THREE ESSAYS IN PUBLIC ECONOMICS

by

James M. Sallee

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

Doctoral Committee:

Professor Joel B. Slemrod, Chair
Professor Rebecca M. Blank
Professor James R. Hines, Jr.
Professor Jeffrey A. Smith

to Reece Jones
and my many other teachers

ACKNOWLEDGEMENTS

I would like to thank my dissertation committee — Joel Slemrod, Rebecca Blank, James Hines and Jeffrey Smith — for excellent mentorship throughout the dissertation process. I would also like to thank the members of an informal graduate student working group to which I belong — Soren Anderson, Brian Cadena, Ben Keys, Brian Kovak and Alex Resch. This working group made a critical contribution to the development of these projects.

Two of the three essays are coauthored. The second essay, on marriage laws, is coauthored with Rebecca M. Blank and Kerwin Kofi Charles. The third essay, on higher education, is coauthored with Alexandra M. Resch and Paul N. Courant. Obviously, this research would not have been possible without their contributions.

I have many people to thank for their assistance with each essay. On the first essay, in addition to the aforementioned, I would like to thank John Bound, Charlie Brown, Dave Cole, Lucas Davis, Dhammika Dharmapala, Wei Fan, Kai-Uwe Kühn, Jim Levinsohn, John List, Walter McManus, Caroline Sallee, Dan Silverman, Chad Syverson, Sarah West, Marina Whitman, Rob Williams, Martin Zimmerman and seminar participants at a number of universities and conferences. Special thanks are due to Wei Fan for significant assistance with the data.

On the second essay, my coauthors and I would like to thank Kristin Essary, Ari Kushner, Emily Beam and David Vorobeychik for excellent assistance in putting together the data. We would also like to thank Esther Duflo and two anonymous referees, along with seminar participants at the University of Chicago, the University of Toronto, the University of Pennsylvania, and Johns Hopkins University for comments.

On the third essay, my coauthors and I would like to Soren Anderson, John Bound, Brian Cadena, Jim Hines, Ben Keys, Mike McPherson, Michael Rothschild, Jeff Smith, Sarah Turner and seminar participants at the University of Virginia and the University of Michigan for their comments and critiques. We also thank Caroline Hoxby and two anonymous referees for excellent comments.

I would also like to acknowledge the financial support of the National Science Foundation and the National Institute for Child Health and Development. I am also very grateful for the work space and collegial atmosphere provided by the National Poverty Center.

Finally, I would like to thank my wife, Caroline Sallee, for steering both of us through the dissertation maelstrom with remarkable poise. Without her strong hand, I would have surely smashed against the rocks.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Introduction

This dissertation consists of three distinct essays. They are not closely related to each other, and, as such represent the breadth of my interests and the freedom granted to me by my committee, for which I am grateful. All three essays may broadly be characterized as public economics, and all three demonstrate my interest in applying economic reasoning to the design and evaluation of public policy. While each essay studies a different topic and uses distinct methods, all three commonly address how public policy can have unexpected consequences and how the design of policy and research should take this into account.

The first essay is an empirical examination of how tax subsidies for hybrid vehicles influenced the market price for those vehicles. As such, it represents a traditional tax incidence analysis. In addition to the substantive investigation, the essay makes a methodological contribution by demonstrating how partial identification (bounding) techniques can be employed to assess the effects of a subsidy when economic agents are able to manipulate whether or not they receive the subsidy. The empirical evidence suggests that consumers received the vast majority of the tax subsidy, which stands in contrast to the predictions of standard economic theory. The essay goes on to describe an economic model in which this result is possible and to analyze the implications for the economics of taxation.

The second essay, written with Rebecca M. Blank and Kerwin Kofi Charles, is an empirical examination of how individuals responded to state laws governing the minimum age at marriage. Our principle finding is that a substantial number of young people wishing to marry were able to avoid restrictive laws by traveling to a nearby state to marry or simply by misrepresenting their age to a government official. This has an important implication for administrative data, which is based on marriage certificates collected from local governments. We conclude that retrospective survey data collected decades later provides a better record of actual marriage behavior than contemporaneous administrative records. We argue that our findings can serve as a broad cautionary tale for empirical researchers who may otherwise assume that administrative data is the best source.

The third essay, written with Alexandra M. Resch and Paul N. Courant, is a theoretical examination of how government should provide higher education. The goal of the paper is to provide an intuitive model of the optimal provision of resources and allocation of students within a system of higher education. In contrast to previous research, our model is driven entirely by an assumption about the complementarity between resources and student ability. The essay provides a straightforward model with intuitive implications about optimal provision rules that we believe will prove useful for empirical researchers and policy-makers.

CHAPTER II

The Incidence of Tax Incentives for Hybrid Vehicles

2.1 Introduction

Federal and state governments in the United States have responded to rising concerns about the consequences of oil consumption in part by introducing tax subsidies for new vehicles that feature fuel-efficient technologies, including gas-electric hybrids. These policies aim to reduce oil consumption in the personal transportation sector, which accounts for 40% of gasoline consumption and 20% of greenhouse gas emissions (Environmental Protection Agency, 2007). At the federal level, the Energy Policy Act of 2005 introduced a substantial personal income tax credit for hybrids. At the state level, thirteen states have passed tax incentives for hybrids, and many others have considered similar actions. In this paper, I determine who benefits from tax incentives for hybrids using transaction level microdata to estimate the incidence of existing subsidies for the Toyota Prius.

I find that consumers capture nearly all of the benefits of tax subsidies. Transaction prices for the Prius did not change following changes of up to $2,650 in the value of the subsidy. Since consumers receive the subsidy directly from the government after their purchase, constant transaction prices imply that consumers capture the benefits of government intervention. This finding has implications for the evaluation of existing and future policies. It also has broader implications for the study of tax incidence because the result, when combined with additional facts about the market for the Prius, is difficult to explain in standard models and may therefore point to phenomena that influence tax incidence but are typically ignored in research.

This paper makes several contributions to public economics. First, the case of the Prius poses a challenge to standard incidence models. Through much of the sample period, the Prius was in excess demand, and consumers had to wait in a queue for several weeks to make a purchase. Given these wait lists, it is difficult to construct a conventional model that allows for a situation in which consumers capture the whole benefit from a tax subsidy. I argue that the most compelling explanation is that Toyota believed that charging higher prices while the subsidy was in place would reduce demand for hybrids in the future, even though higher prices would not change the quantity sold. I construct a stylized model in which a temporary tax has no effect on price when the producer is capacity constrained and current prices influence future demand. Current prices may influence future demand if car shopping is a costly search process and consumers use past prices to form beliefs about the prices they will be offered in the future. Alternatively, current prices may influence future demand for "behavioral" reasons, including a concern for fairness. Under either interpretation, the modeling exercise suggests that familiar features of taxation, such as the equivalence of incidence regardless of who remits a tax, may not hold in this environment. Such considerations are most likely to influence tax incidence in markets for new products or "green" products and in markets in which final prices are negotiated.

Second, the paper proposes a method for establishing bounds on the incidence of a tax when tax changes are anticipated. The federal tax credit introduced sharp changes in the value of the subsidy, and the date of these changes was known in advance. This induced a timing response among consumers, who moved their purchases into the more favorable tax windows. If consumers who adjust their timing in response to a tax change differ systematically from those who do not, then a comparison of prices before and after the tax change will give a biased estimate of tax incidence. To account for this, I employ a novel procedure to estimate an upper bound on price changes. The estimator developed here is similar to the bounding estimator of Lee (2005) in that it provides worst-case bounds to account for missing information, but Lee (2005) addresses a situation with a randomly assigned treatment and missing outcome data, whereas the estimator developed here applies to situations with complete outcome data and selection into the treated category. Anticipated tax changes are a common feature of policy; examples include changes to the capital gains tax rates in the Tax Reform Act of 1986, the tax holiday on repatriated corporate earnings included in the American Jobs Creation Act of 2004, and subsidies for energy saving home improvements in the Energy Policy Act of 2005. Economists have noted the importance of timing responses and the challenge they pose to estimation of permanent responses (Slemrod, 1992); the bounding method used in this paper may be useful to future research that analyzes tax policies, or other policies, with this structure.

Third, this paper is one of the first to analyze hybrid vehicles from an economic perspective. While hybrid vehicles remain a small part of the new car market (about 1.5% of all new cars sold in 2006 were hybrids), their market share has grown nearly 100% per year since 2001. Despite their growing importance to environmental policy, almost no economic research exists on hybrids. The exceptions are a pair of papers that document the correlation between Prius ownership and Green Party registration and other signals of environmental preferences (Kahn, 2007a,b) and a paper that correlates state tax incentives with sales volumes (Gallagher and Muehlegger, 2007).

Finally, this research fills a gap in policy evaluation. Through the third quarter of 2007, the hybrid vehicle tax credit cost the federal government about $785 million, with $394 million going to Priuses.[1] Knowing who benefited from these subsidies is necessary not only to evaluate the current policy, but also to inform future legislative action.[2] Many politicians have called for an increase in the federal credit or for a removal of its phase-out provision.[3] Hybrid tax incentives also remain an active policy consideration at the state level.[4] Moreover, future subsidies for other advanced technologies are likely to take a similar form.

The balance of the paper is organized as follows. Section 2.2 describes important features of the automobile market, with specific attention to the case of the Prius. Section 2.3 details the federal tax incentives, and section 2.4 describes the data used. Section 2.5 includes the principal analysis. Estimates provided in this section demonstrate that there was a clear timing response, indicating that agents were aware of the tax laws, but there was little or no price response, indicating that consumers captured the subsidy. Section 2.6 considers the possibility that there was heterogeneity among those responding on the timing margin and provides tax estimates that account for this. Overall, these results confirm that the subsidy primarily benefitted consumers. In section 2.7 a different empirical methodology is

[1] Author's calculations, based on sales data from *Automotive News*, assuming an 85% take-up rate. Hybrids produced by General Motors are not included because they report only aggregated numbers that include both hybrids and conventional models. This should have little impact on the estimate, however, because General Motors has sold very few hybrids.

[2] A full policy evaluation should consider not only incidence, but also the effect on total units sold. I return briefly to this issue in the conclusion.

[3] President Bush publicly supports the removal of the cap (Associated Press, 2006). In 2005, Rahm Emmanuel (D-IL) introduced a bill in the House of Representatives to increase the credit for domestic hybrids; Evan Bayh (D-IN), Joe Lieberman (D-CT) and Sam Brownback (R-KS) introduced a senate bill in 2006 to remove the phase-out provision from the existing credit; and a 2007 bill introduced in the House would give a $4,000 tax credit to plug-in hybrids (Union of Concerned Scientists, 2007).

[4] The list of states that have introduced a hybrid tax incentive bill in 2007 alone includes, but is not limited to, California, Florida, Hawaii, Iowa, Illinois, Indiana, Kansas, Kentucky, Massachusetts, Missouri, New Jersey, Ohio, Rhode Island, Tennessee, Vermont and Wisconsin (Union of Concerned Scientists, 2007).

used to estimate the incidence of state tax incentives. This yields a very similar estimate of incidence. In section 2.8 I describe why these results pose a challenge to standard models of tax incidence and outline an alternative model, in which current prices are assumed to influence future demand, that could explain the result. Section 2.9 concludes.

2.2 Some Facts About the New Vehicle Market and the Prius

The objective of this paper is to estimate how prices respond to changes in a tax subsidy. This requires accurate measurement of the relevant prices. In the new vehicle market, there are several different prices — the suggested retail price, the dealer invoice price and the final transaction price. In this section, I define these prices and describe how they vary across vehicles and over time. Knowing how these prices typically change is important for constructing counterfactual prices that would be expected in the absence of tax change. I also describe the history of the Prius, which is unusual in that it was in excess demand for several years. This excess demand is key to the final interpretation of the empirical estimates.

The price posted on the sticker of a vehicle is known as the Manufacturer's Suggested Retail Price (MSRP). The manufacturer sets the MSRP and rarely changes this during the model year. The MSRP also does not typically vary across localities, except through differences in destination charges (shipping fees). The MSRP must, by federal law, be posted on all vehicles.[5] Furthermore, a car may not be sold for more than the MSRP unless the dealer attaches a new sticker that includes an additional cost, called an "additional dealer markup" or "additional dealer profit." Dealers can and occasionally do post additional markups on vehicles in high demand.

[5] As part of the Dealer's Day in Court Act of 1956, sticker prices were made mandatory in order to curb a practice called "price-packing," in which customers were offered a high price on their trade-in. The trade-in premium was packaged into the price of the new car, for which the consumer had no reference price (Time Magazine, 1958).

The posting requirement, however, makes it difficult for an individual salesperson to negotiate a price above MSRP if management prefers not to add a markup. The posting requirement also makes it transparent to consumers that the markup is a price increase that is not a reflection of dealer cost.

The manufacturer charges dealers what is known as the dealer invoice price for each vehicle. Franchise law dictates that dealers be charged the same price for identical goods, with concession to cost differences, including local advertising fees or delivery fees.[6] As a result, the dealer invoice price for a specific vehicle varies little across dealerships. The dealer invoice price is also constant over the model year cycle, in most cases. Dealerships pay the invoice price for a vehicle when it is shipped to them. This means that dealers have already paid for all of the vehicles in inventory. Instead of changing the invoice price, manufacturers offer cash incentives to dealerships when they want to lower the effective price. This allows manufacturers to change the price of all vehicles in inventory, which have already been sold to dealers.

The true, economic transaction price to a consumer of a vehicle depends on not just the final price that appears on the contract, but also on manufacturer incentives and the trade-in allowance. Manufacturers do not change the MSRP within a model year, but they directly influence the transaction price by offering cash rebates and low interest financing to consumers. Furthermore, two identical cars sold at the same contract price can represent two different real outcomes, if the consumers receive different trade-in allowances for a given trade-in. The trade-in over or under allowance is the difference between the amount listed for the trade-in on the contract and the dealer's actual estimated cash value, which is used for entering the vehicle into their

[6]Manufacturers do create some price dispersion through the use of bonuses and rewards for meeting certain benchmarks and through variation in advertising fees across localities.

own inventory. To accurately measure the price of a new vehicle, one must account for cash rebates, the value of low-interest financing and the trade-in allowance. Fortunately, the data described in the next section allow me to account for all of these things at the level of the individual transaction.

At any point in time, final transactions prices vary across vehicles because idiosyncratic bargaining outcomes cause the dealer's margin to vary from vehicle to vehicle and from customer to customer. Over time, some movements in the average transaction price are regular and predictable. When a new vintage of a model is introduced, typically in the early fall, it begins selling at a relatively high price. In most cases, the price declines steadily over the model year, with prices falling at a 9.2% annual rate, on average (Copeland, Dunn and Hall, 2005).[7] Automobile sales also follow a pronounced weekly and monthly cycle. Friday and Saturday see the most sales; Monday and Thursday sales are higher than sales on Tuesday and Wednesday; and very few cars are sold on Sundays. There are also more cars sold at the end of the month, as manufacturers seek to hit sales volume targets. These selling cycles translate to less dramatic price cycles, with small price discounts measured on weekends and at the end of the month.[8] The data used in this paper include the exact date of each transaction, which enables me to control for these regular price movements.

Macroeconomic shocks also influence vehicle prices, and they may influence different vehicles in different ways. For example, the price of gasoline has different effects on vehicle demand according to fuel economy. Previous research has shown that vehicles with different fuel economy have different responses to gasoline price changes in both sales volume (Linn and Klier, 2007; West, 2007) and price (McManus, 2007). Research has also shown that gasoline prices affect used vehicle prices differently by

[7] The run of a vintage is known as the model year. It typically begins in the fall and ends after about 18 months.

[8] Busse, Silva-Risso and Zettelmeyer (2006) report a weekend price discount of $28 and an end of month discount of $56 in their main regressions.

fuel economy class (Kahn, 1986). It is therefore important to account for time period effects and the price of gasoline in estimation.

The Toyota Prius

The Toyota Prius is a parallel gas-electric hybrid car. A gas-electric hybrid vehicle has both a gasoline engine and an on-board electric power supply. In a parallel hybrid, both the electric motor and the gas engine are capable of supplying power directly to the wheels. Hybrid technology increases fuel economy by using the electric motor to power the vehicle at low speeds and by recharging the battery with energy generated in the braking process (Alson, Ellies and Ganss, 2005). The Prius is not plugged into an electrical outlet, and it is fueled by conventional gasoline.

The Prius was introduced on a small scale in the United States in model year 2001. The second generation Prius, which introduced the current, distinctive body style, debuted in model year 2004. Shortly after its introduction, the second generation Prius was in excess demand. This lasted until the 2007 model year became available, which was produced on a greater scale. Toyota originally planned to sell 20,000 Priuses per year in the United States. Through the first three quarters of 2007, they were on pace to sell 170,000.

Excess demand led to wait lists, which will turn out to be an important part of the story of how taxes influenced Prius prices. The microdata used in the remainder of this paper do not have a measure of wait times, but they do include a measure of the "days to turn" — the number of days that a vehicle was in the dealer's possession before being sold. Average days to turn should be inversely related to wait lists: the more common are wait lists the faster vehicles will turn over.

The median days to turn for all vehicles in the market is close to 60 days. Figure 2.1 shows the mean and median days to turn of the Prius for the 2003 to 2007 model years. The break between the first and second generation Priuses is apparent in the data. At the end of its run, the first generation Prius sold very slowly. The second generation Prius sold at a remarkable rate for several years. Between the end of 2003 and the fourth quarter of 2006, the Prius turned over around every 4 or 5 days. The 2007 model year began at this same low rate, and rose, starting in November 2006. This may indicate the end of wait lists in many locations. Average days to turn rose further through the new year, and then fell, stabilizing at a new rate of 10 days. The rate of turnover for the second generation Prius at its highest point was only one-third the market average.

This suggests that wait lists appeared at the end of 2003 and faded in late 2006. Two other sources of evidence corroborate the notion that the wait lists began to recede in November 2006. First, media reports declaring the end of wait lists began appearing in November 2006 (e.g., Woodyard (2006)). According to these reports, industry data indicated a rise in the days to turn of Priuses, and analysts inferred from this a reduction in wait lists; and, Toyota executives reported that, with the 2007 model year line, they had enough vehicles to meet demand.

Second, a small sample of data is available from HybridCars.com, which invited hybrid buyers to report their date of purchase, wait times and transaction prices in a public forum. Between April 2005 and August 2007, 253 Prius buyers reported their wait times. Figure 2.2 shows the weekly average reported wait times for this sample. The average wait time in this sample fell to about 2 weeks in September and October 2006, and then went to zero starting in November. Thus, this small sample of wait time data supports the conclusions from the more comprehensive turnover data.

There are several reasons to analyze the Prius in order to understand the hybrid market and the effect of tax incentives. First, the Prius, by itself, represents more than half of the entire hybrid market. Second, the Prius is the most distinct and well-known hybrid vehicle. Third, the market for the Prius is more mature than other hybrids, since it has been through several model year cycles. Finally, unlike most other hybrids, the Prius experienced three policy shifts in the sample period. I move now to a discussion of these policy changes.

2.3 Tax Incentives for Hybrid Vehicles

The federal government has subsidized hybrids through the individual income tax system for several years. Before 2006, the clean fuel vehicle deduction allowed consumers to deduct the "incremental cost" of a clean technology, including hybrids, from their income. This was an "above the line" deduction that could be claimed even if the taxpayer was not using itemized deductions. The incremental cost was capped at $2,000, and all available hybrids qualified for this maximum amount.[9] The Energy Policy Act of 2005, passed in August 2005, eliminated the deduction completely for tax year 2006, and replaced it with the more generous hybrid vehicle tax credit. The new law allowed a tax *credit* of up to $3,400, based not on incremental costs, but on estimated fuel savings. The value of the credit is based on two components: a fuel economy credit and a conservation credit. The fuel economy credit is worth up to $2,400 and is determined by the percentage gain of the hybrid in city rated fuel economy over a model year 2002 benchmark. The conservation credit is worth up to $1,000 and is based on the total number of gallons of gasoline saved over a 120,000 mile vehicle lifetime, as compared to the same benchmark. To qualify for the credit, hybrids also must meet stricter emissions standards. Manufacturers are required to

[9]The IRS received many questions about how to calculate incremental cost. In October 2002, the IRS decided to allow filers to use incremental cost estimates provided by the auto manufacturers (Internal Revenue Service, 2002). The result was that all hybrid vehicles qualified for the full $2,000 maximum.

send applications for each model year to the IRS, and the IRS then sets the exact benefit.[10]

Not every person who purchases a hybrid is eligible to receive the benefit. The credit does not offset tax obligations for those paying the Alternative Minimum Tax. Furthermore, the credit is non-refundable. (Note that these same restrictions effectively applied to the clean fuel vehicle deduction in earlier years.) If a vehicle is leased, the lessor may claim the credit, not the lessee. The seller may claim the credit if the buyer is a tax-exempt entity. The credit is available to business purchasers.

The Energy Policy Act of 2005 also included a phase-out provision. The provision is triggered when a manufacturer sells 60,000 eligible vehicles. The credit is unchanged in the quarter in which the 60,000th vehicle is sold and in the next quarter. After that, the credit falls to 50% of its original value for the next 2 quarters, then 25% for another half year, and then expires completely.

The phase-out was allegedly designed to prevent foreign automakers from benefiting more than domestic automakers over the life of the program (Lazzari, 2006; Leonhardt, 2006). Toyota hit the 60,000 mark in the second quarter of 2006, triggering a cut in credit amounts that began on October 1, 2006. The benefit fell again on April 1, 2007, and it expired completely on October 1, 2007. Honda's phase-out began on January 1, 2008. No domestic automaker has reached the cap.

The sharp date changes provide a natural experiment for the analysis of the influence of the tax policy on the hybrid market. A Prius purchased on or before December 31, 2005 was eligible for a $2,000 deduction (worth at least $500 for households in a middle income bracket and up to $720 for the highest income individuals in 2000). A Prius purchased between January 1, 2006 and September 30, 2006 was eligible for

[10]The certification process has created information lags, which are described in greater detail in section 2.8.

a $3,150 credit. A new Prius purchased between October 1, 2006 and March 31, 2007 was eligible for $1,575. A Prius sold between April 1, 2007 and October 1, 2007 garnered a credit of only $787.50. Table 2.1 summarizes the variation in the federal policy. I use this variation in subsidy size over time to identify the incidence of the federal income tax credit. A number of states have also legislated a state income tax credit or a sales tax exemption for hybrids. These policies are described in more detail in section 2.7.

2.4 A Description of the Data

This paper uses data from J.D. Power & Associates' Power Information Network (henceforth PIN).[11] J.D. Power collects transaction details directly from a large sample of dealers in all major markets in the United States. The sample includes about 15% of all new car sales to final consumers. PIN data include the price of each vehicle sold, the exact date of the sale, financing details, cash rebates and the truncated Vehicle Identification Number (VIN). The data also include age, sex and state of residence of purchasers.[12] The version of the PIN data released to me suppresses the personal identifying information of consumers and the dealer information that would allow identification of specific dealerships, to preserve confidentiality. I restrict the sample to purchased (not leased) vehicles with complete price information.[13] The sample excludes fleet sales. The final sample begins in the fall of 2002 and ends in May 2007.

The VIN of the new vehicle identifies the make, model, model year, and a vari-

[11]Standard data sources that include vehicle purchases and demographic detail will generally have too few hybrids for analysis. There are roughly 100 million households in the United States, and about 16 million new vehicles are sold annually. The Consumer Expenditure Survey (CEX) has about 100,000 consumer units per year. Since the Prius comprises about 1% of the market, one expects roughly 160 Priuses per year in the CEX. In addition, standard data sets have a substantial lag. The most recent CEX is from 2005, and the most recent National Household Travel Survey is from 2001.

[12]Sex is imputed from first names, and ambiguous first names are thus missing.

[13]Less than 3.5% of Priuses in the sample were leased.

ety of other characteristics. The VIN includes information on engine displacement, cylinders, transmission, doors, body type and trim level, but it does not detail all available options (e.g., sun roofs or stereo systems). All factory and dealer installed options, advertising fees and delivery fees, are reflected in a measure called the dealer cost, which is essential for allowing price comparisons of similarly equipped vehicles. Note that dealer cost is not identical to dealer invoice (see the appendix for details).

