standards for automobiles. Much of the work in this dissertation focused on rates of efficiency technology adoption, the feasibility of different levels of standards, and the effect of the standards on the rate of technology change. However, recent rulemakings establishing stricter federal CAFE and GHG emissions standards have been determined at least as much by the political actions and interactions of a bloc of states led by California, the automotive industry, and multiple federal agencies. To understand how such outcomes emerge, it is helpful to consider the costs to automobile manufacturers of building to stricter standards versus the benefits of having only a single standard to consider. I have presented a simple quantitative model that illustrates three potential sets of incentives facing manufacturers in the context of “nested” state and federal standards, where one region of a broader market imposes a stricter standard than is required across the market as a whole. This model offers an explanation for why the automobile industry so stridently opposed California’s authority to set automotive GHG standards, only to embrace those same standards nationwide once California’s authority was upheld.

A simple implementation of this model using parameter values characteristic of the midsize car market illustrates the three possible outcomes. First, if the state standards are not too strict, manufacturers may simply choose to build to the tighter (adopting state) standard everywhere. Second, if the adopting state standards are more strict, manufacturers may find it more profitable to build separate variants of their product for the adopting and non-adopting states than to build to the stricter standard everywhere. However, depending on the cost of developing separate variants, it may be more profitable still for them to lobby for nationwide adoption of the stricter standard. In this case, all firms may be better off building a single version of their product nationwide, as long as their competitors are compelled to do the same. Finally, if the nested standards are very strict, then it may be most profitable for all firms to build separate variants of their products for the two regions. In this case, the incentive for nationwide adoption of the more stringent standard is undercut, and adopting states alone end up bearing the costs of the stricter standards. Worst of all, if the manufacturers were already constrained in the non-adopting states, then the sale of lower-emitting vehicles in the adopting states may enable the sale of higher-emitting vehicles in the non-adopting states, partially or completely offsetting the emission reductions expected in the adopting states. These results suggest that the stringency of California’s GHG standards was in the second category — though a more complete model, accounting in particular for heterogeneous consumer tastes and product characteristics, would be warranted to explore this question in more depth.

Bibliography

ACC (2011). Chemistry & light vehicles. Technical report, American Chemistry Council.

An, F. and DeCicco, J. (2007). Trends in technical efficiency trade-offs for the us light vehicle fleet. *SAE Technical Paper*, (2007–01–1325).

Bandivadekar, A., Bodek, K., Cheah, L., Evans, C., Groode, T., Heywood, J., Kasseris, E., Kromer, M., and Weiss, M. (2008). On the road in 2035: Reducing transportation's petroleum consumption and ghg emissions. Technical Report LFEE 2008-05 RP, Massachusetts Institute of Technology.

Berry, I. M. (2010). The effects of driving style and vehicle performance on the real-world fuel consumption of us light-duty vehicles. Master's thesis, Massachusetts Institute of Technology.

Blonigen, B. B., Knittel, C. R., and Soderbery, A. (2013). Keeping it fresh: Strategic product redesigns and welfare.

Bunch, D. S., Greene, D. L., Lipman, T., Martin, E., and Shaheen, S. (2011). Potential design, implementation, and benefits of a feebate program for new passenger vehicles in california: interim statement of research findings. Technical Report UCD 08-312, University of California, Davis: Institute of Transportation Studies. Prepared for State of California Air Resources Board and the California Environmental Protection Agency.

Carle, D. and Blount, G. (1999). The suitability of aluminium as an alternative material for car bodies. *Materials & design*, 20(5):267–272.

Cheah, L. (2010). *Cars on a diet: the material and energy impacts of passenger vehicle weight reduction in the US*. PhD thesis, Massachusetts Institute of Technology.

Cheah, L. W., Bandivadekar, A. P., Bodek, K. M., Kasseris, E. P., and Heywood, J. B. (2009). The trade-off between automobile acceleration performance, weight, and fuel consumption. *SAE International Journal of Fuels and Lubricants*, 1(1):771–777.

Chon, D. M. and Heywood, J. (2000). Performance scaling of spark-ignition engines: Correlation and historical analysis of production engine data. *SAE Technical Paper*, (2000–01–0565).