The data include a number of additional variables. These include the trade-in vintage and odometer. Also important is the "days to turn" — the number of days that a vehicle was on the dealer's lot before being sold. Transaction details also include information on service contracts, interest rates and other loan details.

Two additional pieces of information come from other sources. I merge official fuel economy ratings from the Environmental Protection Agency's fuel economy guide according to make, model, model year, cylinders, displacement, transmission and trim level. I also use weekly tax inclusive national retail gasoline prices provided by the Energy Information Administration.

As described in section 2.2, the final transaction price of a new car must account for several things. To construct the necessary measure, I begin with the transaction price inclusive of factory installed options but exclusive of taxes, fees, service contracts and after-market options. The excluded items do not generally influence the resale value of the vehicle, and service contracts and after-market options are often negotiated after the final price of the vehicle is determined. I adjust the transaction price inclusive of factory installed options to account for the trade-in allowance, manufacturer's direct to customer cash rebates, and financing incentives. To construct the value of the financing incentive, I assume a 4% annual nominal discount rate and estimate the difference in the present discounted value of the loan actually observed in the data (given the loan amount, interest rate and term) to that same value using the Federal Reserve's 48-month car loan interest rate series.[14] I call this the "incentive adjusted price," or just "price".

This measure of price will vary with the set of options installed on a vehicle, creating significant price variation. For example, the 2007 Prius has a base MSRP of $22,175, but the premium options package adds $6,350 to the retail sticker price. If the number of Priuses that have this options package changes from week to week, average prices will change, even if the prices of identical vehicles do not.

To construct a measure that accounts for the value of options, I also use a measure which subtracts dealer cost from the incentive adjusted price. I call this the "incentive adjusted markup," or just "markup". Note that this is a level amount, not a percentage. Note also that this markup does not distinguish between revenue collected by the dealer and the manufacturer. For example, a dealer's gross profit could rise when the markup falls, if the manufacturer offered a direct to consumer incentive that lowered the final price by less than the full incentive amount.

Table 2.2 shows the sample size, price and markup for the final sample of 64,706 Priuses by model year. The 2007 model year features a premium trim level "Touring" package, which is listed separately in table 2.2. The last column in each panel shows totals for the entire model year. For example, the sample contains 2,381 model year 2003 Priuses. They sold for an average price of $21,068, which was $611 above the average dealer cost.

The main estimates in this paper focus on transactions that occurred within two weeks on either side of each federal tax change. Table 2.2 breaks down the number of observations, average prices and average markups in each of these windows by

[14]This price adjustment methodology follows Corrado, Dunn and Otoo (2006).

model year and trim level. The first two tax changes are close to changes in model years. The 2006 model year Prius began selling in November 2005, and the 2007 vintage began selling in September 2006. The latter is evidenced in table 2.2 by the sharp drop in 2006 model year Priuses between the end of September and beginning of October.

Since it is critical to compare identical vehicles before and after each tax change, I estimate the effect of the first tax change using only 2006 model year Priuses, and I estimate the effect of the other two changes using only 2007 Priuses with the base trim level. These samples are indicated by shading in table 2.2. Thus, for example, there are 513 Priuses in the sample in the last two weeks of December 2005 and 1,007 in the first two weeks of January 2006, but I use only 433 from December and 925 from January in the principal analysis, to restrict comparisons to a single model year. Except where noted, analysis of all the model years and trim levels together does not change any result significantly. Having described the market, the policy variation and the data, I now proceed to a description of how consumers responded to tax incentives.

2.5 Responses to Changes in the Federal Tax Credit

Transactions surrounding each of the three changes in the federal tax credit reveal two key facts. First, there was a clear shifting of transactions into the tax preferred time period in each case. Second, Priuses purchased in the higher subsidy period, near a tax change, were sold for about as much as Priuses purchased in the lower subsidy window near that change. Since consumers later receive a tax break from the government, constant transaction prices imply that net of tax consumer prices move by the full amount of the subsidy. Together, these two facts suggest that agents were aware of the tax policy and that consumers capture most of the benefit.

Agents Shifted Transactions Into High Subsidy Time Periods

The sales distributions surrounding the policy changes reveal strategic shifting. Shifting is limited, however, to the two weeks before and after each tax change. Figure 2.3 shows the distribution of sales over the 62 days in December 2005 and January 2006 for both the Prius and for non-hybrid Toyota sedans.[15] Priuses purchased in December were eligible for the deduction (worth up to $720), and January purchases were eligible for the credit (worth up to $3150). Relative to Toyota sedans, the Prius was sold more heavily at the beginning of January than at the end of December, suggesting that transactions were shifted into January, where the subsidy was greater.[16]

To ensure that the extra January Prius sales are a result of the tax and not a regular time pattern specific to the Prius, we can examine the sales distribution over the same dates in a year with no tax change. Figure 2.4 shows the distribution of sales of the Prius versus non-hybrid Toyotas for December 2004 and January 2005, where there was no change in tax incentives. These data do not show a similar spike in January sales.

Figures 2.5 and 2.6 show the corresponding distributions for September and October. As expected, Prius sales were abnormally high in September 2006, just before the tax credit phased out. In this instance, Prius transactions were moved to an earlier date, whereas the December and January picture shows that Prius transactions were moved to a later date. This symmetry of response bolsters the conclusion that differences between sales patterns of the Prius and non-hybrid Toyota sedans were

[15]Results are very similar if the entire sedan market is used instead of Toyota sedans.

[16]January 1, 2006 was a Sunday, accounting for the low sales. The selling cycle described in section 2.2 is evident in these figures.

driven by tax changes. Figures 2.7 and 2.8 provide the same information for March and April 2007, and lead to the same conclusion.

Table 2.3 quantifies the shift in purchase timing. Relative to the monthly sales volume, the shift is large — about 15% of monthly sales are shifted on average over the three events. In addition, the relative shift is larger for the largest tax change and smallest for the smallest tax change. Furthermore, the sales distribution occurred entirely with the two-week window around each change. The distributions are very similar at the beginning and end of the two-month windows. Given the significant amount of money at stake, it is perhaps surprising that more shifting did not occur.

Average Prices Were Non-Responsive: Descriptive Evidence

Figures 2.9 through 2.11 show seven-day moving averages of prices and markups in the two months surrounding each tax change.[17] The samples in these figures are restricted to a single model year and trim level to avoid introducing price changes that are due to model year seams. Overall, these pictures provide no evidence that prices moved in the anticipated direction in response to tax changes, but there are important differences between prices and markups. Before discussing the graphical evidence, I discuss the interpretation of a discrepancy between price and markup.

Price versus Markup

If average prices and markups fail to move together, it must be because (a) the average value of installed options changed or (b) Toyota changed the dealer invoice (the price it charged dealers) for identically equipped vehicles. Around the first two tax changes, dealer invoice did not change, which means that differences between prices and markups are driven by options packages and that the preferred measure

[17]The moving average is symmetric, using three days on either side of the date. To avoid smoothing a discontinuity, the moving average does not include values from both tax regimes at any point.

for incidence analysis is the markup. Ultimately, the same is true of the third tax event, but price changes in that case are more complicated, as explained below.

Recall from section 2.2 that dealer invoice prices rarely change during the model year, except through dealer cash incentives. Industry sources that report weekly on these incentives indicate that no dealer incentives were in place for the Prius during any of the tax windows.[18] The Toyota dealers that I spoke with indicated that they have never received dealer cash incentives on the Prius. In addition, the distributions of costs in the sample show evidence of changes in the proportion of different options packages, but generally support the notion that Priuses with the same options package cost the same around each of the first two tax windows (see Appendix for details).

Dealer cost did, however, change around the third tax date. In February and March of 2007, Toyota introduced low-interest financing for the Prius. These incentives were eliminated on April 2, 2007, just as the tax credit dropped from $1575 to $787.50.[19] Toyota simultaneously announced an "economic savings bonus" for the Toyota Prius which took effect on April 2. This "economic savings bonus" lowered the price of each options package by as little as $600 or as much as $2000. The dealer invoice for each options package was also changed as part of this program.

To determine whether consumers paid more or less for similarly equipped vehicles in April 2007, as compared to March, we need to know how much each April Prius would have cost dealers had it been sold in March, before the price change. To

[18] Specifically, I examined the "Dealer Incentives" table in relevant issues of *Automotive News*, which indicate no dealer incentives for the Prius in December 2005, January 2006, September 2006, October 2006 or March 2007 (Automotive News, 2005 - 2007). *Automotive News* does report an incentive in April 2007, but this may have indicated either the "economic savings bonus" or a reimbursement for the costs of using a Prius as a loaner vehicle if the customer later bought a Prius (Rechtin, 2007), in which case the incentive would have influenced a tiny fraction of vehicles. In a personal correspondence, an *Automotive News* employee has indicated to me that they believe no dealer cash incentive existed in April 2007, though they were unable to explain what was reported in the magazine at that time.

[19] April 1, 2007 was a Sunday.

do this, I use dealer invoice prices for each options package from Edmunds.com to identify which options package each vehicle most likely had. I then adjust for the change in invoice price so that a Prius with the same options package in March 2007 and April 2007 have the same dealer cost. Details of this procedure are in the Appendix. Costs adjusted for the invoice price change are used throughout the remainder of the paper.

Graphical Results

With the distinction between prices and markups in mind, I move now to a description of the price response to tax changes. Figure 2.9 shows prices and profits for December 2005 (when the deduction was worth up to $720) and January 2006 (when the credit was worth up to $3150). If sellers captured part of the benefit, transaction prices should *rise* after the vertical bar in Figure 2.9. Instead, a small jump downward is evident in the price series, and the markup is unchanged. Since dealer cost was constant, this means that the tax subsidy was captured by consumers. The drop in the price series was due to a fall in the average value of installed options.

Figure 2.10 shows prices and markups for September 2006 (when the credit was worth up to $3150) and October 2006 (when the credit was worth up to $1575). There is no evidence of a price or markup jump around this tax change. Within a few weeks of the tax date, the markup is quite steady. Note that the series do not start at the beginning of September in figure 2.10. This is due to the model year seam. Toyota began selling the 2007 Prius in the middle of September. While there are not very many days available before the tax change, there are about 500 Priuses in the sample, as shown in table 2.2. Many 2006 Priuses were also sold in September, but only about 100 were sold in October, limiting the potential for analysis of the 2006 vehicles.

Figure 2.11 shows prices and markups, using the cost adjusted for the "economic savings bonus," for March 2007 (when the credit was worth $1575) and April 2007 (when the credit was worth $787.50). In this case, there is evidence of a price *increase* when the tax credit is reduced — in contrast to our expectation that a decrease in tax credits would lead to a fall in price. This is due to the fact that the financing incentives in place in March were worth more than the "economic savings bonus."

Two explanations for this unexpected price change are possible. First, I may have miscalculated the value of financing incentives, which require an assumption about the discount rate and the use of a counterfactual interest rate. In the econometric results below, I recalculate incidence estimates using a variety of assumptions. When a 6% or 8% discount rate is used, there is no significant jump in the markup.

Alternatively, consumers may have miscalculated the value of financing incentives. Average prices that do not account for financing incentives fall through this tax change. The price increase evident in Figure 2.11 is entirely driven by the elimination of low-interest financing. Thus, it is possible that consumers mistakenly perceived the "economic savings bonus," which was a visible change in price posted on the sticker of the vehicle, as a price cut that offset the change in tax credit, but failed to properly take into account the price increase embedded in the financing rates.

Other research bolsters the plausibility of consumer miscalculation. Busse, Simester and Zettelmeyer (2007) find striking evidence of consumer misunderstanding surrounding the employee pricing discount promotions of domestic vehicle manufacturers in 2005. They show that, in response to announced promotions in which all consumer were eligible to purchase cars at the employee discounted price, sales surged dramatically (though temporarily), but prices *rose* on average. This was because the manufacturer's incentives in place prior to the promotion were so large that they provided a bigger discount than employee pricing in many cases. The authors point to experimental research that shows that the announcement of a sale can sometimes have an effect on demand, even if prices are not, in fact, discounted (Inman, McAlister and Hoyer, 1990).

In sum, there is considerable graphical evidence that consumers captured the full tax subsidy because markups did not jump in response to tax changes. There are also interesting differences between prices and markups, which indicate the importance of accounting for options packages. The only evidence of a change in markups is around the third tax event, at which point prices moved in the "wrong" direction. This may be the result of imperfect calculation of the value of financing incentives by consumers, which were changed on the same day that the tax changed.

Average Prices Were Non-Responsive: Regression Evidence

In this section, the conclusions drawn from the graphical evidence described above are confirmed via linear regression. These regressions show no evidence of a statistically significant, positive relationship between prices and the value of the federal tax subsidy. It is thus impossible to reject the hypothesis that consumers captured the entire credit.

Table 2.4 reports estimates from regressions of the following form:

$$Price_{ij} = \frac{\alpha_1}{\Delta_1}H_1 + \gamma_1 W_1 + \frac{\alpha_2}{\Delta_2}H_2 + \gamma_2 W_2 + \frac{\alpha_3}{\Delta_3}H_3 + \gamma_3 W_3 + \mathbf{X_{ij}}\beta + \mu_{\mathbf{j}} + \varepsilon_{ij},$$

where i indexes an individual and j a model year and trim level interaction, the other subscripts denote the three tax changes (1=December 2005 and January 2006; 2=September 2006 and October 2006; 3=March 2007 and April 2007), Δ is the dollar amount of the tax change in each case, H is a dummy equal to 1 if the transaction is for the analyzed model year (see table 2.2) and occurs in the *high* tax side within two weeks of a change, W is a dummy equal to 1 if the transaction occurs on either side within two weeks of the change, $\mathbf{X}$ is a vector of controls, μ is a vector of dummies for each model year and trim level, and ε is an error term.[20] Thus, for example, H_1 is coded as 1 if and only if the transaction is for a model year 2006 Prius sold in the first two weeks of January 2006. W_1 is coded as 1 if and only if the transaction is for a model year 2006 Prius sold in either the last two weeks of December 2005 or the first two weeks of January 2006. The dependent variable is the incentive adjusted price paid by consumers.

The sample includes all model year 2003 to 2007 Priuses (N=64,706). The Priuses not in the treatment windows add precision to the estimation of the coefficients on the control variables. The vector of controls includes vehicle cost, day of the week dummies, a dummy for the last five days in a month, state dummies, the retail price of gasoline and a quadratic trend in the length of time that type of Prius was available on the market.

The α/Δ's are the coefficients of interest. Since the α/Δ's are coefficients on the high subsidy time period dummies, if there is a price change in the expected direction, the α/Δ coefficients will all be *positive*. Furthermore, since the high subsidy period dummies are divided by the dollar amount of the tax change, the estimates may be interpreted as the *change in price per dollar of tax change*. The standard tax incidence model thus predicts that $0 \leq \alpha/\Delta \leq 1$. In calculating the size of each tax change, I use the maximum credit amount (\$3150, \$1575 and \$787.50) and I assume a 25% marginal tax rate for the deduction, meaning it is assumed to have been worth \$500.[21]

[20] Here, and elsewhere, the i and j subscripts on the H and W dummies are suppressed to avoid clutter.

[21] The 25% tax rate fits the most appropriate income range. For example, married couples filing jointly with adjusted gross income between \$59,400 and \$119,950 have a marginal tax rate of 25%. This range includes the mean self-reported income of Prius buyers in 2005 (\$87,500) from marketing research data (CNW Research, 2007). In 2005, 21% of all tax filers had a marginal tax rate of 25%, and only 6.4% had a marginal rate above this (see Table

Table 2.4 reports the α/Δ coefficients for several specifications. The first column contains only the model year and trim level dummies as controls. It is, therefore, simply the difference in means of vehicle prices in the 2 weeks with a higher subsidy value and the 2 weeks with a lower subsidy value, adjacent to each tax change, scaled by the change in subsidy value. The second column adjusts for options package composition by including the dealer's cost of the vehicle as a regressor. The third column includes vehicle cost, and adds all of the other controls.

As an example of how to interpret the coefficients, consider the estimate in the first row and second column: controlling for dealer cost, prices fell by 3.7 cents for every increased dollar of subsidy when the tax credit was introduced in January 2006. Given the standard error of 2.4 cents, this is not statistically distinguishable from zero. Since dealer cost was constant in the tax windows (argued above), the specifications that control for dealer cost (columns 2 and 3) are preferred. In these specifications, the price effects for the first two changes are statistically indistinguishable from zero, and are tightly estimated. The third tax change has a statistically significant sign in the unexpected direction. As noted below, the statistical significance of this coefficient's difference from zero is sensitive to alternative definitions of financing subsidies.

In sum, regression evidence indicates that the subsidies did not generate any statistically significant upward price movement in the price that consumers paid, so that one cannot reject the hypothesis that consumers captured the entire subsidy. The possibility that these estimates are biased downwards because "savvier" car buyers were more likely to move their transaction into the higher subsidy tax window is explored in the next section.

3.4, Internal Revenue Service (2007a)). The top marginal rate in 2005 was 35%, which would raise the value of the subsidy to $700.

2.6 Accounting for Heterogeneity in Strategic Timing

The incidence estimates of the previous section are based on a comparison of transactions before and after each tax change. Such a procedure could lead to biased estimates if consumers who reacted to tax changes by moving their date of purchase differ systematically from those who did not. In this section, I provide evidence that there was a difference between consumers who purchased in the high and low tax windows, discuss a procedure for establishing an upper bound on the price effect of taxes when this type of heterogeneity is present, and generate estimates of this upper bound in the case of the Prius. For the first two tax changes, these upper bounds are fairly tight, which confirms that consumers captured a significant majority of the tax benefits.

Evidence of Heterogeneity in Strategic Timing

Even though the data contain relatively few buyer demographic variables, the transaction details are sufficiently rich to reveal differences between consumers who purchased in the high subsidy window and the low subsidy window. Table 2.5 shows mean characteristics of transactions that occurred within two weeks on either side of each tax change. The table also reports an estimate of the difference in mean characteristics of high subsidy and low subsidy transactions taken from a regression of the following form:

$$Characteristic_{ij} = \beta H + \gamma_1 W_1 + \gamma_2 W_2 + \gamma_3 W_3 + \varepsilon_{ij},$$

where i indexes individual consumers and j model year and trim level interactions, H is a dummy equal to one if the Prius is the model year analyzed and the transaction is within two weeks of a tax change on the high subsidy side, and W are dummies equal to one if the vehicle is the model year analyzed and the transaction is within two weeks on either side of each of the three tax events.[22]

Table 2.5 shows that high subsidy consumers were less likely to trade in a vehicle, required a smaller down payment, were less likely to buy a service contract, generated less service contract income for dealers, paid lower total expenses on after-market options, taxes and fees, and were less likely to accept an interest rate that exceeded the buy rate – as compared to low subsidy consumers.[23] High and low subsidy transactions are statistically indistinguishable from each other in age, sex, trade-in vintage, trade-in value, amount financed and whether or not the transaction included life insurance.[24]

Thus, high subsidy car buyers differ from their counterparts in ways that suggest they are better negotiators. Intuitively, consumers with excellent knowledge of the car market are probably less likely to agree to a transaction that allows dealerships to generate income through service contracts, interest rate markups and after-market options. In addition, previous research has found that consumers who trade in a vehicle pay more for their new car (Scott Morton, Zettelmeyer and Silva-Risso, 2001). Thus, table 2.5 demonstrates not only that a difference exists between high and low subsidy buyers, but also that consumers who purchase in the high tax window are likely to be better negotiators, which means that they will pay less for an identical car than their counterparts. If true, then the estimates in table 2.4 underestimate the true tax effect, since high subsidy buyers (a) likely pay less on average than low subsidy buyers within any given tax regime and (b) are more likely to buy a car in

[22]All results are very similar if (a) the high and low transactions are all grouped together without separate dummies, (b) Prius transactions from a previous year with no tax change are used to construct a difference-in-differences estimate, or (c) non-hybrid Toyota sedans are used as a comparison group to construct a difference-in-differences estimate.

[23]The buy rate is the rate that the financing agency quotes to the dealership. Dealers sometimes sign a contract for a higher rate than the buy rate, allowing them to make income on the difference.

[24]Dealerships offer life insurance that offsets debt on a car loan in the event of death.

the high tax benefit period.

Bounding Estimates of the Price Effect of the Tax Credit

To understand the implications of heterogeneous shifting for the estimation of tax incidence, consider the following model. Suppose that there are two types of Prius buyers, "movers" and "stayers". Each individual has an exogenously assigned ideal date of purchase drawn from a distribution common to both movers and stayers. If a mover has an ideal date of purchase near a tax change, that individual will move their transaction into the high subsidy time period. In contrast, a stayer makes their purchase on the ideal date, regardless of pending tax changes. Suppose also that, consistent with the evidence in table 2.5, movers are better negotiators, on average. Under this assumption, a comparison of transactions on either side of a tax change would result in a biased estimate of the tax effect because it compares buyers with different average bargaining abilities.

The estimate obtained by comparing transactions on either side of a tax change can be written and decomposed as follows. Let m denote movers and s stayers, and let τ^H denote the high tax subsidy time period and τ^L the low subsidy time period. Denoting the price as p, the standard estimate is:

$$\begin{aligned} E[p|\tau^H] - E[p|\tau^L] &= \rho E[p|\tau^H, m] + (1-\rho)E[p|\tau^H, s] - E[p|\tau^L, s] \\ &= \rho(E[p|\tau^H, s] - \zeta) + (1-\rho)E[p|\tau^H, s] - E[p|\tau^L, s] \\ &= \theta - \rho\zeta, \end{aligned}$$

where E is the expectations operator, ρ is the percentage of those in the high subsidy period that are movers, $\theta \equiv E[p|\tau^H, s] - E[p|\tau^L, s]$ is the treatment effect of the tax on stayers and $\zeta \equiv E[p|\tau^H, s] - E[p|\tau^H, m]$ is the difference in mean price paid between stayers and movers in the high subsidy period. If movers are, on average, better negotiators than stayers, then $\zeta > 0$ and the observed effect is a downwards biased estimate of θ.

This simple model demonstrates three important things. First, the parameter that could potentially be uncovered is the treatment effect on stayers (θ). Since movers are, by definition, observed in only one period, one cannot hope to directly recover the tax effect on movers, without auxiliary assumptions. Second, if $\zeta > 0$, then the standard before/after comparison (as in table 2.4) is an underestimate of the tax effect on stayers. Third, two pieces of information are required to recover θ: the percentage of movers in the high subsidy period (ρ) and the premium paid by stayers (ζ).

Even if a defensible estimate of ζ is not available, however, one can still bound θ under the following assumptions. First, suppose that the lowest price obtained by a stayer is higher than the highest price paid by a mover. This is stated in assumption 1:

Assumption 1. $\min[p|\tau^H, s] \geq \max[p|\tau^H, m]$.

Under this assumption, one can identify all the stayers in the high subsidy period, given information about how many stayers and movers are present in the period, because they will be those with the highest prices. This is an extreme assumption, since it supposes that all movers are better negotiators than all stayers.

Also required is an estimate of the proportion of stayers present in the high subsidy period (ρ). A plausible estimate of the number of stayers in the high tax period is the number of stayers in the low tax period. This would be true if stayers purchase vehicles at a constant rate within four weeks of each tax change. This is stated in assumption 2 as:

Assumption 2. $n_s^H = n_s^L$,

where n is used to denote the size of the sample in each period.

Combining assumptions 1 and 2, we can identify the prices paid by stayers in each high tax period — they are simply the n_s^L highest prices. Given the extremity of assumption 1, the mean of the n_s^L highest prices in the high subsidy period should *overestimate* the mean price of stayers in the high subsidy period. Combining this overestimate with the underestimate from the standard before/after comparison provides a set of bounds on the true tax effect on stayers.

This methodology is an example of partial identification in the presence of corrupted data, which is analyzed by Horowitz and Manski (1995) and extended by Horowitz and Manski (1997). The corrupted data model supposes that the observed data are a mixture of true data and noise. Given information about the probability of noise, Horowitz and Manski (1995) show how to construct bounds on various parameters. In the present case, prices observed in a tax-favored window are a mixture of true data (prices paid by stayers) and noise (prices paid by movers). Assumption 2 is necessary to provide an estimate of the probability of noise. This methodology has not been previously applied to the literature on tax incidence, and, more generally, has rarely been applied (an important exception is Hotz, Mullin and Sanders (1997)).

The estimator employed here is also related to the bounding technique developed in Lee (2005), which can also be cast as a corrupted data problem. Lee (2005) addresses a situation where the effect of a randomly assigned treatment is only partially identified because some outcome data are missing. In contrast, the approach developed here addresses a situation where selection into the treatment group may occur on unobservable characteristics, but all outcomes are observed.

Next, if the tax effect on stayers is the same as the tax effect on movers, then we have identified the full parameter of interest. This is stated in assumption 3:

Assumption 3. $E[p|\tau^H, m] - E[p|\tau^L, m] = E[p|\tau^H, s] - E[p|\tau^L, s]$.

An individual consumer need not be aware of or responsive to the tax credit in order for it to influence their final price. This is the essence of an incidence model — any individual's price depends not on their own reaction to the tax, but rather on the reactions of others. One way to think about assumption 3 is the following. Overall supply and demand (which is influenced by the tax regime) determine some reference price in each period. Those with high bargaining ability pay some amount less than this reference price, and those with low bargaining ability pay some amount more. In such a case assumption 3 is natural; people at all bargaining abilities experience a shift in the reference price, due to the shift in demand.

Table 2.6 shows results from this exercise, with varying controls. The first three columns match the control specifications in table 2.4. Thus, the lower bound reported in these columns is the corresponding point estimate in table 2.4. The fourth column contains a variety of additional controls that may indicate individual bargaining ability: sex, age, total after-market options, a dummy for an APR above the buy rate, a dummy for life insurance, a dummy for the presence of a trade-in, trade-in actual value, trade-in vintage, and a dummy for a service contract. The reported standard error on the lower bound is a standard estimate from a regression. The standard error on the upper bound is obtained via nonparametric bootstrap, with 5,000 repetitions.

The upper bound is constructed by collecting the n highest residuals from the high subsidy period from the regression that generates the lower bound, where n is the number of observations in the corresponding low subsidy window. The mean of these residuals is scaled by the size of the tax change and added to the lower bound to generate the upper bound. Given the scaling, both the lower and upper bound may be interpreted as dollar price changes per dollar of tax change.

Table 2.6 demonstrates that, even under the extreme assumption regarding heterogeneous shifting, large price responses to the tax change can be ruled out in the first two tax changes, whenever options are controlled (columns 2, 3 and 4). According to column 3, the upper bound estimates for the first two tax changes indicate that consumers got *at least* 73% of the gains around the first tax change and 92% of the gains around the second.

The estimated bounds for the third tax event, however, have little bite. Mechanically, this is due to the larger price variation observed in April 2007. This larger variation is partly due to the financing incentives: as long as some consumers pay cash, the presence of low-interest financing will increase price variation. Economic explanations, which revolve around waiting lists, of the varying results across tax changes are explored in detail below. To preview, a likely explanation is that there were wait lists for the Prius during the first and second tax change, which meant that prices were inflexible. Wait lists had disappeared by the third tax change, making price movements possible, which induced greater variation and potentially a larger tax effect.

In sum, analysis of the response to the federal tax credit suggests that subsidies had very little effect on prices paid by consumers to dealers. Thus the subsidy, during the first two tax changes, almost exclusively benefitted consumers. Before moving to an analysis of state tax incentives, I provide some robustness checks on these results.

Robustness

The bounding estimates in table 2.6 use only one model year and trim level to analyze each tax change, employ a particular definition of financing incentives and do not account for aggregate shocks to the market directly, other than controlling for the price of gasoline. In this subsection, I briefly address each of these issues and discuss how they influence the range of possible estimates of tax incidence.

Accounting for Other Model Years

The first two tax changes occur near seams in the model year. In the estimation above, I use only one model year and trim level in the estimation in order to avoid conflating tax effects with price changes due to the model year seam. This will influence an upper bound estimate if it influences the relative number of observations in the high and low tax period.

To understand how the upper bound estimator is sensitive to relative sample sizes, consider the following extremes. Suppose that there was only one Prius purchased in the low subsidy window around a tax change. In that case, the upper bound estimate would be the highest priced Prius in the high subsidy window minus the price of the lone Prius from the low subsidy period. Alternatively, suppose that there were exactly as many Priuses in the low subsidy window as the high subsidy window. In this case, the upper bound estimate would be identical to the lower bound estimate, since the relative sample sizes indicate that no shifting took place. The choice to use only a single model year, therefore, has consequences for upper bound estimates to the extent that using only a single model year changes the relative sample sizes in the tax windows.

Rather than just lumping all of the Priuses together to address this concern, which might conflate compositional changes with tax effects, I adjust the upper bound estimate to account for the proportion of Priuses sold in each window. For example, in the second tax change, the upper bound estimate in table 2.2 uses the 383 highest prices from late September (high subsidy) to form an estimate of the September price for stayers, because there were 383 Priuses in the October (low subsidy) window. Instead of using that number, I use the proportion of Priuses of all model years that were sold in September to determine how many September sales represent stayers. Since 495 Priuses of all vintages were sold in October and 1,203 were sold in September, the October sales were 41% as many as the September sales. Using that proportion, the number of stayers estimated to be in the September sample is 189. This will mechanically raise the estimated September price of stayers.

Table 2.7 replicates the bounding results with this alternative methodology. As expected, the upper bound on the second tax event rises noticeably. Even so, however, it is an informative upper bound, which says that, even under the most conservative assumptions, consumers captured *at least* half of the tax subsidy. As is evident from a comparison of 2.6 and 2.7, this procedure actually tightens the bound on the first tax change, and it has a minimal influence on the third tax change.

Calculation of Financing Incentives

The measurement of financing incentives necessarily involves assumptions about the discount rate and the counterfactual interest rate. In the case of the first two tax changes, these assumptions are unlikely to exert much influence on estimation, since there was no significant change in financing incentives near the tax date. In the third instance, however, financing incentives changed, making the estimates more sensitive to the assumptions used.

Recall that financing incentives are estimated by comparing the APR that a person receives on their vehicle loan from the lending agency with a prevailing market APR, based on the Federal Reserve's survey of commercial lenders. I calculate the monthly payments, given the amount of the loan and the term, using both interest rates, and then compute the present discounted value of the difference in these monthly payment streams using a discount rate of 4%. In the robustness checks, I vary this discount rate, and recompute the bounding estimates using a discount rate of 2%, 4%, 6% and 8%.

The Federal Reserve's benchmark interest rate is a market wide average. It may not be an excellent counterfactual for any given individual, but it should be accurate for the average car buyer. If Prius buyers, however, are different *on average* than other car buyers in terms of credit, then this market average will be a biased estimate for the average Prius buyer. To account for this possibility, I calculate the difference between the average APR of all Prius buyers and the market rate, during the wait list time period, at which time one expects that Toyota was not subsidizing loans. For the 2004, 2005 and 2006 model years, Prius buyers received an average APR that was one-half point below the market rate, suggesting that they have better credit than the average new car buyer. I construct a modified counterfactual market interest rate by subtracting this premium from the Federal Reserve's survey rate.

One other approach is suggested by Corrado et al. (2006). They use data from the Survey of Consumer Finances to estimate the determinants of the interest rate obtained on new car loans. They report that each year of age is associated with a .05 point reduction in the interest rate, and use this to adjust the counterfactual interest rate so that a person at the mean age receives the unadjusted rate, and those above or below receive a modified rate. I construct a modified counterfactual market interest rate by following their procedure.[25]

Table 2.8 reports a series of bounding estimates that vary the discount rate and use the original measure of financing incentives, the measure that adjusts for the Prius premium and the age adjusted measure. The table reports bounding estimates from a specification with control variables, which matches the third column in table 2.6. The estimates in bold face are identical to the estimates from that table.

As expected, the bounds on the first and second tax change are affected very little by the discount rate or the interest rate measure used. The upper bound does rise somewhat when the age adjusted measure is used. This is not surprising, since the upper bound will rise whenever variance increases, and the adjustment procedure will increase variation by modifying significantly the counterfactual interest rate at very high or low ages. The statistically significant negative sign of the lower bound on the third tax change is eliminated at higher discount rates, but it is affected relatively little by the choice of interest rate measure. Thus, it is possible that the significant, negative sign is the result of the choice of discount rate. The high upper bound, however, is qualitatively unaffected.

Time Period Effects

The estimates above are based on prices observed in different time periods. By using a narrow window around each tax change, the estimates mitigate the probability that macroeconomic shocks caused Prius price movements. The tax changes, however, are each on the first day of a new quarter. Automobile manufacturers are sometimes willing to cut prices at the end of a quarter in order to meet sales targets, so it is worth investigating whether or not other models showed significant prices changes around these dates.

[25] I assign observations with missing age the unadjusted Federal Reserve interest rate.

Table 2.9 replicates the before / after comparison of column 3 from table 2.4 using several samples of non-hybrid vehicles, in order to estimate whether or not there were unusual price movements in the car market as a whole. The samples include non-hybrid Toyota sedans (6 models), the set of best-selling sedans in the same vehicle class from each nameplate (28 models, including all 6 Toyotas), and a subset who have an EPA rated combined fuel economy above 29 miles per gallon (8 models). The coefficients are *not* scaled by tax changes, so they are interpreted as dollar changes around the tax date. For example, the first column says that, on average, among all non-hybrid Toyota sedans, prices fell by $50 from late December 2005 to early January 2006.

Overall, table 2.9 gives little reason to believe that time effects significantly influence the tax estimates for the Prius. Only three of nine coefficients are statistically significantly different from zero, and all are small in magnitude. Adjusting the tax estimates for the Prius according to even the largest of these estimates would have no impact on the main conclusions of this paper.

Another way to look for time period effects is to look at each model individually. Table 2.10 shows the estimated time period effects from the same specification for every vehicle manufactured by Toyota, including vehicles sold under the Lexus and Scion brands. Several of the Lexus models sell for more than twice the amount of the Prius. Given this variance across models, table 2.10 uses estimates on log price to make coefficients more comparable. Statistically significant time effect estimates are labeled with an asterisk. There is no overall pattern in table 2.9 that would suggest a caution for the Prius estimates. While there are some statistically significant, small price movements, there is little action among sedans, and many of the price movements in any given window are of opposing sign.

In sum, the principal conclusions of the bounding estimates in table 2.6 are robust to concerns regarding the use of a single model year, assumptions made in the construction of financing incentives, and time period effects in the broader new vehicle market. The only potentially important deviation from the baseline estimates is that the alternative methodology for estimating the number of stayers present in the high tax window does raise the upper bound on the estimate for the second tax change.

Having established that consumers captured the majority of the gains of federal subsidies, at least around the first two tax changes, I move now to analysis of the incidence of state tax incentives. The analysis of state tax incentives necessarily involves a different research design, but it leads to the same conclusion. Consumers captured the vast majority of state tax incentives as well.

2.7 Estimation of the Effect of State Incentives

In several states, federal tax subsidies for hybrid vehicles were supplemented by state policies. In this section, I use a state panel research design to estimate the incidence of state tax incentives for the Prius. I find that, as in the case of the federal tax credit, consumers captured nearly all of the benefits from state tax incentives.[26]

State policies do not lend themselves to narrow analysis of transactions around tax changes because sample sizes for all but the largest states are too small. Since state laws change at different times, however, a state panel research design is possible. This design has two significant advantages that complement the federal analysis. First, the estimating equation can include general time period effects, which alleviates any potential concern about macroeconomic shocks. Second, given credible time period controls, it becomes more palatable to use observations further away from the tax

[26]Note that, if the law of one price prevails across states, then the effect of a state tax incentive might be different than a federal tax incentive. States demonstrate significant differences in prices, however, suggesting that prices are not smoothed completely across borders.

change. Thus, in the state panel regressions below, I use data from the entire life of each tax program, not just transactions immediately around a policy change.

Table 2.11 lists the states with tax incentives, along with the type of incentive, the amount and the effective dates. Twelve states have passed tax incentives that subsidize the Prius. Connecticut, the District of Columbia, Maryland, Maine, New Mexico, and New York had or have a full or partial sales tax exemption. Colorado, Louisiana, New York, Oregon, Pennsylvania, South Carolina, and West Virginia had or have a state income tax credit.[27] Colorado and West Virginia had the largest incentives, which were worth more than the maximum federal credit for some model years. Sales tax exemptions in Connecticut and the District of Columbia, as well as the credits in Oregon and New York were also worth more than a thousand dollars for most Priuses.[28]

Table 2.12 reports regressions of the form:

$$Price_{ist} = \lambda\tau_{st} + \mathbf{X_{ist}}\beta + \gamma_t + \delta_s + \mu_j + \varepsilon_{ist}$$

where i indexes an individual, s a state and t a time period, $Price$ is the incentive adjusted price, τ is the state tax incentive, $\mathbf{X}$ is a vector of controls, γ are week dummies, δ are state dummies, μ are model year cross trim level dummies, and ε is the error term. The parameter of interest is λ. The controls include dealer's cost, day of the week dummies, a dummy for the five days at the end of a month and a quadratic trend in time on the market.[29] The regressions are run on the same sample

[27] Utah had a credit for which the Prius did not qualify because it does not have a non-hybrid version.

[28] Some states also passed laws that allow hybrids to use car pool (High Occupancy Vehicle) lanes, regardless of the number of passengers. Estimating the effect of these incentives on Prius prices directly is challenging, because such a policy should affect prices as soon as it is expected to be effective and the effective dates for most states is imprecise. In most cases, a bill was in process for some time in the legislature. In several cases, after the bill became law, uncertainty remained about the policy because states required a waiver from the federal government to allow vehicle access on federal interstate highways. The states with HOV policies were not the same as the states with tax incentives, so HOV policies pose a problem only to the extent that they introduce bias into the estimation of time effects. The most concerning case is California, which is large enough to impact regression estimates. Results reported below are robust to including separate dummies for California for when the HOV policy was active and when it was not.

[29] Note that a variable indicating the value of the federal tax credit would be perfectly collinear with time period

of Priuses used above.

In these regressions, the coefficient on the tax variable is identified by states that experience a policy change in the sample period. Of the twelve states with tax incentives for the Prius, only eight experienced a change in the sample period: Connecticut, the District of Columbia, Maine, Maryland, New York, Pennsylvania, South Carolina and West Virginia. The final two columns of table 2.11 show how many Priuses are in the sample when the state policy was in effect and when it was not.[30]

Table 2.12 reports the results. The point estimate in column one says that, for a one dollar increase in a state tax incentive, the price of a Prius rises by five cents. Whether or not the cost of the Prius is included as a regressor, the point estimate is indistinguishable from zero. The upper bound on the 95% confidence interval is .21. These results are robust to the exclusion of any of the states with tax incentives, ensuring that the result is not driven by a single state with a large sample or a big tax change.

These estimates do not explicitly control for the heterogeneous shifting behavior addressed in the bounding exercise in section 2.6. Any concerns about this heterogeneity, however, are greatly mitigated by the fact that the state panel estimates do not just use transactions very close to tax changes. If heterogeneous shifting takes place only close to changes, then shifting will be too small to play a significant role in these estimates. In the case of the federal policy, table 2.3 showed that extra Priuses were sold only within two weeks of each tax change. It is quite likely, therefore, that only a small number of transactions were shifted by state incentives, where less money was at stake, relative to the total number of Priuses sold in those states.

dummies, since the federal credit only varies over time.

[30] New Mexico's policy changed in the time window, but there are no New Mexico Priuses in the sample before the change.

The point estimates in table 2.12 are small enough that a zero effect cannot be ruled out, but precise enough to rule out large price movements. Thus, these results corroborate the conclusions drawn from the federal policy: consumers capture the vast majority of tax credits for the Prius. Having established that consumers captured all, or nearly all, of both state and federal tax incentives for the Prius, I move now to an interpretation of this empirical result.

2.8 Why Did Prius Prices Not Respond to Tax Changes?

The empirical estimates in this paper show that the price paid to dealers for the Prius was not responsive to changes in tax subsidies. In this section, I offer an interpretation of this result. In the standard static, competitive tax incidence model, this result indicates that supply is quite elastic, relative to demand. I argue that this interpretation is inconsistent with other facts about the market, particularly the existence of wait lists. Moreover, in any standard model, if Toyota was capacity constrained, they should be expected to appropriate the gains from a subsidy. The empirical result, which indicates precisely the opposite, is therefore puzzling.

One possible explanation is that Toyota believed that future demand for hybrids would be diminished if they charged market clearing prices for the Prius during the period of capacity constraints. If so, then an increase in demand due to a temporary tax credit may have no effect on price, even if the seller is capacity constrained. I outline a two period model that assumes that first period prices shift the second period demand curve. I show how a tax subsidy in the first period can be captured by consumers, even if the seller faces a binding production constraint. I then discuss market characteristics that could give rise to such a demand system, including directed search and a concern for fairness. The section concludes by drawing lessons for the analysis of tax incidence in other markets.

Standard Models are Inconsistent with the Case of the Prius

In incidence analysis, the default procedure is to translate a price change into a statement about relative elasticities of supply and demand. The most basic, textbook model of incidence is a static, partial equilibrium model of a competitive market. Analyzing this case is important because it underlies the default interpretation of empirical estimates. The derivation of incidence begins with definitions of the elasticity of supply and demand and the market clearing condition. Log-linearized, this leads to the familiar system

$$\hat{Q^D} = \eta^D(\hat{p} + \hat{\tau})$$
$$\hat{Q^S} = \eta^S \hat{p}$$
$$\hat{Q^S} = \hat{Q^D},$$

where I have assumed that the consumer remits the tax and used the notation of Fullerton and Metcalf (2002): $\hat{x} = \frac{\partial x}{x}$, η is the own price elasticity, Q is quantity, p is producer price, τ is tax and S denotes supply and D demand. Substitution leads to the expression:

$$\frac{\hat{p}}{\hat{\tau}} = \frac{\eta^D}{\eta^S - \eta^D}.$$

In this simple model, a zero price effect ($\hat{p} = 0$ for $\hat{\tau} \neq 0$) can result from two extreme cases: (a) perfectly inelastic demand ($\eta^D = 0$), or (b) perfectly elastic supply ($\eta^S = \infty$).

Neither of these cases is a plausible explanation for the Prius, given additional facts about the market. First, it is unlikely that demand is extremely inelastic in such a narrow market. The subsidy in question applies only to hybrid vehicles, for which conventional vehicles are generally a close substitute. Furthermore, since Toyota is a monopoly seller of the Prius, if $\eta^D = 0$ at the chosen quantity, Toyota could raise the price of all Priuses without reducing quantity sold.

Second, the existence of wait lists makes elastic supply implausible. If supply were perfectly elastic, then a tax subsidy would cause a sharp increase in the quantity sold. Since Toyota in fact had a capacity constraint, supply must have been inelastic.[31] The existence of wait lists indicates that Toyota did not simply increase production to satisfy increased demand at a constant price. Moreover, in such a simple static model, a capacity constraint would imply a perfectly *inelastic* supply curve. This would imply that *suppliers* would capture the entire subsidy, which is precisely the opposite of the empirical facts.

Relaxing Competitive and Dynamic Assumptions Does Not Explain the Result

The incidence expression discussed above is based on a static, competitive market and a permanent tax change. This is a useful benchmark, but the assumptions involved are poor approximations of the vehicle market. Relaxing restrictive assumptions, however, will not explain all of the facts of the Prius case unless the model can explain wait lists. In general, adding market power and dynamic considerations can create additional conditions under which consumers will capture the entire benefit of a tax credit. These results, however, will not hold if there is a wait list, so that sellers can raise prices without lowering quantity sold.

The non-response of Prius transaction prices to tax changes is therefore a puzzle. If Toyota faced a capacity constraint, why were they unable to appropriate all of the gains from an increase in demand? I argue below that prices did not respond to

[31] Toyota never issued a general press release citing a supply shortage, but there was a known shortage of hybrid batteries reported in early 2005 (Truett, 2005), and media reports have quoted Toyota spokespeople claiming that the shortage of Priuses was due to production and supply constraints (e.g., Benton (2005); Bailey (2007)).

the introduction of the tax credit because Toyota had a vested interest in keeping consumer prices below the market clearing level, in order to increase demand for hybrids in future periods. Before analyzing this possibility, I provide evidence that the pricing puzzle is not due to widespread ignorance or ineligibility.

Awareness and Eligibility Do Not Explain the Result

One would not expect to observe price changes in response to an income tax credit if consumers were unaware of the credit or if they were ineligible to receive the benefits. In this section, I argue that awareness and eligibility may have had some effect on the magnitude of tax responses, but that the effect was likely to have been small.

Tax provisions are often complicated and obscure. As a result, individuals may not respond to a tax incentive because they are ignorant of it or are confused by its implications.[32] If consumers did not know about the hybrid vehicle tax credit, or misunderstood the phase-out schedule, one might expect constant transaction prices. Moreover, slow information dissemination by the Internal Revenue Service may have contributed to confusion regarding the hybrid vehicle tax credit. The tax credit took effect on January 1, 2006, but the IRS did not officially certify the Toyota Prius or announce the size of the credit until April 7, 2006 (Internal Revenue Service, 2006a). The first phase-out of the credit began on October 1, 2006, but the official IRS statement that the phase-out would begin was not issued until September 20 (Internal Revenue Service, 2006b).

In spite of the slow dissemination of official information, accurate details about

[32] Confusion may prevail even when the stakes are large. For example, Romich and Weisner (2000) and Smeeding, Ross and O'Connor (2000) find that a large percentage of Earned Income Tax Credit (EITC) recipients cannot distinguish between refunds due to over-withholding and EITC payments. And, Jappelli and Pistaferri (2007) argue that tax reforms in Italy had no discernible effect on the mortgage market because people did not understand how the changes influenced the implicit subsidy of mortgages.

the credit were available on the Internet and in other media. The American Council for an Energy-Efficient Economy released a report estimating eligibility and credit amounts in the fall of 2005 (American Council for an Energy-Efficient Economy, 2005). Other media cited this report, which provided very accurate predictions regarding credit amounts. In terms of the phase-out, Toyota sold over 42,000 hybrids in the first quarter of 2006. Thus, they were certain to reach 60,000 hybrids in the second quarter. The phase-out implications were well understood by the media and correctly reported long before the IRS pronouncement. This information was further disseminated by individuals in Internet chat rooms, which frequently included discussions of the tax credit.[33] In addition, evidence from the Internet shows that consumers sought this information. Figure 2.12 shows the relative density of searches on Google for the term “hybrid tax credit,” which reveals search spikes prior to the credit’s introduction and phase-out, and at other logical times.

All of this suggests that many consumers were informed about the credit, but it does not imply that all consumers were informed. One reason to believe that nearly all consumers were informed is that dealers had an incentive to inform them. All automakers have a system for transmitting information to dealerships about incentives and policies. Dealerships then inform salespeople during regular meetings. The Toyota dealers I spoke with confirmed that Toyota distributed information about the tax credit several weeks prior to each change, which ensures that dealers were informed. In order for dealers to potentially capture part of the tax subsidy, they would have to make sure that customers knew they would later receive a benefit on their tax returns.

Even if consumers were well informed, the tax credit may not have influenced

[33]Two relevant examples are PriusChat.com and the town hall section of Edmunds.com, http://townhall-talk.edmunds.com/direct/view/.f0a4a16.

prices if very few were eligible to receive the benefit. Recall that taxpayers subject to the Alternative Minimum Tax (AMT) and those with no tax obligation will not benefit from the credit. To determine what percentage of consumers were eligible for benefits, one would ideally like to compare take-up rates in tax return data to the total sales of hybrids. Unfortunately, the Statistics of Income (the division of the IRS that reports this type of data) has no available data on vehicle credits from 2006 tax returns at this time. Furthermore, they have no data on the tax deduction from earlier years, because the deduction was added as a miscellaneous item in a line on the form 1040, and the coding of the data did not disaggregate hybrid deductions from several other items.[34]

In the absence of data from the IRS, one might attempt to estimate the proportion of Prius buyers who were not eligible by estimating the tax situation from detailed income information, but the data to do this with precision are not available.[35] CNW Research, a marketing research firm in the auto industry, does report the mean self-reported income and home ownership rates of Prius owners. In 2006, they report the mean income of Prius buyers to be $88,750, compared to $60,750 for new car buyers across the whole vehicle market. In 2006, 77.1% of Prius buyers reported owning a home, compared to 57.9% for the whole market (CNW Research, 2007). Given these numbers and the fact that the Prius costs around $25,000, it seems unlikely that very many Prius buyers have too little tax obligation to gain from the non-refundable credit.