Crabb, J. and Johnson, D. (2010). Fueling innovation: The impact of oil prices and cafe standards on energy-efficient automotive technology. *The Energy Journal*, 31(1):199–216.

Das, S., Curlee, T., and Schexnayder, S. (1997). Materials used in new generation vehicles: supplies, shifts, and supporting infrastructure. Technical Report ORNL/TM-13491, Oak Ridge National Lab., TN (United States).

DeCicco, J. and Ross, M. (1993). An updated assessment of the near-term potential for improving automotive fuel economy. Technical report, American Council for an Energy-Efficient Economy, Washington, DC.

DeCicco, J. M. (2010). A fuel efficiency horizon for u.s. automobiles. Technical report, The Energy Foundation.

Delucchi, M. A. and Lipman, T. E. (2001). An analysis of the retail and lifecycle cost of battery-powered electric vehicles. *Transportation Research Part D: Transport and Environment*, 6(6):371–404.

DOT (2004). Cost and weight added by the federal motor vehicle safety standards for model years 1968-2001 in passenger cars and light trucks. Technical Report DOT HS 809 834, U.S. Department of Transportation.

Dupnick, E. and Graham, J. (1996). The demand for advanced materials in the automotive industry: Projections for the next decade. In *AIP conference proceedings*, volume 361, page 163.

EIA (2011). Annual energy outlook 2011, with projections to 2035. Technical Report DOE/EIA-0383(2011), U.S. Energy Information Administration.

EIA (2013). March 2013 monthly energy review. Technical Report DOE/EIA-0035(2013/03), U.S. Energy Information Administration.

Elter, V.-C. and Castedello, M. (2012). Cost of capital study 2011/2012: Developments in volatile markets. Technical report, KPMG.

EPA (2009a). California state motor vehicle pollution control standards; notice of decision granting a waiver of clean air act preemption for californias 2009 and subsequent model year greenhouse gas emission standards for new motor vehicles. *Federal Register*, 129(199):32744–32784.

EPA (2009b). Regulatory announcement: Epa will propose historic greenhouse gas emissions standards for light-duty vehicles. Technical Report EPA-420-F-09-028, Office of Transportation and Air Quality.

EPA (2010). Light-duty automotive technology, carbon dioxide emissions, and fuel economy trends: 1975 through 2010. Technical Report EPA-420-R-10-023, U.S. Environmental Protection Agency.

EPA (2012). Light-duty automotive technology, carbon dioxide emissions, and fuel economy trends: 1975 through 2011. Technical Report EPA-420-R-12-001a, U.S. Environmental Protection Agency.

EPA (2013a). Inventory of u.s. greenhouse gas emissions and sinks: 1990–2011. Technical Report EPA 430-R-13-001, U.S. Environmental Protection Agency, 1200 Pennsylvania Ave., N.W. Washington, DC 20460.

EPA (2013b). Light-duty automotive technology, carbon dioxide emissions, and fuel economy trends: 1975 through 2012. Technical Report EPA-420-R-13-001, U.S. Environmental Protection Agency.

EPA (2013c). Regulations & standards: Presidential announcements and letters of support.

EPA and NHTSA (2010). Light-duty vehicle greenhouse gas emission standards and corporate average fuel economy standards; final rule. *Federal Register*, 75(88):25324–25728.

EPA and NHTSA (2012). 2017 and later model year light-duty vehicle greenhouse gas emissions and corporate average fuel economy standards: Final rule. *Federal Register*, 77(199):62623–63200.

FHWA (2010). Highway statistics series.

Fischer, C., Harrington, W., and Parry, I. (2007). Should automobile fuel economy standards be tightened? *The Energy Journal*, 28:1–29.

Goulder, L. H., Jacobsen, M. R., and Van Benthem, A. A. (2012). Unintended consequences from nested state and federal regulations: The case of the pavley greenhouse-gas-per-mile limits. *Journal of Environmental Economics and Management*, 63(2):187–207.

Goulder, L. H. and Stavins, R. N. (2011). Challenges from state-federal interactions in us climate change policy. *The American Economic Review*, 101(3):253–257.

Greene, D. (1990). CAFE or price?: An analysis of the effects of federal fuel economy regulations and gasoline price on new car mpg, 1978-89. *The Energy Journal*, 11(3):37–57.