Despite the fact that Prius owners are relatively wealthy and relatively likely to own a home, the limitations imposed by the AMT probably did not affect many

[34] I have been told by people at the Treasury Department that no one has calculated the revenue expenditures from the clean-fuel tax deduction using actual tax return data.

[35] As mentioned above, publicly available data like the Consumer Expenditure Survey and the National Household Travel Survey are not recent enough and contain too small a sample to be useful.

Prius buyers because the AMT applied to a low percentage of tax returns. In 2005, the most recent year for which official AMT data is available, in only five states were more than 5% of returns affected by the AMT (Internal Revenue Service, 2007b). The highest rate was New Jersey, where 6.8% of all filers were affected by the AMT. Thus, even if Prius buyers were twice as likely to be influenced by the AMT as the rate in the highest state, 86% would be unaffected.

Furthermore, many AMT payers are uncertain of their AMT status before filing. In addition, it appears that many consumers affected by the AMT were surprised to find that they could not benefit from the credit. The discussions at PriusChat.com and Edmunds.com on the Prius include a number of references to the AMT. Most posters who were affected by the AMT seemed surprised that they did not receive the full benefit. Thus, the most likely form of misinformation actually works in favor of creating a price and sales response, because many ineligible consumers apparently expected to receive the credit.

In sum, there are several reasons to believe that awareness and eligibility could only have played a small role in creating a zero price response for the Prius. The data indicated that many consumers understood the policy, and it is reasonable to believe that nearly all consumers would have been informed by dealers. The vast majority of Prius buyers were likely eligible for the credit, and the most likely type of confusion, regarding eligibility and the AMT, probably made many ineligible consumers believe that they were eligible, which would have amplified the price response. Since awareness and eligibility do not explain the price result, the puzzle remains. I move now to a stylized model that describes conditions that could give rise to a zero price response.

An Explanation, and a Stylized Model that Encapsulates It

Standard models cannot explain all of the facts of the Prius market unless they can explain the presence of wait lists. The conventional explanation among industry watchers for wait lists is the following. Toyota was surprised by the strength of demand for the second generation Prius, which led to shortages beginning in 2003. They did not want prices to rise sufficiently to clear the market because they had a long-term interest in establishing hybrids as an affordable, cost efficient technology in the eyes of consumers. As the symbol of hybrid technology, they did not want the Prius to be perceived as a vehicle available only to the wealthy. In addition, they wanted the Prius to be perceived as a car comparable in cost to midsize sedans like the Camry and Accord, because they believed that, over time, that is where they would price the Prius. In other words, Toyota believed that if prices rose to clear the market during the period of excess demand, then demand for hybrids in the future would be lower.

Toyota could not meet the higher than expected demand by increasing production because they faced capacity constraints. In the automobile industry, there are very large fixed costs to production, and there are potential bottlenecks in parts suppliers, especially if a vehicle needs an uncommon part like a hybrid battery which is produced by only one supplier. Prior to the 2007 model year, all Priuses were assembled at one plant, the Tsutsumi Plant in Toyota City, Japan. When Toyota ramped up production significantly for 2007, it required the opening of a new assembly line in a different plant. This suggests that there were indeed large capital investments that could not be quickly changed to accommodate greater than forecasted demand.

Toyota is heavily invested in hybrid technologies. Toyota sold its one millionth hybrid in June 2007, and it announced that it expects to sell one million a year by

2010, making hybrids 10% of Toyota's global sales (Healey, 2007). Hybrids already represent 10% of Toyota sales in the United States, as of 2007. In 2003, Toyota pledged to introduce a hybrid version of all of their vehicles. They have backed away from this recently, but still plan to double the number of models available and make a hybrid version available in every market segment within a few years (Healey, 2007). Toyota currently holds over 1,000 patents on hybrid related technologies, and Nissan, Ford and Fuji Heavy Industries have agreements to use Toyota's techniques (Rowley, 2006). This suggests that, if Toyota did believe that higher Prius prices in 2005 would jeopardize hybrid demand in later years, then they had a lot at stake.

Model

To understand how such a concern could influence tax incidence, I develop a simple two period example in which a monopolistic seller faces a first period capacity constraint and first period price is assumed to influence second period demand. I show that a first period tax credit may have no influence on the price. This example highlights several key issues, including the distinction between tax exclusive and inclusive prices, which I return to in the discussion.

For ease of exposition, I assume a linear functional form for demand. I posit a multiplicative relationship between first period prices and second period demand. The benefit of analyzing this case is that closed-form optimal values take very simple functional forms, and comparative statics are then immediately transparent. Let D denote demand, p price, τ the tax subsidy, and denote second period variables with a prime ($'$). Then

$$D = \alpha(A - (p - \tau))$$

$$D' = (\alpha' - .5\beta p^2)(A - (p' - \tau')),$$

where A, A', α, and α' are positive scalars. We can interpret A, A', α, and α' as shifters of the size of the market in each period. β is a scalar that describes how first period price influences second period demand. If β is zero, then the model collapses to a standard case. If β is positive, then high first period prices shrink future demand. The setup assumes that consumers remit the tax and that second period demand is influenced by the subsidy exclusive price. The latter assumption is substantive, as will be noted below. As a regularity condition, assume that $\alpha' - .5\beta p^2 > 0$ for the optimal p.

The firm produces Q units at constant marginal cost c in both periods, but it faces a capacity constraint in the first period, at which time it can produce only up to N units.[36] Thus, it maximizes profits (π) by solving the following problem:[37]

$$\max_{p,p'} \quad \pi = (p-c)Q + (p'-c)(\alpha' - .5\beta p^2)(A - (p' - \tau'))$$

$$\text{where} \quad Q = \begin{cases} \alpha(A-(p-\tau)) & \text{if} \quad \alpha(A-(p-\tau)) \le N \\ N & \text{if} \quad \alpha(A-(p-\tau)) > N. \end{cases}$$

Consider first the case where the capacity constraint is not binding, so that $Q = D = \alpha(A - (p - \tau))$. The first-order conditions produce a system of two equations. Substitution yields the optimal values — labeled $*$ for optimality and u for unconstrained:

$$p_u^* = \frac{A + \tau + c}{2 + \frac{\beta}{\alpha}\left(\frac{A' + \tau' - c}{2}\right)^2}$$
$$p_u'^* = \frac{A' + \tau' + c}{2}.$$

If $\beta = 0$, which means that first period price has no effect on second period demand, as would be true in a standard model, then $p_u^* = \frac{A+\tau+c}{2}$, which is the monopolist's

[36] The capacity constraint is exogenous in this example. It can be thought of as arising from a choice made in a prior period, in which the producer must pay fixed costs to build capacity but is uncertain of demand.

[37] A discount factor could be added without changing the results, or, equivalently, α' and β can be interpreted as including the discount factor.

price in a static one period model and is symmetric to the second period price. If β is positive, then the second term in the denominator is also positive. As one would expect, the larger is the second period effect on demand (β), the lower is the optimal first period price. Note that, because the scale effect on the market is multiplicative, the optimal second period price is unchanged by first period prices.

In the constrained case, the monopolist's optimal price is below the price that creates N units of demand. Because the monopolist is capacity constrained, it will sell N units at the optimal price, but the market will not clear. There will be excess demand. In this case, the optimal prices may be written as follows, with N denoting the capacity constrained case:

$$\begin{aligned} p_N^* &= \frac{N}{\beta\left(\frac{A'+\tau'-c}{2}\right)^2} \\ p_N'^* &= \frac{A'+\tau'+c}{2}. \end{aligned}$$

As expected, given the independence of second period price from scale effects, the capacity constraint has no effect on second period prices, so that $p_N'^* = p_u'^*$. The first period price changes, however, so that all first period terms no longer influence first period price. The parameters of the first period demand curve are no longer relevant because the seller has chosen a price / quantity bundle that is not on the demand curve.

The main point of this exercise is that, in the capacity constrained case, a shift in the first period demand curve caused by a tax credit does not have an effect on the price charged by a monopolist, unless it causes the monopolist to shift from the constrained quantity to an unconstrained one. The relevant comparative statics from this simple example, assuming that $\beta > 0$, are:

$$\frac{1}{2} \geq \frac{\partial p_u^*}{\partial \tau} > 0 \qquad \frac{\partial p_N^*}{\partial \tau} = 0$$
$$\frac{\partial p_u'^*}{\partial \tau'} = \frac{1}{2} \qquad \frac{\partial p_N'^*}{\partial \tau'} = \frac{1}{2}$$
$$\frac{\partial p_u^*}{\partial \tau'} < 0 \qquad \frac{\partial p_N^*}{\partial \tau'} < 0$$
$$\frac{\partial^2 p_u^*}{\partial \tau \partial \beta} < 0 \qquad \frac{\partial^2 p_u^*}{\partial \tau \partial A'} < 0.$$

The first period tax subsidy has a positive effect on first period price whenever the producer is unconstrained, but it has no effect on price in the constrained case. Second period prices rise 50 cents on the the dollar in response to second period tax subsidy changes. Note that this is identical to the comparative static that would arise in the first period if $\beta = 0$, so that there was no price effect on second period demand. Thus, the effect of first period price on second period demand mutes the tax response in the first period. We see also that an increase in the second period subsidy will *decrease* first period price. The reason is that an increase in the second period subsidy represents a shift in the second period demand curve that makes that market more profitable, and the value of expanding that market through low first period prices expands. Finally, note that the first period tax effect in the unconstrained case gets smaller as β and A' rise, because a larger β indicates a bigger effect on second period market size and A' indicates a more profitable second period market. Moreover, as $A' \to \infty$, $\frac{\partial p_u^*}{\partial \tau} \to 0$; i.e., as the second period market gets sufficiently large, the first period tax effect will shrink to zero.

This modeling exercise suggests a set of conditions under which a temporary tax credit for the Toyota Prius might not influence the current price charged by Toyota. This would be true if (a) Toyota was capacity constrained when the subsidy was available and (b) by raising price while the tax credit was available, Toyota would hurt the market for their product in the future. Even if they were not capacity constrained, if the size of the future market is sufficiently large, the tax effect in the first period may approach zero.

The key assumption made here is that first period, *subsidy exclusive* prices have an influence on demand in the second period. If second period demand is influenced by the tax inclusive price, then the tax pass through would be different. In the constrained case, the pass through would be 1, not 0. The intuition is simple; in the constrained case, the firm has chosen the optimal scaling factor for the second period market, and it wants to keep this same optimal factor after a change in the subsidy. To do so, it must completely offset the subsidy increase. Thus, which price influences future demand is an important distinction to keep in mind for the discussion of what behavior might give rise to such a stylized model. I move now to a discussion of factors that could give rise to the demand system analyzed in this example.

Directed Search May Generate Such a Result

In this and the next subsection, I describe two different approaches to explaining why current subsidy exclusive prices might influence future demand. In this section, I describe how costly search can create this link. In the next section, I discuss the economic literature on fairness. The description in this section can be thought of as rational in the neoclassical sense, whereas the next section relies on what might be called behavioral considerations.

Consider the following description of the car buying process. Consumers do not know the final transaction price of a vehicle or how much they like a particular model until they invest time researching, test driving, obtaining a price quote from a dealer, or even bargaining over price. Consumers do not know the exact price or how much they will like a particular model before incurring these time costs, but they have beliefs about how these outcomes are distributed. Given these beliefs, they choose a model to research, test drive and price and thus receive a draw from the distribution. They can either accept this draw and buy the vehicle or obtain another draw by searching again — over the same model or a different model. At some point, the consumer will choose to end the search process, and will purchase the vehicle that represents the best option from the set of draws.

The model of Weitzman (1979) describes a search problem with all of these features.[38] In his model, an agent faces a set of alternatives, which each have a unique search cost and a distribution of possible outcomes. After searching among several alternatives, the agent chooses the best option. The optimal search algorithm describes the order in which the options are searched and the optimal stopping rule. Weitzman (1979) shows that all of the alternatives can be described by their "reservation value", which is a relatively simple function of the payoff distribution and costs. Agents search the option with the highest reservation value, and they stop when a realized value of one of the searched options exceeds the reservation value of all remaining options.[39]

Consumer *beliefs* about the price distribution of a model, therefore, may influence final demand for that model because it influences the model's reservation value. Higher expected prices will lower the reservation value of a vehicle, which will move it further down the search queue. Since consumers who achieve a good realization from an early search will stop searching and purchase a vehicle, fewer consumers will search over a vehicle when it is moved down the queue. As a result, demand will fall.

[38] I am not the first to suggest that the Weitzman model is a good model of the car buying process, see Moorthy, Ratchford and Talukdar (1997).

[39] Weitzman uses the term "reservation price," which I have modified to "reservation value" to avoid confusing it with transaction prices. One of the key points of this model is that reservation values are not the same as expected values, even with risk neutral agents. An option that has a higher expected value than another may nevertheless have a lower reservation value.

Note that, even if the actual distribution of prices is constant, a change in consumer beliefs will shift the demand curve inward.

Such a characterization of the car buying process only makes sense if final transaction prices are not costlessly observable. Final transaction prices are not immediately observable in the new vehicle market. List prices are obtainable at minimal cost, but final transaction prices vary significantly across geography, over time and according to idiosyncratic factors, like trade-ins and credit ratings.

Surveys indicate that consumers, on average, consider only about three models before making a new vehicle purchase (Ratchford and Srinivasan, 1993; Moorthy et al., 1997; Ratchford, Lee and Talukdar, 2003). As a result, search order will have a very large impact on demand, which in turn implies that automobile manufacturers should be very concerned with factors that determine consumer's "shopping lists," including past transaction prices.

Past transaction prices will influence future demand for a particular model if past transaction prices are known to consumers and used to form their beliefs about the distribution of prices. Past transaction prices are collected by consumer services agencies, like Edmunds.com and Consumer Reports, and reported to consumers. Past transaction prices are also posted on Internet websites, and they are transmitted via word of mouth within social networks. Thus, it is likely that past prices are used in forming beliefs about future prices.

In terms of the stylized model, this implies that $\beta > 0$, so that higher current prices reduce the number of consumers who search a particular model, and thus reduces future demand. When a new product, like a hybrid vehicle, is introduced on the market, consumers will be uncertain of its price. If early transaction prices are very high, consumers may expect future prices to be high, and this will shift inward the future demand curve.[40] Thus, an automaker who makes a capital investment that limits the production of a model and then finds that demand is higher than expected faces a choice. If they plan to make a bigger capital investment in the future, so that future sales will rise and the market clearing price will fall, should they let prices rise to clear the market in the short run? If so, demand for the vehicle in the future may fall, if consumers interpret high prices as indicative of future prices.

Might consumers be able to infer that high early prices are an overestimate of future prices? If consumers knew that high early prices were the result of a capacity constraint, or if they understood the true cost of hybrid components, then they might not interpret high early prices as a signal of high future prices in equilibrium. Consumers are unlikely to be informed about production constraints and technology costs in the automobile industry, because much of this information is considered proprietary and is held secret. None of the industry analysts I spoke with claimed to know with any precision how much hybrid components cost Toyota. It seems unlikely, therefore, that consumers are able to estimate future hybrid costs based on fundamentals. Instead, they probably use current and past prices to form beliefs about future hybrid prices.

Tax Equivalence: Which Price Influences Future Demand?

Directed search explains the tax incidence of the Prius only if consumer beliefs are influenced by *subsidy exclusive* past prices. In the stylized model, this means that second period demand was shifted by p, not $p - \tau$. Evidence presented above suggests that hybrid buyers knew of the tax credit and understood its implications. This does not, however, imply that potential future buyers know anything about it. Even if those who purchased while the credit was in effect were perfectly well

[40]Note that high prices could also be interpreted in some cases as a signal of quality. I abstract from such concerns in the present discussion, but this could be built into a more formal model.

informed, the price transmitted to future buyers may be the subsidy exclusive price. Prices posted on internet chat rooms, Edmunds.com, Consumer Reports and other such sites are almost always tax exclusive prices. Individuals sometimes report "out the door" prices which include taxes and fees, but, since sales taxes vary across localities, reporting agencies typically report tax exclusive prices. In the case of the hybrid tax credit, these agencies did not subtract out the credit when reporting transaction prices.

Note also that, if Toyota is concerned with how the subsidy exclusive price influences future demand, then a subsidy paid directly to Toyota might have had a different incidence. This is in contrast to standard models, but it is a natural consequence of the stylized example, given the assumption that the relevant signal is tax exclusive.

If consumers do know about the tax credit and attempt to adjust their forecast of future prices accordingly, they will have to have a method for distinguishing between a higher price due to taxes and a change to Toyota's long term projected costs. The costs of a new technology are uncertain, even to the manufacturer. Automakers are continually updating their beliefs about the long run costs and viability of varying hybrid systems. If Toyota adjusted upward their long run cost projections, this may look just like an increase due to a temporary tax credit to consumers. Given this ambiguity, it may be profit-maximizing to ration quantities in early periods. A temporary increase in demand due to a tax credit may therefore have no effect on price because the manufacturer wishes to signal a particular future price to the market.

The Role of Dealerships

Once Toyota sells a vehicle to a dealer, the dealer owns it, and Toyota is prohibited from dictating the retail price. Even if Toyota wanted to keep prices below market clearing levels, why would dealers not raise prices? The most obvious answer is that, while Toyota is legally prohibited from setting retail prices, they nevertheless have enormous leverage over dealerships and can punish dealerships who do not comply with their sales strategy. According to Toyota dealers that I spoke with, Priuses were allocated based on how fast they were sold by dealerships. Thus, if a dealership raised the price of a Prius, so that it cleared the local market, eliminated the wait list, and Priuses took longer to sell, then that dealership would see the arrival rate of future Priuses decline. This would have been quite costly, since the Prius, even if it sold at sticker price with no markup, was grossing far more than any similarly priced vehicle. Thus, even if dealers had no independent reason to hold prices below market clearing levels, Toyota was probably able to incentivize them to comply with their pricing scheme.

Concerns for Fairness May Generate Such a Result

A starkly different interpretation of the stylized model is based on behavioral considerations, including a concern for fairness. If consumers are willing to punish firms that are perceived as unfair, then prices may not respond to a tax subsidy if such a price change is deemed unfair. While this is unnatural in a neoclassical model, there is a substantial body of research that documents concerns for fairness and a willingness to punish. In terms of the example above, a firm violating fairness in the first period will shift the demand curve downward in the second period (i.e., $\beta > 0$).

Behavioral considerations may also explain the case of the Prius if consumers are subject to framing biases, in the spirit of Kahneman and Tversky (1979). In the car market, vehicles are commonly classified into segments (sedans, compact cars, midsize SUVs, etc.) and price levels (economy, luxury, etc.). Vehicles are frequently compared to other vehicles in their "class" in terms of their reliability and performance. If framing matters, consumers may rate a car differently depending on the quality of other vehicles in its class.

The Prius is typically classified as a compact sedan, but it is the most expensive compact sedan, and is more expensive than many midsize sedans. If prices had risen during the wait list period, the Prius would have approached the price of some *luxury* sedans, which tend to have significantly better performance and features. Consumers may have deemed the Prius to be of inferior quality if they compared it to these luxury sedans, and this reputation may have influenced future demand. If so, then Toyota may have wished to hold prices below market clearing levels to maintain a comparison with compact sedans. An increase in demand due to the tax credit might therefore not have influenced price.

Behavioral concerns may also explain why *dealers* did not raise prices to clear the market. Kahneman, Knetsch and Thaler (1986) asked consumers whether or not it was fair for an automobile retailer to raise prices by $200 in response to a shortage for a popular model. In one version, respondents were told that the dealer had been selling the car at the list price; the increase would therefore require charging more than the sticker price. In a second version, respondents were told that the dealer had been selling the car at $200 below list price; the increase would therefore lead to a new price equal to the list price. In the first case, 71% of respondents said the price increase was unfair, but only 42% thought the price increase was unfair in the second case. (The difference is statistically significant.)

This is direct evidence that the sticker price is a salient reference point. Dealers may have felt that charging above the sticker price would be perceived as unfair. Car dealers may be especially concerned about such reputations because many consumers distrust them. If dealerships experience a discrete reputation cost when charging above the sticker price, a marginal increase in demand due to a temporary tax credit may induce no transaction price response. Note also that this means that the sticker price is a way for Toyota to influence retail prices, if the sticker price acts as a ceiling on the retail price.

More generally, consumers appear to rate the fairness of a price change differently, depending on the perceived cause of the change. Kahneman et al. (1986) ask several questions that demonstrate the difference in perceived fairness between price increases due to temporary shortages versus changes in cost. In one example, they ask respondents whether or not it is fair for a hardware store to increase the price of snow shovels the day after a snowstorm; 82% say that this is unfair. Frey and Pommerehne (1993) follow a similar method and find that 80% of respondents believe that price increases made to clear markets with short-run excess demand are unfair.

Toyota may be especially concerned about fairness perceptions in regards to the Prius because the Prius has generated public goodwill for Toyota. For many, the Prius is a metonym for Toyota, while the Hummer is a metonym for General Motors. For those concerned with the automobile industry's role in carbon emissions, this reflects positively on Toyota and negatively on GM. It may be partly because of this gap in perception that there was no outcry when Toyota passed GM in global sales in the first quarter of 2007. This marked the first time in 80 years that GM was surpassed, but there appeared to be more celebration of Toyota's leadership in fuel economy than nostalgia for the American company's age of dominance in the public discourse. In other circumstances, it is easy to imagine Toyota's ascension as a spark for policy intervention to aid the domestic industry. This is not directly a reason why Toyota would not charge more for a Prius, but it is a reason why Toyota has a significant investment in the goodwill afforded by the Prius. If sharp price changes in response to taxes jeopardized that goodwill, Toyota might be better served to forgo the revenue from a higher price.

Tax Equivalence and Fairness

A concern for fairness may also break tax equivalence, if consumers view a subsidy rebated to them differently than a subsidy rebated to a seller. This line of reasoning suggests an interpretation of the results in Busse et al. (2006), who find that consumers capture roughly all of a customer cash rebate, but less than half of a dealer cash incentive. They attribute this result to information asymmetry, but it is also consistent with consumers feeling some "ownership" over a rebate given directly to them. If consumers feel entitled to a cash rebate or tax credit "given to them," then incidence may depend on which side of the market remits the tax.

Willingness to Punish is Key

We should expect firms to respond to fairness criteria only if future demand responds to violations of fairness norms. Some evidence regarding consumer willingness to punish firms that offend fairness norms comes from similar survey questions (Kahneman et al., 1986; Campbell, 1999). A willingness to punish has also been shown in experiments (see Fehr and Gächter (2000) for a discussion of the evidence).

Another piece of evidence regarding the value of fairness perceptions to a firm comes from a recent episode in the mobile phone market. In July 2007, Apple introduced the iPhone, a high priced mobile phone. The 8-gigabyte iPhone debuted at a price of $599, and people queued for days to buy one on the release date. Only two months later, in September, Apple dropped the retail price by a third to $399. Consumers who had purchased the identical phone in the first two months were outraged. The public relations pressure was sufficient to move Apple CEO Steve Jobs to post an open letter apologizing to customers and to provide all early purchasers with a $100 store credit. This rebate was essentially a $75 million check written to consumers to recover goodwill.[41]

In sum, a link between current prices and future demand can explain the tax incidence observed in the data. This link could be forged by imperfect information, which forces consumers to use past prices as a signal of future prices. Or, it could be created by behavioral considerations, including fairness perceptions. I next present additional evidence from the data that links the incidence result to the wait lists, and then conclude with lessons for future research on tax incidence.

Evidence in Support of the Importance of Wait Lists

The modeling exercise suggests that incidence may depend on whether or not a product is in excess demand. Several pieces of evidence from the hybrid market are consistent with this. First, in the case of the Prius, the difference between the first two and the final tax change is consistent with incidence differing in the presence of wait lists. Second, Prius prices were more sensitive to other price determinants when there was not a wait list, which supports the notion that prices were inflexible during the 2004 to 2006 model years. Third, evidence from other hybrid models confirms the case of the Prius — those vehicles that had wait lists look similar to the Prius in their response to taxation.

The bounding exercise produced a much tighter bound on the first two tax events

[41] In early September, Apple claimed to have sold its millionth iPhone. As a conservative estimate, if Apple had to rebate $100 on only 750,000 units, that would work out to be $75 million.

than the second (see table 2.6). Thus, it may be the case that the pass through effect of the final tax change, at which time there was no wait list, was larger than the first two events. Toyota's introduction of the economic savings bonus program is also suggestive of a response to taxation, even though the price influence of this program was offset by the change in low interest financing. Recall again the argument that consumers may have perceived a price cut that offset the tax change, even if this was not true, as suggested by (Busse et al., 2007).[42]

In addition, the price of the Prius was more sensitive to other factors during the wait list period. Table 2.13 shows regression results from a model of Prius prices in which the end of the month dummy, the price of gasoline and weeks that the model has been available are interacted with a wait list dummy. For this specification, I assume that all 2004, 2005 and 2006 model year Priuses were on a wait list, and 2002 and 2007 models were not. All three interactions are statistically significant, demonstrating that the Prius was more sensitive to these factors when there was no wait list in place. This corroborates the notion that the price of the Prius was inflexible during wait list periods, which explains why tax changes would not influence prices.

Finally, bounds estimated for the price effect of tax changes on other hybrids are broadly consistent with zero effects for vehicles on wait lists, and not for other models. Only two non-Toyota models were present in reasonable numbers before the tax policy was initially introduced. These are the Ford Escape and the Honda Civic. The Ford Escape did not sell at an especially fast turnover at any point in its history. The Honda Civic, however, looks much like the Prius at the end of 2005 and beginning of 2006, when the tax credit was introduced. The 2006 Honda Civic

[42]Note that the estimates from the state tax changes are almost entirely based on the wait list period. Thus, the state variation is not helpful in directly comparing the wait list period to the non-wait list period, but it does confirm that the low pass through result is due to variation in the wait list period.

debuted a fully redesigned body style, and both the conventional and hybrid versions of the Civic sold at a very rapid rate, making them look much like the Prius during the tax change.

There are two other Toyota vehicles with sufficient available data, the Highlander and the Camry. The Highlander Hybrid has sold like a normal vehicle for its entire history. The Camry was introduced shortly before the first tax phase-out, and was in excess demand at that time. It was still selling at a very fast rate, and was unavailable on lots in many locations, at the second phase-out.

Table 2.14 reports the incidence bounds for these other vehicles using the same specification as column three in table 2.6, which is reported again for comparison. The Escape and Highlander, which did not experience excess demand, show higher upper bounds. In two out of three cases for the Highlander and in the case of the Escape the lower bound estimate is also positive, unlike the Prius. These estimates are, however, far less precise than the Prius estimates. The Civic and Camry, which were in excess demand at the tax changes, both have negative lower bounds, with varying tightness on the upper bound, just like the Prius.

Overall, this evidence indicates that Prius prices were inflexible during the wait list period, and that this inflexibility led to no change in transaction prices in the face of tax changes. Thus, as an empirical matter, the incidence of a tax may be quite different when a product is in excess demand. This is consistent with the stylized example described above. Several other lessons for the economics of tax incidence are suggested in the next section.

Lessons for the Economics of Tax Incidence

The model above demonstrates that, when current prices influence future demand, a model of tax incidence may generate predictions that deviate significantly from standard models. In general, if future demand is inversely related to current tax exclusive prices, then consumers will bear more of the burden of a temporary tax (or gain more of the benefit of a subsidy) than would otherwise be the case.

In addition, the incidence of a tax introduced when a product is experiencing a shortage may be different than the same tax placed on the same product at a different point in time. Temporary and permanent taxes may also be expected to have different effects in the presence of a binding capacity constraint.

Furthermore, if the current *tax exclusive* price is the price that determines future demand, then the model suggests that a subsidy to consumers may have a different incidence than a subsidy to producers. The economics of taxation has recently shown renewed interest in cases under which various tax equivalence theorems fail to hold. Slemrod (2007) analyzes a partial equilibrium model that characterizes conditions under which tax incidence varies, according to who remits a tax. This is closely related to the labor supply model of Slemrod (2001). The driving force behind incidence asymmetry in these models is tax avoidance, compliance and enforcement, which depend on the system of remittances. This line of reasoning is pursued in both Slemrod (2007) and Kopczuk and Slemrod (2006) as an explanation for the real world discrepancies between theoretically equivalent retail sales taxes and value-added taxes. A recent paper by Chetty, Looney and Kroft (2007) explores a different failure of tax equivalence. In a field experiment, they find that demand falls for groceries that are posted with tax inclusive prices. They interpret this result in a bounded rationality framework, in which consumers incur cognitive costs when calculating tax-inclusive prices. In the car market, Busse et al. (2006) argue that information asymmetry generates differences in the incidence of manufacturer incentives — which act like a subsidy — depending on whether incentives are rebated to consumers or dealers. The case of the Prius highlights conditions under which remittance matters, not because of avoidance, bounded rationality or information asymmetry, but instead because tax exclusive prices influence future demand.

The case of the Prius highlights several conditions that increase the likelihood that current prices will influence future demand. First, the Prius was a new product and the manufacturer was uncertain about consumer demand. This led to a conservative capital investment, which led to a binding capacity constraint. It also meant that consumers were unsure of the price. Second, the Prius generated goodwill among consumers towards Toyota. Toyota may have been especially reluctant to raise prices temporarily, if it might jeopardize goodwill. Third, it is costly to learn about prices in the car market. As a result, past prices, which are reported on the Internet and in consumer buying guides, hold considerable value in informing car buyers on how to direct their search for a new vehicle. Future research that determines which of these conditions is pivotal, and that distinguishes between "rational" and "behavioral" explanations, could shed considerable light on how taxes affect behavior in a variety of markets.

Many "green" products are likely to feature these conditions. Many firms now feature products that are environmentally friendly for marketing and public relations reasons. And, many "green" products are new products, for which the overall market demand is uncertain. As government agencies seek to encourage environmentally friendly practices through the tax system, the potential for incidence asymmetry and the incidence results for the Prius should be kept in mind.

Immediate examples of similar products are other advanced technology vehicles, such as plug-in hybrids, hydrogen fuel cell vehicles, clean diesels and solar or electric powered cars. Many of these advanced technology vehicles are already subsidized by state and federal policies, and they are likely to receive further subsidies in the future. A reasonable expectation for any of these products, if they become popular, is that early adopters would get a windfall gain from government subsidies. For example, a new line of electric sports cars, made by a firm called Tesla Motors, will go on sale during the 2008 model year. The first year's entire production line has already been sold ahead of time. The vehicles will likely qualify for preexisting tax incentives, but this will not affect the price paid by consumers to Tesla.

2.9 Conclusions and Directions for Future Research

This paper uses transaction level data on new vehicle purchases to assemble several pieces of evidence which indicate that consumers captured the significant majority of the benefits from tax subsidies for the Toyota Prius. The federal tax credit for hybrids, which was introduced by the Energy Policy Act of 2005, created three sharp changes in the value of federal tax subsidies. Incidence estimates based on comparing transaction prices just before and just after each tax change show that subsidy exclusive transaction prices moved very little, if at all, which implies that consumers captured the bulk of the subsidy.

Many transactions were shifted across time in order to maximize tax benefits. This implies that consumers were aware of the tax policy, and it raises the possibility that consumers who purchase just before or just after a tax change differ systematically from each other, because some choose to move their transaction to gain the higher tax subsidy but others do not. The paper develops a method of bounding the effect of the tax in the presence of this type of heterogeneity, which may be useful to incidence analysis whenever tax changes are anticipated. This methodology yields informative upper bounds on the first two tax changes and verifies the conclusion that consumers captured the majority of the subsidy.

An analysis of the incidence of state tax incentives corroborates this result. The paper uses a state panel research design and finds that consumers capture all, or nearly all, of state tax benefits. The empirical finding that consumers capture the bulk of state and federal tax benefits has important implications for policy and for future research on tax incidence.

The federal government spent approximately $400 million on the tax credit for the Prius. The incidence estimate in this paper implies that this money went directly to Prius owners, not to Toyota or to dealers. A full evaluation of the policy would include not only an incidence estimate, but also an estimate of the effect of the credit on the total number of Priuses sold. A full analysis of the latter question requires a structural approach, but institutional facts give reason to doubt that the production response was significant. Given that there were wait lists prior to the 2006 model year, Toyota probably did not produce any "extra" Priuses as a result of the policy, since they appear to have been capacity constrained before the law was passed. Toyota made a major capital investment prior to the 2007 model year, but this investment was made with the knowledge that the tax credit was already scheduled to phase-out. Thus, the credit probably had little impact on the capacity expansion.

As mentioned in the introduction, Gallagher and Muehlegger (2007) find that state sales volumes rose in response to tax credits and other incentives. One cannot readily interpret their estimates as a *production* response, however, since their reduced form methodology does not distinguish total sales effects from timing effects or geographic reallocation. Toyota's response to a state tax policy may be to divert Priuses that would have been sent to a different state into the state with the incentive. Future research that estimates a production response to the federal credit or untangles geographic shifts in sales from the overall production response would be an important complement to the price analysis performed in this paper.

The finding that consumers captured a significant majority of tax incentives is important beyond the scope of immediate policy analysis because the result is difficult to explain in standard models of tax incidence. The paper argues that Prius buyers captured the benefits of tax incentives because Toyota believe that raising prices would jeopardize future demand for hybrids. The paper provides a theoretical example in which future demand depends on current prices, which might make a seller with a capacity constraint willing to hold prices below the market clearing level, in order to increase future demand.

In such a model, the standard result that who remits a tax does not influence incidence may not hold. There is growing interest in public economics regarding such deviations from traditional tax theory, which may be guided by the case of the Prius. The characteristics of the Prius market that likely contributed to this situation were that it was a new product, that learning new vehicle prices involves costly search, and that consumer good-will toward the company was tied to the product. Many "green" products are new and are used by firms to generate goodwill for a brand, and many such products are likely to be subsidized by governments aiming to combat rising environmental and security concerns raised by the consumption of oil. Such policies should be designed with potential incidence asymmetry in mind. In addition, future research should consider such policies a fruitful area for finding deviations from traditional tax theory which may deepen our understanding of how non-standard considerations determine how taxes affect real world behavior.

2.10 Appendix

This appendix provides details about the measure of dealer cost used throughout the paper, analyzes the stability of this measure and documents the adjustment made to account for the price change instituted by Toyota in April 2007. The results presented in the paper above are sensitive to whether or not dealer cost is controlled. Thus, it is important to determine whether or not dealer cost varies over time because Toyota charges dealers a different amount for identical cars, or because the percentage of cars with more expensive options fluctuates. Toyota also made a change to invoice and sticker prices in April 2007, which it called an “economic savings bonus.” The base price of the vehicle did not change, but the invoice and sticker price of each options package did. Thus, to compare identical vehicles before and after the tax change, it is necessary to adjust dealer cost for cars sold in April, which requires knowing which options package each vehicle had. Since their is no indicator for options packages in the data, I develop a method for assigning options packages to each car by examining the dealer cost data.

In the case of the Prius, the vast majority of the variation in dealer cost comes from the factory-installed options package. The 2006 Prius had eight options packages, and the 2007 Prius had six. Each vehicle is shipped from the factory with one of these options packages installed, adding up to $6,000 to dealer cost. If the data included an indicator for each options package, one could simply see if Priuses with any particular options package cost the same amount before and after each tax change. Even without this, if the dealer cost variable in the data was influenced only by the options package, then the dealer cost measure could easily be matched to outside information on the invoice price of each package. There are other factors that influence the dealer cost, however, which introduce noise into this procedure.

Dealer invoice of the base model and each options package is readily available from online sources, which acquire this information and provide it to consumers for purposes of price negotiation. The challenge is to match this precisely to the measure of dealer cost in the data. The measure of dealer cost available in the data includes factory installed options, dealer installed "hard add" options, advertising fees and the delivery charge. It does not include what is known as the holdback.

Factory-installed options include not only the options package, but also small additional features, like floor mats, a cargo net or a satellite radio. Thus, two Priuses with the same options package may differ in price because one includes a cargo net, and the other does not. Dealer cost will also reflect "hard adds," features that add to the resale value of the car, such as a roof rack or upgraded tires. It is impossible to determine from cost data alone which of these features is present in a vehicle, and these factors therefore introduce noise into the process.

Recall from section 2.2 that dealer invoice prices for identical items do not vary across localities, in accordance with franchise law, but geographic variation does arise in small measure through fees. Toyota adds $55 to the delivery fee for Alabama, Georgia, Florida, North Carolina and South Carolina. Since the delivery charge is a known function of state, I can easily account for this in the calculations. There may also be some geographic variation in document preparation fees, however, which will introduce a small amount of noise. More substantial variation may come through advertising fees, which are typically between 1% and 3% of the base MSRP for a vehicle, depending on the cost of advertising in the local market. There is no reliable source of information that indicates the variation across geography in advertising fees. My accumulation of anecdotal evidence, however, suggests that most Priuses were charged very close to 2% of base MSRP.

The final important discrepancy between dealer invoice and dealer cost is what is known as the holdback. The holdback is a rebate that the dealer receives from the manufacturer after a car is sold. In the case of Toyota, the dealer holdback is known to be 2% of the base MSRP. Thus, dealer invoice prices quoted in consumer resource guides are 2% higher than the actual final cost to the dealer. Since the measure of cost in the data nets out holdback, it is necessary to subtract holdback from the invoice price to make them comparable. Since the cost measure includes advertising fees, however, which are estimated to also be roughly 2% of base MSRP, I let these two omissions cancel each other out. This will be correct on average, but it will induce variation, to the extent that advertising fees are above or below 2% in a particular location.

Thus, several considerations create small amounts of variation in the dealer cost measure. This variation makes it impossible to match exactly each vehicle to an appropriate invoice price and thereby infer the options package. Fortunately, the noise appears to be small enough that the imputation is meaningful. In figure 2.13, I show a kernel smoothed density of the dealer cost of all 2007 Priuses in the sample between September 2006 and March 2007. The dealer cost variable is adjusted for the $55 surcharge on deliveries in the relevant states. The vertical bars are at invoice prices, estimated from Edmunds.com data as follows: I sum the base price, options package price and delivery fee for each vehicle. Based on my reading of numerous consumer accounts, I add the dealer invoice price of floor mats to all but the least expensive options package, since all consumers report floor mats being included on all Priuses with options package 2 or greater in these years. Overall, figure 2.13 shows remarkable consistency between high points in the density, and the invoice prices. If advertising fees are on average 2% (and thereby offsetting the holdback on average), we would expect modal prices to be somewhat to the right of each line, since small additional options would add weight to the distribution asymmetrically.

Now, I move specifically to a discussion of how cost may vary around each tax change. Recall from section 2.2 that the dealer's invoice price of a vehicle rarely changes during the model year cycle. When these prices do change, this is usually a part of a public campaign, as in the case of the "economic savings bonus" for the Prius in April 2007. Instead, the manufacturer typically manipulates price by offering cash incentives to dealers. *Automotive News* reports on these types of incentives for all models weekly, and it reports no such dealer cash incentives for the Prius at any point, with the exception of a cash incentive that appears to have affected a very small proportion of vehicles (see section 2.2 for details and citations). Thus, we are looking for confirmation that there were no other cost changes.

Figures 2.14 to 2.16 show kernel density plots of vehicle cost for the two months surrounding each tax change. Looking at the first two tax changes, it is apparent that there was a composition shift. There were clearly more richly equipped Priuses sold in December 2005 than January 2006; and there were clearly fewer sparsely equipped Priuses sold in September 2006 than October 2006. In addition, there appear to be spikes in prices that line up at nearly identical values in each case. This suggests that, for these two experiments, there was a noticeable composition change, but that identically optioned cars cost dealers the same amount before and after the tax changes.

As expected, figure 2.16 tells a different story. In this case, there appears to be both a shift in composition and a dramatic shift downward in cost from March to April 2007. This coincides with the introduction of Toyota's sale. Since the composition does not appear stable, it is important to compare cost adjusted consumer prices to estimate a tax effect. To do so, I assign an options package value to each vehicle from April 2007 and then use the cost measure that would have been in effect for that vehicle, had it been sold in March. Note that this procedure will blur dealer's gross profit and confuse the allocation of benefits between Toyota and dealers, but, if the assignment is correct, it will accurately reflect consumer prices.

Table 2.15 shows the MSRP and dealer invoice prices for each options package before and after the price change, according to Edmunds.com. I assign each Prius an options package as follows. I add the base price, the delivery charge ($660), and the cost of floor mats ($141) to each options package according to April prices. Because locations with lower than average advertising fees will have costs below this number, I assign cut-off values at $100 below this number. Any Prius with a price greater than this number, but less than the next highest cut-off is assigned to that options package. Having assigned options packages, I then add the cost difference back into dealer cost. If the vehicles were correctly assigned, then the dealer cost will now be identical to the dealer cost in March for the same vehicle.

Figure 2.17 shows the adjusted April cost distribution against the March distribution. These now look remarkably similar in the location of peaks. This highlights again that there was a composition change, with more expensive options packages being sold in April. Overall, this similarity of the March and April distributions suggests that the adjustment procedure did a good job assigning options values based on cost.

Figures 2.14 to 2.17 provide visual evidence to support the industry reports that suggest costs were stable through the tax windows, once one has accounted for the April price cut. A statistical test of this proposition can be performed as follows.

First, I use the same procedure as described above to assign options packages to each vehicle based on cost to assign all 2006 and 2007 Priuses an options package. Then, I run a regression of the form:

$$Cost_{ij} = \alpha H + \psi_j + \varepsilon_{ij}$$

separately for each four week period surrounding a tax change, where H is a dummy for the high tax portion of that window and ψ are options package fixed effects. The regressions exclude a small number of vehicles that lie below the estimated minimum invoice price and vehicles more than \$1,000 above the maximum. Table 2.16 reports the coefficients and standard errors on α. Sensitivity analysis showed that some results are sensitive to the size of the downward adjustment for minimum prices in a tier. I thus report a variety of values.

Overall, this table confirms the industry reports and visual evidence. Dealer cost did not change within the tax windows, except for the known sale in April. Column 1 shows that, in only 1 of 6 cases is the difference in January and December costs statistically significant. In no case is the September to October difference significant. In the March versus adjusted April comparison (column 3), 3 of 6 cases yield a significant difference. This may be due to imperfections in the cost adjustment, or it may be a statistical anomaly. As expected, if this procedure is used to compare the March and the unadjusted April numbers, there is a vast difference in costs.

2.11 Figures and Tables

Figure 2.1: Average Days to Turn of Toyota Prius

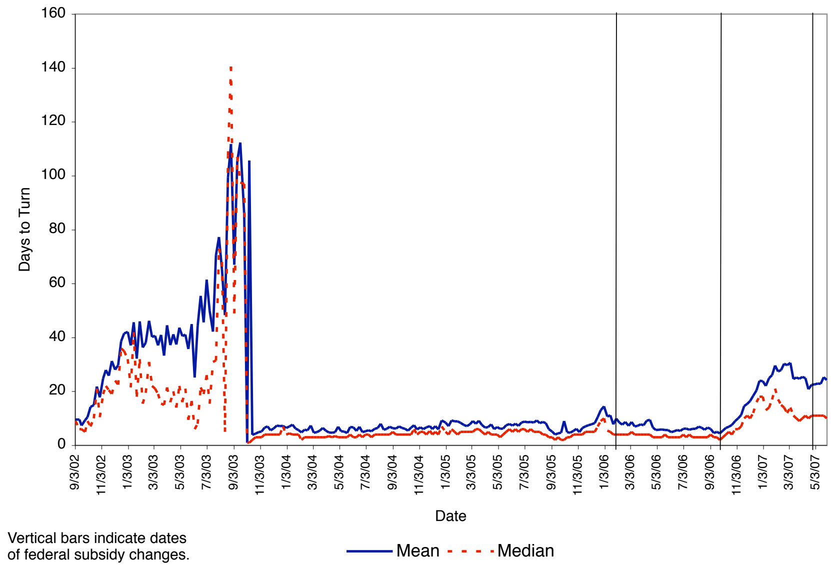

Figure 2.2: Mean Reported Length of Wait List for Priuses by Month

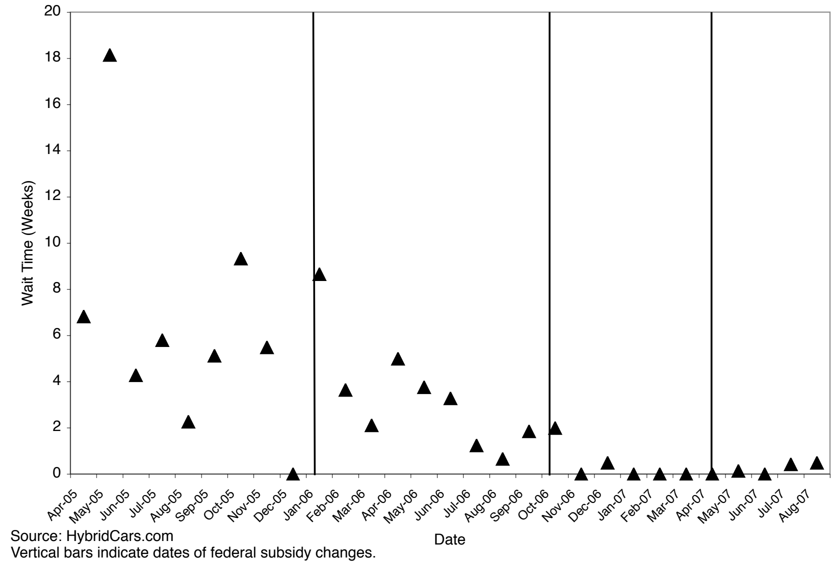

Figure 2.3: Distribution of Sales in December 2005 and January 2006, Prius and Non-Hybrid Toyota Sedans (First Tax Change)

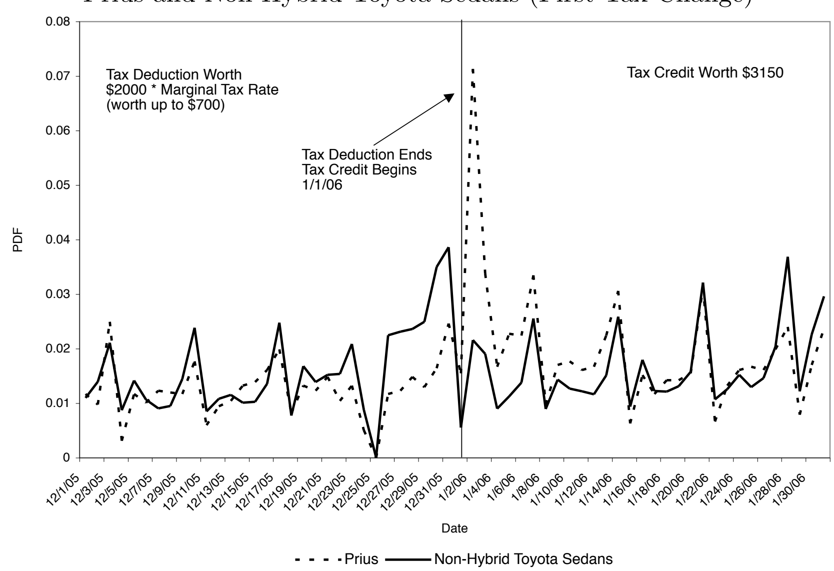

Figure 2.4: Distribution of Sales in December 2004 and January 2005, Prius and Non-Hybrid Toyota Sedans (No Tax Change)