Greene, D. (2001). TAFV alternative fuels and vehicles choice model documentation. Technical Report ORNL/TM-2001/134, Oak Ridge National Laboratory (US).

Greene, D. L. and Fan, Y. (1994). Transportation energy efficiency trends, 1972–1992. Technical report, Oak Ridge National Lab., TN (United States).

Greene, D. L., German, J., and Delucchi, M. A. (2009). *Reducing climate impacts in the transportation sector*, chapter 11, pages 181–229. Springer.

Greene, D. L., Patterson, P. D., Singh, M., and Li, J. (2005). Feebates, rebates and gas-guzzler taxes: a study of incentives for increased fuel economy. *Energy Policy*, 33(6):757–775.

Gutherie, A. (1978). Fairmont/zephyr-engineered for lightweight and improved fuel economy. Technical Report 780134, Society of Automotive Engineers, 400 Commonwealth Dr, Warrendale, PA, 15096, USA,.

Hill, K., Szakaly, S., and Edwards, M. (2007). How automakers plan their products: A primer for policymakers on automotive industry business planning. Technical report, Economics and Business Group, Center for Automotive Research, Ann Arbor, MI.

Jacobsen, M. (2012). Evaluating us fuel economy standards in a model with producer and household heterogeneity. `http://econ.ucsd.edu/~m3jacobs/Jacobsen_CAFE.pdf`.

Jaffe, A. and Stavins, R. (1994). Energy efficiency investments and public policy. *The Energy Journal*, 15(2):43–43.

Kasseris, E. P. (2006). Comparative analysis of automotive powertrain choices for the near to mid-term future. Master's thesis, Massachusetts Institute of Technology.

Kleit, A. (2004). Impacts of long-range increases in the fuel economy (cafe) standard. *Economic Inquiry*, 42(2):279–294.

Knittel, C. R. (2011). Automobiles on steroids: Product attribute trade-offs and technological progress in the automobile sector. *The American Economic Review*, 101(7):3368–3399.

Kromer, M. A. and Heywood, J. B. (2007). Electric powertrains: Opportunities and challenges in the us light-duty vehicle fleet. Technical Report LFEE 2007-03 RP, Massachusetts Institute of Technology: Laboratory for Energy and the Environment.

Luo, A. (2002). Magnesium: current and potential automotive applications. *JOM Journal of the Minerals, Metals and Materials Society*, 54(2):42–48.

Lutsey, N. and Sperling, D. (2005). Energy efficiency, fuel economy, and policy implications. *Transportation Research Record: Journal of the Transportation Research Board*, 1941:8–17.

Lutz, R. (2011). *Car Guys vs. Bean Counters: The Battle for the Soul of American Business*. Portfolio/Penguin, New York, NY.

MacKenzie, D. W. (2009). Trends and drivers of the performance – fuel economy tradeoff in new automobiles. Master's thesis, Massachusetts Institute of Technology.

Malliaris, A., Hsia, H., and Gould, H. (1976). Concise description of auto fuel economy and performance in recent model years. Technical Report 760045, Society of Automotive Engineers, 400 Commonwealth Dr, Warrendale, PA, 15096, USA,.

Mayer, S., Seeds, A., and of Automotive Engineers, S. (1994). BMW's aluminum light-weight prototype car projects: Comparison of aluminum and steel performance. Technical Report 940154, Society of Automotive Engineers, 400 Commonwealth Dr, Warrendale, PA, 15096, USA,.

Michalek, J. J., Papalambros, P. Y., and Skerlos, S. J. (2004). A study of fuel efficiency and emission policy impact on optimal vehicle design decisions. *Journal of Mechanical Design*, 126:1062.

Murphy, D. D. (2004a). *The Business Dynamics of Global Regulatory Competition*, chapter 2, pages 84–117. University of California Press.

Murphy, D. D. (2004b). *The structure of regulatory competition: Corporations and public policies in a global economy*. Oxford University Press.

Murphy, T. (2006). Aluminum claims no. 2 ranking. WardsAuto.com.

NESCAUM (2008). Northeast States for Coordinated Air Use Management: Submission to federal docket id no. NHTSA-2008-0089. `www.nescaum.org/documents/nescaum-cafe-comment-ltr-final-070108.pdf`.