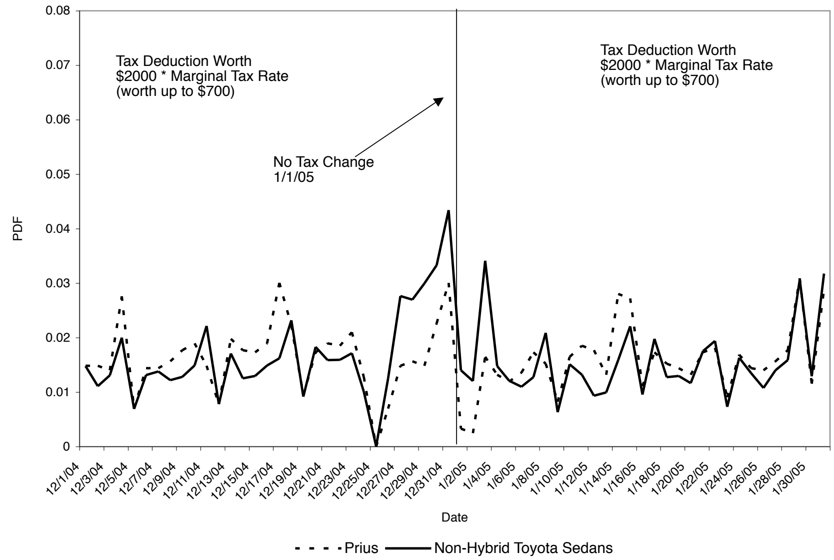

Figure 2.5: Distribution of Sales in September 2006 and October 2006, Prius and Non-Hybrid Toyota Sedans (Second Tax Change)

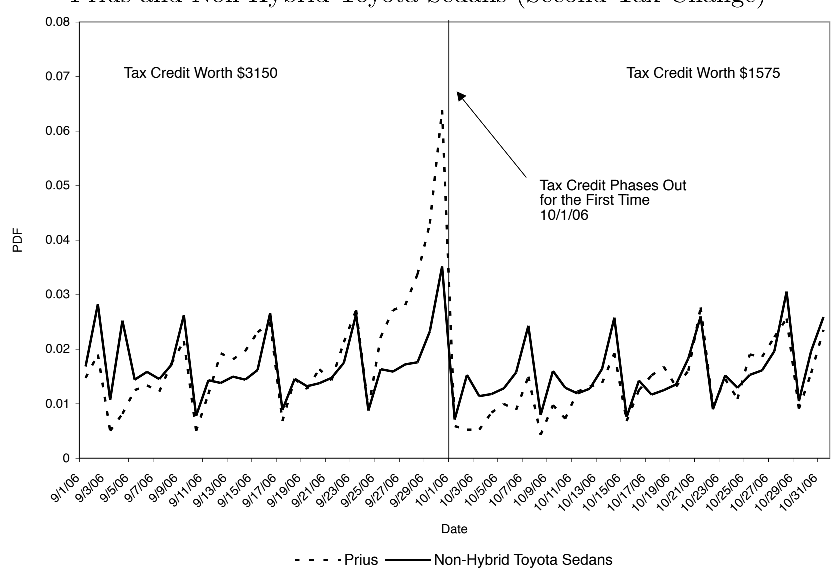

Figure 2.6: Distribution of Sales in September 2005 and October 2005, Prius and Non-Hybrid Toyota Sedans (No Tax Change)

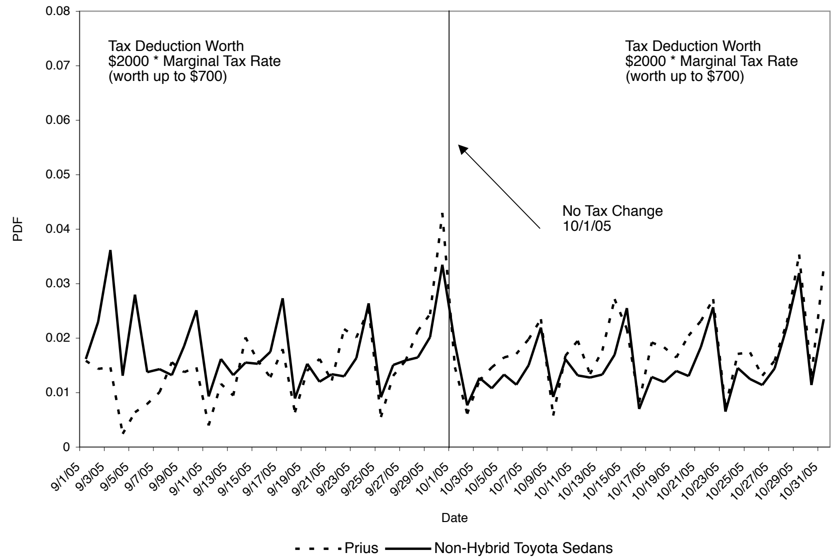

Figure 2.7: Distribution of Sales in March 2007 and April 2007, Prius and Non-Hybrid Toyota Sedans (Third Tax Change)

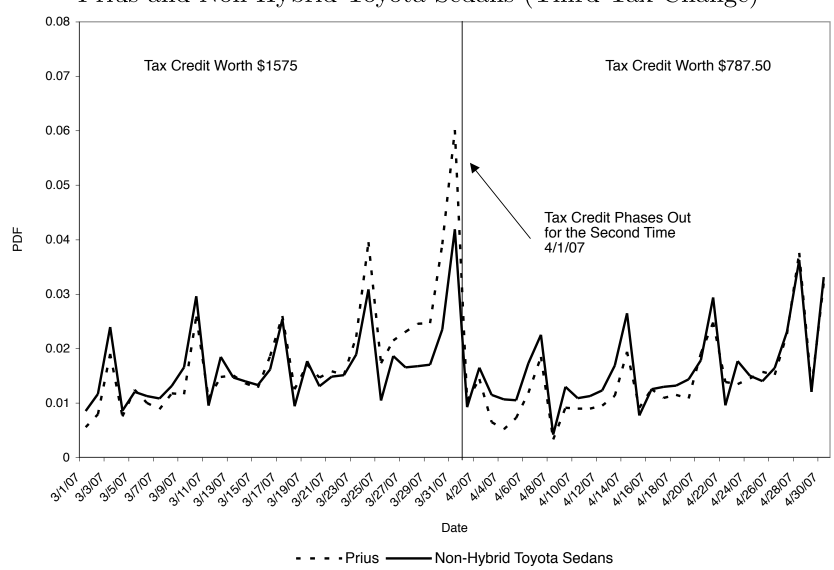

Figure 2.8: Distribution of Sales in March 2006 and April 2006, Prius and Non-Hybrid Toyota Sedans (No Tax Change)

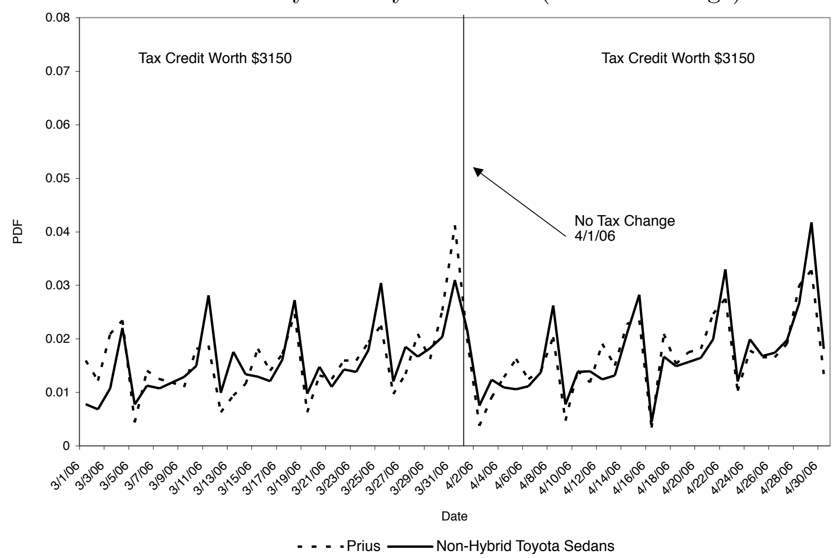

Figure 2.9: Mean Prius Prices, December 2005 and January 2006

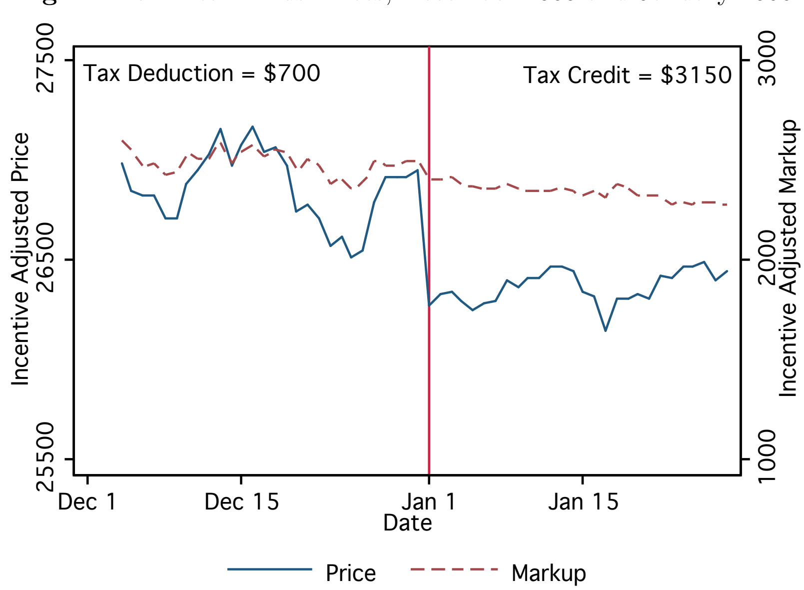

Figure 2.10: Mean Prius Prices, September and October 2006

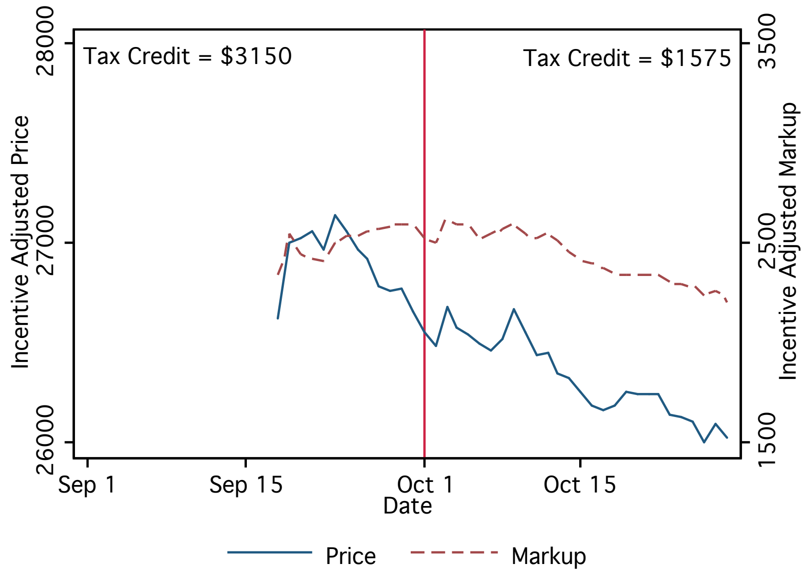

Figure 2.11: Mean Prius Prices, March and April 2007

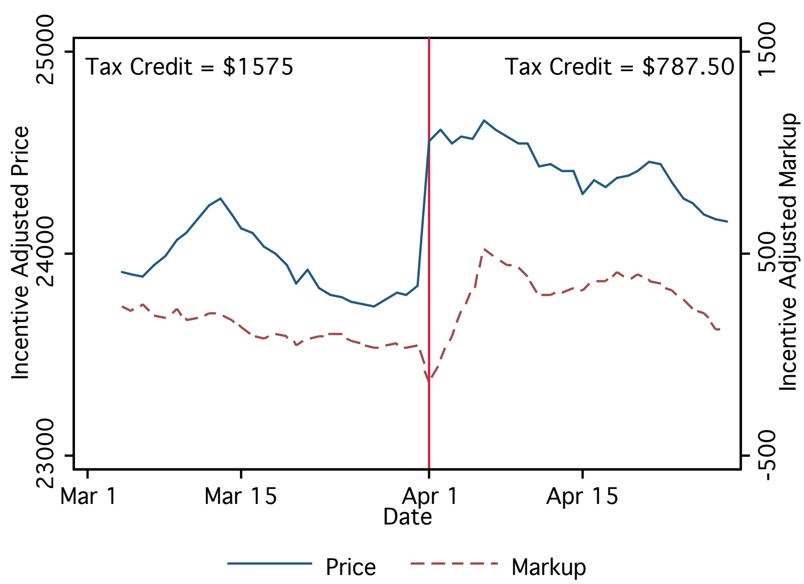

Figure 2.12: Google Trends Chart of Searches for 'Hybrid Tax Credit'

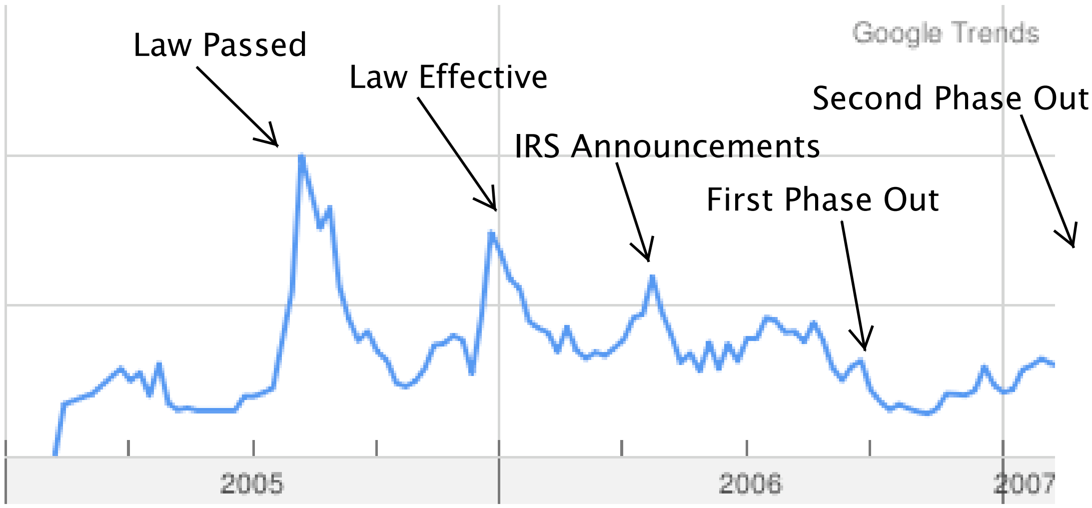

Figure 2.13: Estimated Density of Dealer Cost for Prius, September 2006 to March 2007

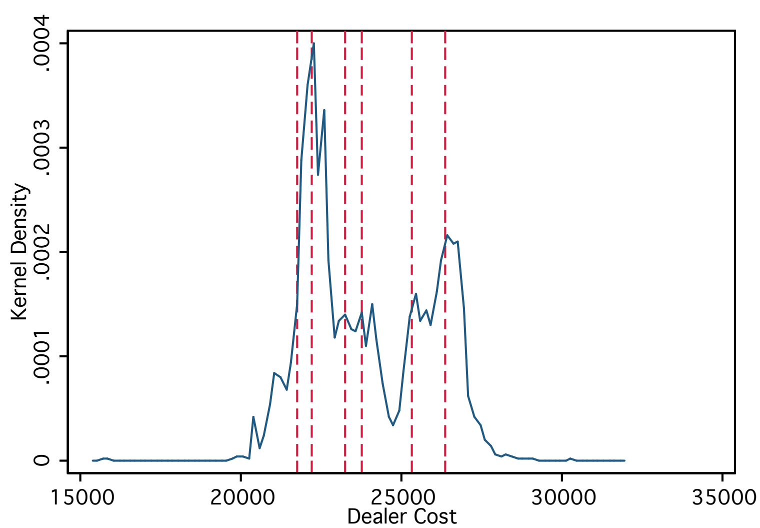

Vertical bars indicate estimated minimum invoice prices for each options package.

Figure 2.14: Estimated Density of Dealer Cost for Prius, December 2005 and January 2006

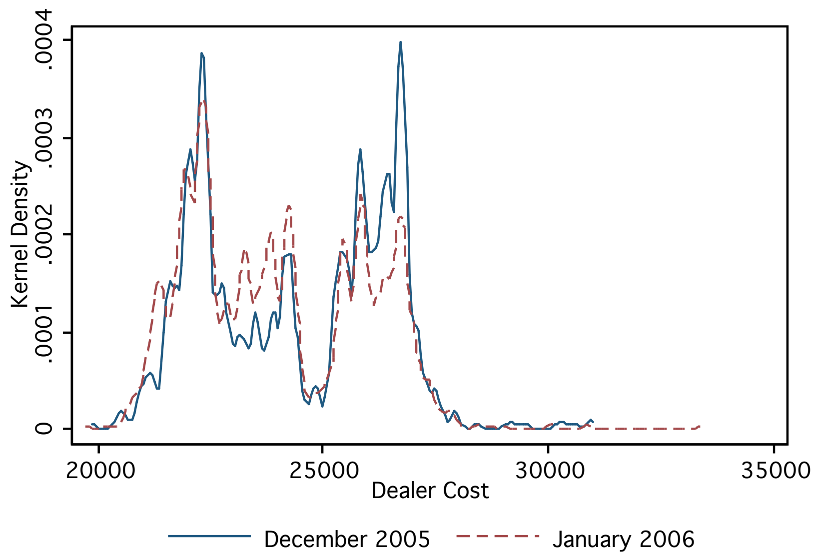

Figure 2.15: Estimated Density of Dealer Cost for Prius, September and October 2006

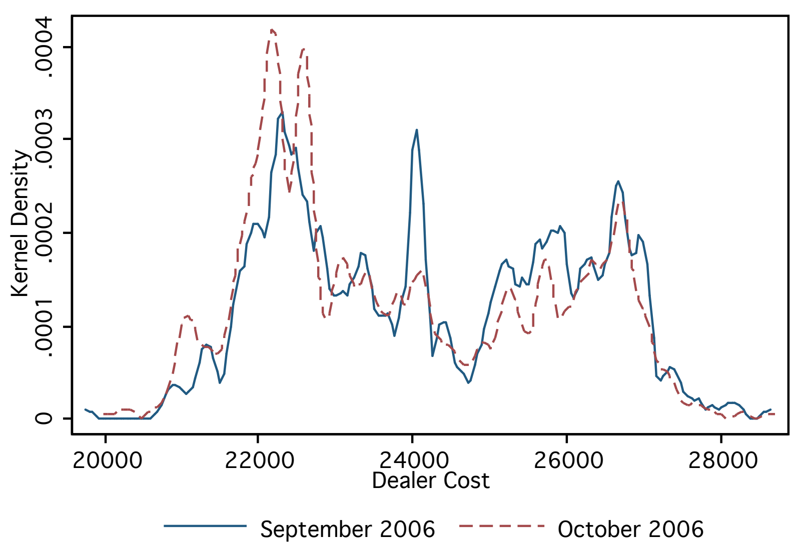

Figure 2.16: Estimated Density of Dealer Cost for Prius, March and Original April 2007

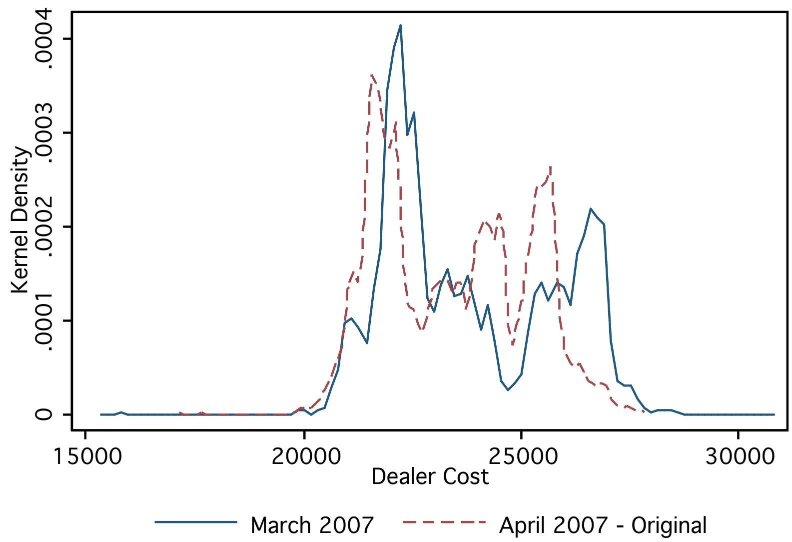

Figure 2.17: Estimated Density of Dealer Cost for Prius, March and Adjusted April 2007

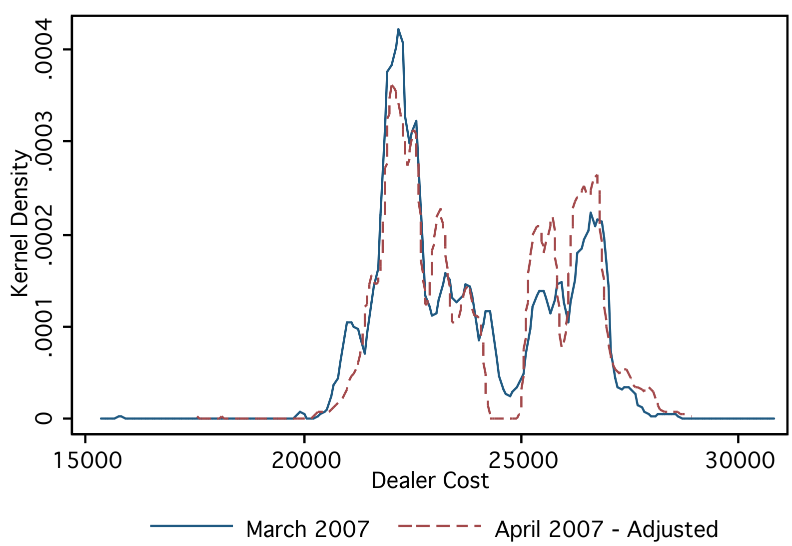

Table 2.1: Variation in Federal Tax Incentives for the Toyota Prius

Date Effective	Tax Incentive
Jan 1 2001 to Dec 31 2005	$2,000 deduction (up to $700 value in 2005)
Jan 1 2006 to Sep 30 2006	$3,150 credit
Oct 1 2006 to Mar 31 2007	$1,575 credit
Apr 1 2007 to Sep 30 2007	$787.50 credit
Oct 1 2007 forward	no credit

Source: Internal Revenue Service

The maximum value of the deduction varies slightly over time because the top marginal tax rate varies.

Table 2.2: Final Sample of Priuses by Model Year
(Sample Size, Mean Incentive Adjusted Price and Incentive Adjusted Markup)

Sample Size

	First Tax Change		Second Tax Change		Third Tax Change		Entire Sample
	Last 2 Weeks	First 2 Weeks	Last 2 Weeks	First 2 Weeks	Last 2 Weeks	First 2 Weeks	
Model Year	December 2005	January 2006	September 2006	October 2006	March 2007	April 2007	Total
2003							2,381
2004							7,776
2005	80	82					20,897
2006	433	925	723	69	2	2	15,715
2007			460	383	1,777	729	16,561
Touring 2007			20	43	125	43	1,376
Total	513	1,007	1,203	495	1,904	774	64,706

Incentive Adjusted Price, Exclusive of Subsidy (Mean)

	First Tax Change		Second Tax Change		Third Tax Change		Entire Sample
	Last 2 Weeks	First 2 Weeks	Last 2 Weeks	First 2 Weeks	Last 2 Weeks	First 2 Weeks	
Model Year	December 2005	January 2006	September 2006	October 2006	March 2007	April 2007	Total
2003							21,068
2004							24,438
2005	25,190	24,163					25,110
2006	26,855	26,340	25,914	25,683	24,854	23,300	26,397
2007			26,811	26,497	23,832	24,512	24,748
Touring 2007			27,917	28,208	26,081	26,378	26,663
Total	26,595	26,163	26,290	26,532	23,981	24,613	25,134

Incentive Adjusted Markup, Exclusive of Subsidy (Mean)

	First Tax Change		Second Tax Change		Third Tax Change		Entire Sample
	Last 2 Weeks	First 2 Weeks	Last 2 Weeks	First 2 Weeks	Last 2 Weeks	First 2 Weeks	
Model Year	December 2005	January 2006	September 2006	October 2006	March 2007	April 2007	Total
2003							611
2004							2,279
2005	1,999	1,465					2,270
2006	2,489	2,366	2,296	2,217	408	-456	2,362
2007			2,558	2,554	59	253	804
Touring 2007			2,771	2,595	199	933	1,330
Total	2,412	2,293	2,404	2,511	69	289	1,837

Shading indicates the model year and trim level used in the primary estimation for each tax change.