Newell, R., Jaffe, A., and Stavins, R. (1999). The induced innovation hypothesis and energy-saving technological change. *The Quarterly Journal of Economics*, 114(3):941–975.

NHTSA (2006). Average fuel economy standards for light trucks model years 20082011: Final rule. Federal Register, 71, 66 / Thursday, April 6, 2006 / Rules and Regulations.

NPC (2012). Advancing technology for americas transportation future. Technical report, National Petroleum Council.

OTA (1991). Improving automobile fuel economy: New standards, new approaches. Technical Report OTA-E-504, U.S. Congress, Office of Technology Assssment.

Pilz, H., Schweighofer, J., and Kletzer, E. (2005). The contribution of plastic products to resource efficiency. Technical report, Plastics Europe - Association of Plastics Manufacturers, Brussels, Belgium.

Popp, D. (2002). Induced innovation and energy prices. *The American Economic Review*, 92(1):160–180.

Roth, R., Dieffenbach, J., Koehr, R., and Isaacs, J. (1998). Economic analysis of the ultra light steel auto body. In *SAE Conference Proceedings P*, pages 285–290. Society of Automotive Engineers.

Salonitis, K., Pandremenos, J., Paralikas, J., and Chryssolouris, G. (2009). Multifunctional materials used in automotive industry: A critical review. In *Engineering Against Fracture*, volume 1, page 59.

Santini, D. and Anderson, J. (1993). Determinants of multiple measures of acceleration. Technical Report 931805, Society of Automotive Engineers, 400 Commonwealth Dr, Warrendale, PA, 15096, USA,.

Schafer, J. L. and Kang, J. (2008). Average causal effects from nonrandomized studies: a practical guide and simulated example. *Psychological Methods*, 13(4):279.

Shiau, C.-S. and Michalek, J. (2007). A game-theoretic approach to finding market equilibria for automotive design under environmental regulation. In *Proceedings of the ASME International Design Engineering Technical Conferences*, volume 6, pages 187–195.

Shiau, C.-S. N., Michalek, J. J., and Hendrickson, C. T. (2009). A structural analysis of vehicle design responses to corporate average fuel economy policy. *Transportation Research Part A: Policy and Practice*, 43(9):814–828.

Stodolsky, F., Vyas, A., Cuenca, R., and Gaines, L. (1995). Life-cycle energy savings potential from aluminum-intensive vehicles. In *1995 T-bonel Life Cycle Conference & Exposition*, Vienna, Austria.

Vogel, D. (1995). *Trading Up: Consumer and Environmental Regulation in a Global Ecoonomy*. Harvard University Press.

Vogel, D. and Kagan, R. (2004). *National Regulations in a Global Economy*, chapter Introduction, pages 1–41. University of California Press.

Whitefoot, K., Fowlie, M., and Skerlos, S. (2011). Product design response to industrial policy: Evaluating fuel economy standards using an engineering model of endogenous product design. *Energy Institute at Haas Working Paper WP-214. Berkeley, CA: University of California Energy Institute.*

Whitefoot, K. S. (2011). *Quantifying the impact of environmental policy on engineering design decisions*. PhD thesis, University of Michigan.

Whitney, D., Peschard, G., and Artzner, D. (1997). Cost performance of automobile engine plants. IMVP Working Paper W-0180a, Massachusetts Institute of Technology.

Wohlecker, R., Wallentowitz, H., Henn, R., and Leyers, J. (2006). Communication module–mass reduction. Technical Report 56690, Forschunggesellschaft Kraftfahrwesen mbH Aachen, Body Department.

Young, Y. (1991). Regression analysis for acceleration performance of light duty vehicles. Technical Report DOT HS 807 763, U.S. Department of Transportation.

Zoepf, S. (2011). Automotive Features: Mass Impact and Deployment Characterization. Master's thesis, Massachusetts Institute of Technology, 77 Massachusetts Ave. Cambridge, MA 02139, USA.

Zoepf, S. and Heywood, J. (2012). Characterizations of deployment rates in automotive technology. *SAE International Journal of Passenger Cars-Electronic and Electrical Systems*, 5(2):541–552.