Table 2.3: Sales Share Differences
(Toyota Prius v. Non-Hybrid Toyota Sedans Around Tax Changes)

	Difference in Sales Share	
	Two Weeks Before Tax Change Minus Two Weeks After Change	Rest of First Month Minus Rest of Second Month
Prius		
December 2005 & January 2006	-18%	-6%
September 2006 & October 2006	20%	-3%
March 2007 & April 2007	20%	-4%
Non-Hybrid Toyota Sedans		
December 2005 & January 2006	6%	-7%
September 2006 & October 2006	4%	0%
March 2007 & April 2007	7%	-3%
	Difference-in-Differences (Prius Minus Non-Hybrid Toyota Sedans)	
December 2005 & January 2006	-24%	2%
September 2006 & October 2006	16%	-3%
March 2007 & April 2007	13%	-2%

All sales shares are based on the total sales over the two month period surrounding the tax change, using all model years and trim levels.

Table 2.4: Incentive Adjusted Price Change of Priuses Per Dollar of Tax Change in Four Weeks Surrounding Tax Change

	No Controls	Composition Adjusted	Controls
January 2006 - December 2005	-0.194	-0.037	-0.022
	(0.054)	(0.024)	(0.024)
September 2006 - October 2006	0.199	-0.010	-0.058
	(0.106)	(0.047)	(0.047)
March 2007 - April 2007	-0.864	-0.207	-0.195
	(0.118)	(0.075)	(0.074)

Dependent variable is the incentive adjusted transaction price of a new Prius, exclusive of tax subsidies.
Heteroskedasticity robust standard errors in parentheses.
Jan-Dec estimated based on 2006 Prius, Sep-Oct and Mar-Apr coefficients based on 2007, base trim Prius.
"No Controls" includes only model year and trim level dummies.
"Composition Adjusted" adds vehicle cost as a regressor.
"Controls" adds vehicle cost, day of week dummies, state dummies, end of month dummy, price of gasoline and quadratic vehicle trend.
N=64,706 overall, with the following sample sizes in each tax window: 433 (Dec), 925 (Jan), 460 (Sep), 383 (Oct), 1,777 (Mar) and 729 (Apr).

Table 2.5: Evidence of Heterogeneous Timing Response,
Mean Consumer and Transaction Characteristics Surrounding Each Tax Change

		High Subsidy Prius Buyers within 14 Days of Change	Low Subsidy Prius Buyers within 14 Days of Change	Difference	Standard Error
Demographics					
Percent Female					
	Jan-Dec	35.1%	31.6%		
	Sept - Oct	37.2%	36.6%	1.07%	(1.62)
	Mar - Apr	36.0%	35.7%		
Age (Years)					
	Jan-Dec	44.7	44.7		
	Sept - Oct	45.3	43.1	-0.67	(0.46)
	Mar - Apr	46.0	45.2		
Trade-In Vehicles					
Percent with Trade-In					
	Jan-Dec	28.3%	30.3%		
	Sept - Oct	22.2%	30.0%	-5.41%	(1.45)
	Mar - Apr	37.5%	43.9%		
Trade-In Vintage (Year)					
	Jan-Dec	2000.2	1999.9		
	Sept - Oct	2001.2	2000.9	-0.13	(0.21)
	Mar - Apr	2000.6	2001.1		
Trade-In Actual Cash Value ($)					
	Jan-Dec	2,677	2,683		
	Sept - Oct	2,217	3,065	-609	(358)
	Mar - Apr	2,952	3,965		
Contract Details					
Total Down ($)					
	Jan-Dec	5,359	6,421		
	Sept - Oct	6,070	5,843	-482	(238)
	Mar - Apr	4,511	4,988		
Amount Financed ($)					
	Jan-Dec	24,387	24,379		
	Sept - Oct	24,400	24,565	3	(272)
	Mar - Apr	23,347	23,299		
Percent Purchased Service Contract					
	Jan-Dec	37.7%	45.0%		
	Sept - Oct	42.4%	48.6%	-4.95%	(1.53)
	Mar - Apr	45.9%	49.0%		
Service Contract Profit for Dealers ($)					
	Jan-Dec	586	738		
	Sept - Oct	771	871	-94	(27)
	Mar - Apr	779	838		
Percent Purchased Life Insurance					
	Jan-Dec	0.54%	0.23%		
	Sept - Oct	0.00%	0.26%	-0.01%	(0.24)
	Mar - Apr	0.45%	0.55%		
Total Cost of After-Market Options and Fees($)					
	Jan-Dec	2,816	3,177		
	Sept - Oct	3,592	3,455	-124	(50.83)
	Mar - Apr	3,067	3,160		
Percent with Buy Rate < APR					
	Jan-Dec	13.6%	14.5%		
	Sept - Oct	13.7%	17.5%	-8.38%	(1.13)
	Mar - Apr	3.5%	18.0%		

The difference estimate is a coefficient on a dummy equal to one if the transaction is within 2 weeks of a tax change, on the high subsidy side, from a regression with dummy variables for each 4 week window.

Sample sizes in each window are as follows: 433 (Dec), 925 (Jan), 460 (Sep), 383 (Oct), 1,777 (Mar) and 729 (Apr).

Table 2.6: Estimated Bounds on the Incentive Adjusted Price Change of Priuses Per Dollar of Tax Change in Four Week Window Surrounding Tax Change

	No Controls		Composition Adjusted		Controls		Extra Controls	
	Lower Bound	Upper Bound	Lower Bound	Upper Bound	Lower Bound	Upper Bound	Lower Bound	Upper Bound
Jan06 - Dec05	-0.194	0.622	-0.037	0.249	-0.022	0.270	-0.006	0.280
	(0.054)	(0.076)	(0.024)	(0.034)	(0.024)	(0.033)	(0.023)	(0.032)
Sep06 - Oct06	0.199	0.581	-0.010	0.127	-0.058	0.082	-0.040	0.105
	(0.106)	(0.182)	(0.047)	(0.064)	(0.047)	(0.065)	(0.045)	(0.063)
Mar06 - Apr06	-0.864	2.043	-0.207	1.311	-0.195	1.264	-0.090	1.263
	(0.118)	(0.173)	(0.075)	(0.109)	(0.074)	(0.107)	(0.069)	(0.098)

Dependent variable is the incentive adjusted transaction price of a new Prius, exclusive of tax subsidies.
Heteroskedasticity robust standard errors in parentheses. Upper bound standard errors are from a nonparametric boostrap with 5,000 repetitions.
Jan-Dec estimated based on 2006 Prius, Sep-Oct and Mar-Apr coefficients based on 2007, base trim Prius.
"No Controls" includes only model year and trim level dummies.
"Composition Adjusted" adds vehicle cost as a regressor.
"Controls" adds vehicle cost, day of week dummies, state dummies, end of month dummy, price of gasoline and quadratic vehicle trend.
"Extra Controls" adds sex, age, total after-market options, a dummy for an APR above the buy rate, a dummy for life insurance, a dummy for the presence of a trade-in, trade-in actual value, trade-in vintage, and a dummy for a service contract.
N=64,706 overall, with the following sample sizes in each tax window: 433 (Dec), 925 (Jan), 460 (Sep), 383 (Oct), 1,777 (Mar) and 729 (Apr).

Table 2.7: Estimated Bounds on the Incentive Adjusted Price Change of Priuses Per Dollar of Tax Change Using Sample Proportions to Reflect All Model Years and Trim Levels

	No Controls		Composition Adjusted		Controls		Extra Controls	
	Lower Bound	Upper Bound	Lower Bound	Upper Bound	Lower Bound	Upper Bound	Lower Bound	Upper Bound
Jan06 - Dec05	-0.194	0.551	-0.037	0.224	-0.022	0.244	-0.006	0.256
	(0.054)	(0.075)	(0.024)	(0.032)	(0.024)	(0.032)	(0.023)	(0.030)
Sep06 - Oct06	0.199	1.767	-0.010	0.550	-0.058	0.518	-0.040	0.516
	(0.106)	(0.120)	(0.047)	(0.078)	(0.047)	(0.077)	(0.045)	(0.072)
Mar06 - Apr06	-0.864	2.067	-0.207	1.324	-0.195	1.277	-0.090	1.274
	(0.118)	(0.172)	(0.075)	(0.110)	(0.074)	(0.108)	(0.069)	(0.099)

Dependent variable is the incentive adjusted transaction price of a new Prius, exclusive of tax subsidies.

Heteroskedasticity robust standard errors in parentheses. Upper bound standard errors are from a nonparametric boostrap with 5,000 repetitions.

Jan-Dec estimated based on 2006 Prius, Sep-Oct and Mar-Apr coefficients based on 2007, base trim Prius.

"No Controls" includes only model year and trim level dummies.

"Composition Adjusted" adds vehicle cost as a regressor.

"Controls" adds vehicle cost, day of week dummies, state dummies, end of month dummy, price of gasoline and quadratic vehicle trend.

a trade-in, trade-in actual value, trade-in vintage, and a dummy for a service contract.

N=64,706 overall, with the following sample sizes in each tax window: 433 (Dec), 925 (Jan), 460 (Sep), 383 (Oct), 1,777 (Mar) and 729 (Apr).

The upper bound is calculated using 471 (Jan), 189 (Sep) and 722 (Mar) prices, to mimic the proportion of Priuses of all model years and trim levels in each period.

Table 2.8: Estimated Bounds on the Incentive Adjusted Price Change of Priuses Per Dollar of Tax Change Using Various Assumptions to Calculate Financing Incentives

	Discount Rate = 2%		Discount Rate = 4%		Discount Rate = 6%		Discount Rate = 8%	
Unadjusted APR	Lower Bound	Upper Bound	Lower Bound	Upper Bound	Lower Bound	Upper Bound	Lower Bound	Upper Bound
Jan06 - Dec05	-0.024	0.296	**-0.022**	**0.270**	-0.021	0.258	-0.021	0.251
	(0.027)	(0.038)	**(0.024)**	**(0.033)**	(0.023)	(0.031)	(0.022)	(0.030)
Sep06 - Oct06	-0.108	0.043	**-0.058**	**0.082**	-0.032	0.103	-0.017	0.117
	(0.055)	(0.075)	**(0.047)**	**(0.065)**	(0.043)	(0.062)	(0.042)	(0.061)
Mar06 - Apr06	-0.558	1.111	**-0.195**	**1.264**	0.014	1.382	0.146	1.467
	(0.085)	(0.125)	**(0.074)**	**(0.107)**	(0.070)	(0.101)	(0.068)	(0.098)
Prius Adjusted APR								
Jan06 - Dec05	-0.026	0.295	-0.024	0.270	-0.022	0.258	-0.022	0.251
	(0.027)	(0.039)	(0.024)	(0.034)	(0.023)	(0.031)	(0.022)	(0.030)
Sep06 - Oct06	-0.114	0.037	-0.062	0.079	-0.034	0.102	-0.018	0.116
	(0.057)	(0.076)	(0.048)	(0.067)	(0.044)	(0.063)	(0.042)	(0.061)
Mar06 - Apr06	-0.523	1.106	-0.167	1.276	0.037	1.397	0.165	1.481
	(0.085)	(0.128)	(0.075)	(0.110)	(0.070)	(0.103)	(0.068)	(0.100)
Age Adjusted APR								
Jan06 - Dec05	-0.015	0.457	-0.018	0.363	-0.019	0.317	-0.019	0.292
	(0.045)	(0.053)	(0.033)	(0.041)	(0.028)	(0.036)	(0.026)	(0.033)
Sep06 - Oct06	-0.042	0.324	-0.019	0.236	-0.005	0.199	0.004	0.181
	(0.089)	(0.125)	(0.064)	(0.094)	(0.053)	(0.080)	(0.048)	(0.072)
Mar06 - Apr06	-0.515	2.245	-0.170	1.941	0.032	1.813	0.160	1.765
	(0.137)	(0.166)	(0.102)	(0.132)	(0.086)	(0.116)	(0.079)	(0.108)

Dependent variable is the incentive adjusted transaction price of a new Prius, exclusive of tax subsidies.
Heteroskedasticity robust standard errors in parentheses. Upper bound standard errors are from a nonparametric boostrap with 5,000 repetitions.
The discount rate is the annual discount rate applied to the loan in order to calculate the present discounted value of interest rate subsidies.
Prius adjusted APR substracts .5 from the Federal Reserve 48-month car loan interest rate to adjust for average Prius premium when there are no subsidies.
Age adjusted APR subratcts .05 from the Federal Reserve rate per year of age above the mean to adjust for average credit differences across different ages.
Specification corresponds to column three ("Controls") from Table 6. Regression includes model year and trim level dummies, vehicle cost, day of week dummies, state dummies, end of month dummy, price of gasoline and quadratic vehicle trend.
Jan-Dec estimated based on 2006 Prius, Sep-Oct and Mar-Apr coefficients based on 2007, base trim Prius.
N=64,706 overall, with the following sample sizes in each tax window: 433 (Dec), 925 (Jan), 460 (Sep), 383 (Oct), 1,777 (Mar) and 729 (Apr).

Table 2.9: Estimated Time Period Effects on Incentive Adjusted Price of Non-Hybrid Sedans in Four Week Window Surrounding Tax Change

	Comparable Sedans	Non-Hybrid Toyota Sedans	Fuel-Efficient Sedans
January 2006 - December 2005	-49.7	79.3	-69.6
	(27.0)	(17.4)	(24.5)
September 2006 - October 2006	-24.4	25.5	-4.5
	(22.9)	(16.2)	(22.8)
March 2007 - April 2007	49.6	-23.6	-12.8
	(23.2)	(16.6)	(22.9)

Dependent variable is the incentive adjusted transaction price of a new sedan.

Heteroskedasticity robust standard errors in parentheses.

Specification corresponds to column three ("Controls") from Table 6. Regression includes model year / trim level / engine type dummies, vehicle cost, day of week dummies, state dummies, end of month dummy, price of gasoline and quadratic vehicle trend.

Comparable vehicles are the 6, Accord, Aveo, Charger, Civic, Fit, Five Hundred, G6, Grand Prix, Impala, Impreza, LaCrosse, Malibu, Maxima, Milan, Montego, Passat, Sable, Sentra, tC, xA, xB and the non-hybrid Toyota sedans.

Non-hybrid Toyota sedans include the Avalon, Camry, Camry Solara, Corolla, Matrix and Yaris.

Fuel efficient sedans include the subset of these vehicles with a combined EPA fuel economy above 29 miles per gallon: Aveo, Civic, Corolla, Fit, Matrix, Sentra, xA and xB.

Overall, N=1,536,075 for comparable sedans, 570,059 for non-hybrid Toyota sedans and 584,418 for fuel-efficient sedans.

Table 2.10: Estimated Log Price Premium For Toyota Models in High Tax Subsidy Window

		First Two Weeks Jan 06 v. Last Two Weeks Dec 05			Last Two Weeks Sep 06 v. Last Two Weeks Oct 06			First Two Weeks Jan 06 v. Last Two Weeks Dec 05		
		Log Premium	Standard Error		Log Premium	Standard Error		Log Premium	Standard Error	
Toyota Nameplate Cars										
	Avalon	0.0026	(0.0026)		-0.0065	(0.0025)	*	0.0026	(0.0030)	
	Camry	-0.0014	(0.0021)		0.0014	(0.0017)		0.0043	(0.0018)	*
	Camry Solara	-0.0056	(0.0062)		-0.0003	(0.0056)		0.0022	(0.0052)	
	Corolla	-0.0083	(0.0024)	*	-0.0011	(0.0023)		0.0019	(0.0022)	
	Matrix	0.0013	(0.0057)		-0.0036	(0.0051)		0.0031	(0.0053)	
	Yaris				0.0015	(0.0035)		-0.0040	(0.0036)	
Toyota Nameplate Light Trucks										
	4Runner	-0.0032	(0.0031)		-0.0020	(0.0033)		0.0036	(0.0033)	
	FJ Cruiser				-0.0044	(0.0037)		-0.0032	(0.0038)	
	Highlander	0.0046	(0.0031)		-0.0019	(0.0028)		0.0014	(0.0030)	
	Land Cruiser	0.0073	(0.0159)		-0.0144	(0.0143)		0.0038	(0.0175)	
	RAV4	-0.0009	(0.0032)		-0.0013	(0.0024)		-0.0049	(0.0022)	*
	Sequoia	0.0069	(0.0044)		-0.0142	(0.0045)	*	0.0056	(0.0053)	
	Sienna	-0.0098	(0.0026)	*	0.0025	(0.0023)		0.0051	(0.0026)	*
	Tacoma	0.0062	(0.0027)	*	0.0017	(0.0025)		0.0018	(0.0024)	
	Tundra	0.0093	(0.0038)	*	-0.0033	(0.0032)		-0.0342	(0.0270)	
Lexus Nameplate Vehicles										
	ES 330	-0.0035	(0.0032)		0.1902	(0.0373)	*			
	ES 350				0.0017	(0.0028)		-0.0029	(0.0028)	
	GS 300	0.0016	(0.0052)		-0.0039	(0.0133)		-0.0621	(0.0387)	
	GS 350				-0.0041	(0.0056)		-0.0054	(0.0060)	
	GX 470	0.0016	(0.0051)		0.0000	(0.0059)		0.0046	(0.0066)	
	IS 250	-0.0019	(0.0051)		0.0006	(0.0051)		-0.0024	(0.0046)	
	IS 300	0.1196	(0.0514)	*						
	IS 350	0.0026	(0.0064)		0.0065	(0.0067)		0.0061	(0.0126)	
	LS 430	-0.0031	(0.0052)		-0.0094	(0.0094)				
	LS 460							-0.0056	(0.0045)	
	LX 470	-0.0085	(0.0094)		-0.0131	(0.0097)		0.0163	(0.0141)	
	RX 330	-0.0024	(0.0029)		0.0609	(0.0216)	*	-0.0439	(0.0525)	
	RX 350				-0.0003	(0.0031)		-0.0062	(0.0111)	
	SC 430	-0.0005	(0.0082)		0.0078	(0.0114)		0.0101	(0.0120)	
Scion Nameplate Vehicles										
	tC	0.0002	(0.0029)		0.0033	(0.0026)		0.0009	(0.0028)	
	xA	0.0026	(0.0047)		-0.0045	(0.0039)		-0.0086	(0.0057)	
	xB	0.0043	(0.0041)		-0.0052	(0.0035)		0.0016	(0.0057)	

Dependent variable is the log of the incentive adjusted transaction price of a new vehicle.
Heteroskedasticity robust standard errors in parentheses. * Denotes significance at 95% level.
Specification corresponds to column three ("Controls") from Table 6. Regression includes model year / trim level / engine type dummies, vehicle cost, day of week dummies, state dummies, end of month dummy, price of gasoline and quadratic vehicle trend.

Table 2.11: State Tax Incentives for the Toyota Prius

State	Subsidy Type	Amount*	Start Date	End Date	Sample Size: Subsidy Off	Sample Size: Subsidy On
Colorado	Income Tax Credit	$3150 to $4622	7/1/00	Still Effective		
Connecticut	Full Sales Tax Exemption	6% ($1500)	10/1/04	Still Effective	47	100
District of Columbia	Full Sales Tax Exemption	7% ($1750)	4/15/05	Still Effective	32	65
Lousiana	Income Tax Credit	2% ($500)	1/1/91	Still Effective		
Maine	Partial Sales Tax Exemption	2.5% ($625)	1/1/97	12/31/05	169	193
Maryland	Partial Sales Tax Exemption	$1000 max	7/1/00	7/1/04	2,487	509
New Mexico	Full Sales Tax Exemption	3% ($750)	7/1/04	Still Effective		
New York	Income Tax Credit	$2000	1/1/01	12/31/04	1,305	214
New York	Partial Sales Tax Exemption	$240	1/1/01	2/28/05		
Oregon	Income Tax Credit	$1500	1/1/98	Still Effective		
Pennsylvania	Rebate	$500	3/25/05	Still Effective	525	1,574
South Carolina	Income Tax Credit	$630	1/1/06	Still Effective	46	182
West Virginia	Income Tax Credit	$3150 to $3750	7/1/97	6/30/06	27	11
Total					4,638	2,848

* For sales tax exemptions, the value of the exemption on a $25,000 car is included in parentheses, for ease of comparison.

Table 2.12: Incentive Adjusted Price Change of Priuses Per Dollar of State Tax Incentive Change

	Specification 1	Specification 2
State Tax Incentive	0.051	-0.022
	(0.081)	(0.088)
Dealer Cost		1.062
		(0.011)
State Fixed Effects	X	X
Week Fixed Effects	X	X
N	64,706	64,706
AdjR2	0.25	0.80

Depedent variable is the incentive adjusted transaction price of a new Prius, exclusive of tax subsidies.
Heteroskedasticity robust standard errors, clustered on state, in parentheses.
Controls include day of week dummies, dummy for end of the month, quadratic in vehicle time and model year and trim level dummies.

Table 2.13: Price Sensitivity of the Prius With and Without a Wait List

	Incentive Adjusted Price of Prius
End of Month	-8.07
	(7.51)
End of Month, No Wait List	-108
	(22)
Price of Gasoline	1.19
	(0.67)
Price of Gasoline, No Wait List	5.78
	(1.12)
Weeks on Market	-0.35
	(0.10)
Weeks on Market, No Wait List	-9.85
	(0.64)
N	64,706
R2	0.80

Dependent variable is the incentive adjusted transaction price of a new Prius, exclusive of tax subsidies.
Heteroskedasticity robust standard errors, clustered on state, in parentheses.
Controls include dealer cost, day of week dummies, quadratic in vehicle time and model year and trim level dummies.
Model year 2004, 2005 and 2006 Priuses are coded as having a wait list.

Table 2.14: Estimated Bounds on the Incentive Adjusted Price Change of Other Hybrids Per Dollar of Tax Change

	Toyota Prius		Honda Civic		Ford Escape		Toyota Highlander		Toyota Camry	
	Lower Bound	Upper Bound	Lower Bound	Upper Bound	Lower Bound	Upper Bound	Lower Bound	Upper Bound	Lower Bound	Upper Bound
Jan06 - Dec05	-0.022	0.270	-0.034	0.348	0.164	0.359	0.023	0.443		
	(0.024)	(0.033)	(0.056)	(0.073)	(0.195)	(0.187)	(0.081)	(0.111)		
Sep06 - Oct06	-0.058	0.082					-0.034	1.061	-0.091	0.557
	(0.047)	(0.065)					(0.166)	(0.234)	(0.076)	(0.106)
Mar06 - Apr06	-0.195	1.264					0.279	2.415	-0.401	0.962
	(0.074)	(0.107)					(0.306)	(0.434)	(0.150)	(0.214)

Dependent variable is the incentive adjusted transaction price of a new vehicle, exclusive of tax subsidies.

Heteroskedasticity robust standard errors in parentheses. Upper bound standard errors are from a nonparametric boostrap with 5,000 repetitions.

Specification corresponds to column three ("Controls") from Table 6. Regression includes model year and trim level dummies, vehicle cost, day of week dummies, state dummies, end of month dummy, price of gasoline and quadratic vehicle trend.

The following model years are used: Civic (2006), Escape (2006), Highlander (2006 for first change and 2007 for second and third) and Camry (2007 for both tax changes).
Prius: N=64,706 overall, with the following sample sizes in each tax window: 433 (Dec), 925 (Jan), 460 (Sep), 383 (Oct), 1,777 (Mar) and 729 (Apr).
Civic: N=27,155 overall, with the following sample sizes in each tax window: 210 (Dec) and 364 (Jan).
Escape: N=6,804 overall, with the following sample sizes in each tax window: 67 (Dec) and 82 (Jan).
Highlander: N=9,707 overall, with the following sample sizes in each tax window: 131 (Dec), 218 (Jan), 181 (Sep), 89 (Oct), 240 (Mar) and 111 (Apr).
Camry: N=9,617 overall, with the following sample sizes in each tax window: 422 (Sep), 187 (Oct), 539 (Mar) and 288 (Apr).

Table 2.15: Price Changes Due to "Economic Savings Bonus" for 2007 Prius

Manufacturer's Suggested Retail Price				
Base		$22,175	22,175	0
Package	1	825	825	0
	2	1,175	575	600
	3	2,555	2,105	450
	4	3,180	2,580	600
	5	5,080	3,280	1,800
	6	6,350	4,550	1,800
Dealer Invoice				
Base		20,419	20,419	0
Package	1	659	659	0
	2	960	518	442
	3	2,018	1,895	123
	4	2,518	2,322	196
	5	4,133	2,952	1,181
	6	5,149	4,095	1,054

Source: Edmunds.com

Table 2.16: Dealer Cost Premium in High Tax Period, Within Assigned Options Package
2007 Prius

Options Package Threshold	Jan 06 (v. Dec 05)	Sep 06 (v. Oct 06)	Adjusted Mar 07 (v. Apr 07)	Unadjusted Mar 07 (v. Apr 07)
$25	-8.47	-13.84	-39.34	619.2
	(18.60)	(22.30)	(14.10)	(15.40)
50	-17.16	-12	-31.37	625.2
	(18.40)	(22.00)	(14.00)	(15.40)
75	-16.98	-12.31	-30.38	624.7
	(18.30)	(22.10)	(13.70)	(15.30)
100	-42.62	-13.94	-26.47	629.9
	(18.00)	(22.10)	(13.60)	(15.30)
125	-24.6	-12.45	-19.98	637.1
	(17.60)	(22.10)	(13.50)	(15.20)
150	-22.19	-13.1	-6.52	651.5
	(17.20)	(22.10)	(13.20)	(15.20)
Fixed Effect for Assigned Options Packages	X	X	X	X

Dependent variable is dealer cost. The reported coefficient is a dummy for the high tax period in each window.
Heteroskedasticity robust standard errors in parentheses.
Each coefficient is from a separate regression, which assigns each vehicle in the sample an options package if it is greater than the base price plus options price minus threshold, but less than the similarly defined cutoff for the next highest options package.
N is approximately 1,330 for column 1, 815 for column 2 and 2,370 for columns 3 and 4. The sample varies slightly because minimum and maximum thresholds for censoring vary slightly by the selected threshold.

CHAPTER III

A Cautionary Tale About the Use of Administrative Data: Evidence from Age of Marriage Laws

3.1 Introduction

This paper presents a cautionary note about the use of administrative data in empirical work, especially work devoted to assessing the effects of laws. The behavioral changes induced by laws often represent important tests of theory and, because they are plausibly exogenous, may be exploited to identify various causal effects. Increasingly, scholars studying the effects of laws, and other empirical questions, rely on administrative data. These data are drawn from official records, and are thus thought to provide more reliable information about the behaviors they record than do general surveys. We argue that this presumption may often be incorrect, with important implications for the validity of empirical analysis.

When a particular behavior has been outlawed, there will be both people who modify their behavior as intended by the law and others who attempt to circumvent the law in various ways. In particular, "non-compliers" can either (a) move to an area where the behavior is not outlawed and engage in it there; or (b) break the law where they are. Both these things bias administrative records towards suggesting larger changes in the behavior addressed by the law than actually occurred.[1]

[1]That some agents might migrate to avoid a law to which they would otherwise be subject has been studied in

In the first case, official records correctly indicate that recorded behavior in an area conform to the area's laws, but they miss the fact that some of the people to whom an area's law was directed actually behaved as they would have without the law by moving somewhere else. In the second case, administrative records might not even accurately reflect behavioral changes within a given area. People who break the law have an incentive to conceal that fact, and that incentive is likely especially large when the threat of a sanction is large — as is true in virtually all cases where administrative data are collected.

These potential problems are rarely discussed in studies using administrative data. Moreover, because researchers are predisposed to believe that laws affect behavior in the intended direction, they are unlikely to be vigilant against the possibility that information from an administrative data source suggests exaggerated effects. In practice the consequences of this bias in administrative data may be severe, and it may lead scholars to dramatically incorrect conclusions about a law's effects. Interestingly, while they have other limitations, data from general surveys may not be subject to the particular problems of administrative data. Research which combines data from different sources, or which uses data from one source to assess the validly of results from another, is thus much more likely to identify parameters of interest than research using one type of data.

To assess the practical importance of these issues we study how minimum age of marriage laws, which were designed to reduce early marriage, affected young people's marriage behavior. Before the 1970s, adolescents wishing to marry faced minimum age of marriage restrictions which varied widely across states. These laws converged dramatically over a short period in the early 1970s, so that by 1975 the minimum

a variety of contexts, including migration for tax-related reasons (Asplund, Friberg and Wilander, 2007; Lovenheim, 2008); migration to obtain legal abortions (Blank, George and London, 1996), or migration to avoid age-of-drinking laws (Kenkel, 1993).

permissible age of marriage was identical across almost all states. These legal changes present an ideal opportunity to study how marriage responds to policy initiatives and could be exploited in instrumental variables or other approaches to provide estimates of how early marriage affects various outcomes.

The "gold-standard" data source for studying the effect of marriage laws would seem to be the official marriage statistics from the National Vital Statistics System. These data are drawn from a state's actual marriage certificates, and represent the official administrative record of marriage collected at the point of marriage. Other surveys, like the decennial Census, enquire about marriage history, but rely on a person's retrospective report about behavior. Surprisingly, we find that for marriages in the 1950s, the Vital Statistics evidence suggests dramatically larger effects of minimum age of marriage laws than are found in Census data.

We examine three possibilities for the differences between the two data sources. First, we show that the discrepancy is not the result of mechanical differences between the data sets or recall bias in Census records. Second, because the Census reports marriages by where people live while the Vital Statistics records marriages by where they occur, the different estimates from the two data sources might actually describe the same underlying behavior, if people travelled to more permissive states to marry. We show that while there was some marriage related "migration" among young people in the relevant cohorts, migration cannot explain the overall discrepancy between the data sources. The third possible explanation is that young people in the 1950s actually married at younger ages than legally permitted in their states but misrepresented their ages on official marriage records. Consistent with this interpretation, we show that most of the differences between retrospective Census and Vital Statistics records disappear by 1970, before the convergence in legal marriage ages across states but after documentary evidence of proof of age, such as driver's licenses and social security cards, became more common.

Our results suggest that minimum age of marriage laws did lower young marriage as intended. However, because in the 1950s many young people circumvented the laws by moving to other states or by misrepresenting their ages on official records, Vital Statistic records indicate much larger changes in marriage than actually occurred. More generally, our results illustrate how data quality might be affected by the incentives that agents face to give misinformation and how that incentive may be particularly high for administrative records. Standard regression or instrumental variables estimates of how behavior changed which rely exclusively on administrative data may thus be very much at variance with changes that actually occurred in the population.

Below, we review the history of minimum age of marriage laws and describe available information about marriage in the Vital Statistics and Census. Section 3.3 presents results from these two data sources, showing that the estimated effects of age of marriage laws differ sharply across them. In Section 3.4 we analyze, in turn, the importance of mechanical issues and recall bias in Census data, marriage migration and systematic misrepresentation in the Vital Statistics data for explaining the discrepancy across the two data sources. Section 3.5 concludes.

3.2 History of Minimum Age Laws and Data on Early Marriage

In 1950, laws governing minimum ages of marriage varied widely across states and across genders.[2] Over the short period from 1971 to 1975, a burst of legislative activity led to a dramatic convergence in laws across states and genders.[3] In 1980,

[2]We gathered information on legislative statutes about age of marriage laws in each state by tracing back in time, through successive amendments, from when most of these laws were initially introduced in the early 1800s.

[3]These changes were likely occasioned by the constitutional amendment in 1971 making 18 the voting age in national elections, and by demands from the Women's Rights movement for legal parity for women in all laws.

an 18 year old man or woman wishing to marry could do so without parental consent in all but three states in the country. Merely ten years earlier, in 1970, an 18 year old man would not have been free to marry without parental consent in more than 40 states. Minimum marriage ages were historically lower for women, but even so about 10 states prohibited 18 year old marriage without parental consent for young women in 1970. Table 3.1 shows some of the variation in minimum ages without parental consent across states and over time.

We study the effects of minimum age of marriage laws in the post-World War II period on marriage behavior.[4] We focus throughout on "non-consent" laws, the minimum age of marriage not requiring parental consent. Generally, the minimum ages in these laws are above 16 and are thus likely to "bite", in the sense that non-trivial amounts of people may actually wish to marry at these ages. Restrictions on marriage at ages as young as 12 or 14 (some of the ages in "consent" laws) probably affected few young people. Furthermore, changes in marriage induced by non-consent laws have nothing to do with the preferences of a third party, such as a parent. How these laws affected the marital behavior of young persons across various cohorts can be investigated with data from two different sources: the National Vital Statistics or the decennial Census.[5] We discuss these in turn.

Vital Statistics marriage data come from information submitted voluntarily by states from 1940-1995. The data are from the marriage certificates filled out by couples just prior to marriage and as a requirement to having the marriage legally

[4]Dahl (2005), who uses laws as an instrument for adult poverty with Census data, is the only other example we are aware of which exploits minimum age of marriage laws. His work differs from ours in several ways. First, Dahl focuses exclusively on women and, because he focuses on marriages with parental consent, he analyzes marriages at much younger ages than we do. Dahl mentions and briefly analyzes the potential for cross-state marriages, but he suggests that this is evidence that state laws are binding. By contrast, we think the fundamental lesson to be learned from marriage migration is that it can create large differences between administrative and survey data. Dahl does not mention the potential for misrepresentation, which is central to our findings.

[5]Apart from the Census and Vital Statistics, the other systematic data source is the June supplement to the Current Population Survey, which includes questions about marriage in some years. Relative to the Census, the June CPS samples are smaller and cover more recent years which are of less interest to us.

registered. During the years it collected marriage data, Vital Statistics saw a large increase in the number of states providing information — from less than 20 in 1950 to virtually all states in 1980. Vital Statistics marriage information is not available for all states in all years and is, for most years, not available in age-disaggregated form.[6] Age-specific information is available for the eighteen reporting states in 1950 (and for several years in the 1940s), but is not available again until 1968. For the remainder of the 1950s and 1960s, Vital Statistics only reports information about the number of marriages for aggregate categories, such as marriages among persons aged 14-19 or 20-24. Because the data are collected at the time of marriage and are official administrative data, the Vital Statistics is typically considered the authoritative source of information on marriage in the U.S.

The decennial Censuses represent an alternative source of data about marriage outcomes for the populations of young people affected by marriage laws. In the 1960, 1970 and 1980 Censuses, respondents were asked about the date of their first marriage and their date of birth. Individuals' age at marriage can be readily calculated from these two pieces of retrospective information. The Census contains very large samples and provides age-disaggregated information about marriage. But retrospective Census information about marriage has the drawback that reports may be subject to recall error. Further, location information in the Census is limited to a person's state of residence and state of birth; hence, where a marriage occurred is not perfectly known. Their differences notwithstanding, if lower non-consent ages lowered the likelihood of early marriage, this effect should be evident and of similar magnitude in both data sources. We examine these patterns in the next section.

[6] Data on marriages before 1968 are available in published annual reports, available online at http://www.cdc.gov/nchs/products/pubs/pubd/vsus/vsus.htm. Microdata files for 1968 to 1995 are maintained by the National Bureau of Economic Research at http://www.nber.org/data/marrdivo.html.

3.3 Estimates of the Effect of Marriage Laws on Early Marriage from Vital Statistics and Census Data

We begin with a regression analysis of marriages from 1951-1979. Vital Statistics information for these years is not age-disaggregated, so we focus on marriage for the 14-19 age group. We convert the marriage counts reported by Vital Statistics to marriage rates by estimating the number of 14-19 year olds in each gender/state/year cell using the decennial Censuses, and either the official inter-censal estimates from the Census department or linear interpolation when these are not available.

We also use retrospective reports from the 5% IPUMS sample of the 1980 Census. To make the results from the two surveys strictly comparable, we collapse age-disaggregated Census data in the relevant cohorts into the same 14-19 age bin observed in the Vital Statistics, and limit the Census data to the set of states (using state of birth) and years for which there is Vital Statistics information.[7] Finally, Census marriage rates for 14-19 year olds are the rate in the IPUMS sample.

On these two separate samples, we estimate event-study regressions of the form

$$\begin{aligned} Y_{gst} =& \beta_0 + \sum_l \delta_1^l I_{gs}^t (t - k_t^* \leq -4) + \sum_l \delta_2^l I_{gs}^t (-3 \leq t - k_t^* \leq -1) \\ &+ \sum_l \delta_3^l I_{gs}^l (1 \leq t - k_t^* \leq 3) + \sum_l \delta_4^l I_{gs}^l (t - k_t^* \geq 4) + \Gamma_t + \Gamma_s + \varepsilon_{gst} \qquad (3.1) \end{aligned}$$

where Y_{gst} is the marriage rate in the 14-19 age group in year t for persons of gender g; Γ_t and Γ_s are, respectively, vectors of year and state fixed effects; and ε_{gst} is a random error term. The variable k_l^* denotes the year in which a particular "non-consent" age law l was changed (generally lowered) for persons of gender g in state s. The indicator variables I_{gs}^l in equation 3.1, denote, respectively, that the year t is four or more years before k_l^*; between 1 and 3 years before k_l^*; between 1 and 3 years

[7] Results from the Census sample using all states are similar. We have also estimated related models in which we use all states and a fully age-disaggregated sample from the Census. These results are qualitatively the same as the Census results shown here and can be seen in the appendix.

after k_l^*; or 4 or more years after k_l^*. The excluded time period in the regressions is the year in which the law changed, so the estimated coefficients δ represent how, after accounting for factors peculiar to a state and for general time effects, average marriage rates for 14-19 year olds in a given period compare to marriage rates in the year that a law changed.

Table 3.2 reports the results for regression equation 3.1 for women and men separately across the two surveys. The regressions include treatment variables for all of the different non-consent laws observed in our data, but the upper panel of the table only reports coefficient estimates associated with lowering non-consent ages from 21 to 18 for men and women, which is overwhelmingly the most common legal change in the data. In the lower panel of the table, we present the estimated change in the rate of marriage rate among 14-19 year olds between the period 1-3 years before the law to the period 1-3 years after the law ($\hat{\delta}_3 - \hat{\delta}_2$) relative to the average rate of marriage among 14-19 year olds in the state. We focus on this estimate because, in our view, only behavioral changes within the six-year period bracketing the change in a law can plausibly be attributed to the law rather than some other factor.

The results from the Vital Statistics indicate that marriage laws significantly affected young marriage rates. For young women, the estimates imply that a change in a state's female non-consent age from 21 to 18 was associated with a 16.5 percent increase in the marriage rate of young women aged 14-19. For men, the estimated effect of a 6.8 percent increase in rates of early marriage is not as large, but is not at all trivial. These estimates are precisely what we would expect if minimum age laws strongly affected young people's marriage decisions.

These patterns differ dramatically from what we find with retrospective Census data. For women, Census results indicate that marriage laws appear to have had little to no effect on the rate of early marriage. The estimates for men are positive, although they are substantially smaller than Vital Statistics results.[8]

Recall that in addition to the aggregated information used in Table 3.2, the Vital Statistics reports age disaggregated marriage data in 1950 and some years after 1968.[9] We use this information to graphically examine differences in the impact of marriage laws across states in 1950 and 1970. To understand the figures, suppose that the number of women (men) of age a wishing to marry in state s at time t is n_{st}^{a}. Let π_a be the share of people who can always get married because of some dispensation like parental or judicial consent. If $\hat{a_s}$ is the age at which persons within the state can marry without parental consent, and if marriage laws are binding, then the share of all marriages occurring in state s in year t to persons of age a is

$$f_{st}(a) = \begin{cases} \frac{\pi_a n_{st}^a}{\sum_{a<\hat{a_s}} \pi_a n_{st}^a + \sum_{a \geq \hat{a_s}} n_{st}^a} & a < \hat{a}_s \\ \frac{n_{st}^a}{\sum_{a<\hat{a_s}} \pi_a n_{st}^a + \sum_{a \geq \hat{a_s}} n_{st}^a} & a \geq \hat{a}_s. \end{cases} \tag{3.2}$$

Equation 3.2 implies that, barring the unrealistic assumption that every young person below the state minimum age can receive some exception or dispensation, the probability density function of ages at marriage should display a discrete jump or "spike" at a state's legal non-consent age. These "spikes" should occur at different ages in different states, depending upon the marriage laws in effect in the state.

Figure 3.1(a) shows the distribution of age at marriage for women in 1950 in the 15 states where the age of marriage without parental consent was 18, using Vital Statistics and Census data. The sharp spike in age 18 marriage in the Vital Statistics is entirely absent from the Census data. Census data also show higher marriage rates

[8]We have tried several alternative specifications to the ones reported, all of which lead to the same conclusion as the results in Table 3.2.

[9]The age disaggregated data from 1950 includes both first marriages and remarriages, and we (necessarily) compare this to first marriage data from the Census. Including remarriages, which are rare at younger ages, has a predictable effect of overestimating marriages at older ages. In the appendix, we perform the best available interpolation of first marriages from the data, and show that, consistent with our expectations, using the remarriage data biases our main estimates against our findings. We thus report the remarriage inclusive data in our main results because we believe it is the conservative choice.

at earlier ages, but from age 19 onward, the two series are quite similar. The six point discrepancy at age 18 marriage across the two surveys is substantial. Figure 3.1(b) shows a similar comparison for the 4 states in 1950 where the age of marriage without consent was 21 for women. Again, there is a much higher incidence of early marriage in 1950 in the Census data. Vital Statistics shows a peak at age 21 which is totally absent in the Census data. The Census and Vital Statistics also show different results for men's marriages in 1950. Figure 3.1(c) shows less of a discrepancy between data sources in states whose non-consent ages were either 18 or 20. Even so, there is a spike in the Vital Statistics at 21, and Census data show a higher prevalence of marriage at younger ages. Figure 3.1(d) plots the distribution of age at first marriage in states with a non-consent age of 21. The Census data does not show the spike at age 21 found in the Vital Statistics. Like the women's data, the Census shows a higher incidence of younger marriages among all men who marry; for ages greater than 22, the Census and Vital Statistics data are identical.[10]

In Figure 3.2 we repeat the exercise above but for marriages in 1970. Strikingly, the figures show that for both men and women, the differences between the two surveys found in 1950 marriages are either entirely absent or else are substantially reduced in 1970. In Table 3.3 we conduct a series of formal tests for the differences in the probability distributions shown in Figures 3.1 and 3.2. The results strongly confirm what is clear in the graphs: Census and Vital Statistics age distributions are statistically different from each other for marriages in 1950, but are statistically the same for marriages in 1970.[11]

[10] As noted above, we have the age of marriage data from Vital Statistics available within states for several years in the 1940s as well as in 1950. We have compared Census and Vital Statistics for these years and see the same pattern of discrepancy as are shown here for 1950.

[11] Treating these age groups as discrete bins, the data on age of marriage in Figure 3.1 may be thought of as a set of binomial distributions. We report a set of pair wise tests of the equality of proportions across legal regimes in Table 3.3, for different age bins. Since we have large samples, the binomial distribution is well approximated by the normal distribution, which implies that the differences are also approximately normal. The difference statistics have a z distribution.

3.4 Alternative Explanations for Differences across Data Sources

Before exploring our preferred explanations for the difference across the two surveys, we briefly explore whether mechanical factors, having to do with how the two data sets are constructed, may account for these discrepancies. The most obvious difference between Vital Statistics and Census marriage data is that Census information about marriage is retrospective. Might recall error associated with these reports explain the patterns documented above? An argument that recall error in Census data explains our findings would have to take one of two forms. First, it could be that when retrospectively reporting their early marriage, Census respondents systematically report having gotten married at younger ages than they actually did. But it is not obvious why the tendency to report an earlier marriage should change discontinuously at particular ages — as would have to be true for the Census and Vital Statistics to agree at older ages. Moreover, a variety of considerations suggests that if people do systematically misreport marriage in a particular direction, they likely retrospectively under report very young marriages in the Census.[12]

A second and more plausible recall error concern is that people make random errors when retrospectively reporting their age at marriage. These errors would tend to flatten the distribution of marriage ages in the Census reports relative to Vital Statistics. Since we use the 1980 Census for both 1950 and 1970 marriage, the recall problems, with the attendant flattening, would be worse for those who married in 1950 than those who married in 1970 — possibly leading to the patterns we estimate.

We address this concern in a variety of robustness exercises, the details of which can be seen in the appendix. Summarizing the results, we find that when we estimate

[12] These factors include the fact that people: (a) may choose not to report a short duration marriage from their early life, as the 1980 Census directs them to do for annulled marriages; and (b) erroneously describe the date of their *current* marriage, despite clear instructions to discuss their first marriage.

the 1950 age of marriage distribution using both the 1960 and 1970 Censuses (instead of 1980) the patterns are graphically very similar to and statistically no different from those shown in Figure 3.1. Indeed, if anything, the point estimates from these alternative exercises suggest that using the 1980 Census may actually lead to a slight underestimate of the difference in early marriages in the Census and Vital Statistics — probably because people are less likely to report early, unsuccessful marriages as time passes. Finally, we think the fact that all our analyses reveal that the largest differences between Census and Vital Statistics data occur exactly at states' non-consent ages is persuasive evidence that recall bias does not explain the differences across the surveys.

A second mechanical concern is that the Census results are faulty because we have to assume that state of birth is the same as state of marriage.[13] If people moved randomly prior to marriage, the errors we make in assigning people's state of marriage would be random, which would tend to smooth the Census data relative to Vital Statistics data. However, this should have been a larger problem in 1970 than in 1950, given higher income levels and more mobility in 1970. Instead, we find that the discrepancies between the two data series vanish rather than increase between 1950 and 1970. Note that we distinguish this point from the systematic migration to which we devote much attention below.

Two final points about the possible importance of mechanical differences between the surveys warrant some discussion. In the Census our estimates of people's age at marriage are based on their reports about the month and year they married, rather than the exact day. We assess the sensitivity of our results to any resulting

[13]We only include people in our Census sample who report being born in the U.S. This excludes U.S. marriages among those who were born outside the U.S. but immigrate prior to marriage; these immigrant marriages are included in the Vital Statistics data. We ignore this discrepancy, given the years we are focusing on are years when immigration into the U.S. is relatively low.

imprecision in the estimated age of marriage, and find that our results are robust to these tests.[14] The other issue is the possibility of attrition bias in the Census, because of death or emigration from the U.S. after marriage. Census data contain no marriage information for these missing persons. Note that the longer the retrospective period of recall, the greater the attrition this will produce in the Census marriage reports. Since we are looking at teenage marriages reported no more than 30 years later, we assume this is not a major problem in our data.

Systematic Marriage Migration

If differences in survey construction do not explain the observed differences in age of marriage across surveys, what does? We speculate that some young people systematically evaded the laws — possibly by traveling to more permissive states to marry. How important a factor is this migration in explaining the patterns shown above? In 1950 where we find the largest differences across the two data sources, available Vital Statistics data does not permit a comparison of out-of-state marriage rates across different age categories. In addition, the set of reporting states is very limited in the early years of the Vital Statistics. These shortcomings notwithstanding, we can nonetheless provide an intuitively appealing assessment of migration's influence.

To see the logic of our approach, consider a state which was a likely *destination* for marriage migration in 1950 and suppose that all cross-state marriage is motivated by legal avoidance. Census records for the state would show the age of marriage distribution of state residents, while Vitals Stats would show the age of marriage distribution for residents and in-migrants. But since all in-migrants would be at

[14] Specifically, we use electronic Vital Statistics records for the 1970s and estimate age of marriage in precisely the same way as the Census (thereby causing some imprecision in the estimated age). We compare the distribution of actual age at marriage (which is known precisely in the Vital Statistics) with the less precise Census-like calculation, and estimate virtually the same age of marriage distribution.

least as old as the non-consent age in the state, the state's Vital Statistics marriage records should show that a discretely larger fraction of marriages in the state occur at precisely the state's non-consent age than do Census data, which measure marriages only for the state's residents. Consider next a state which is a likely *source* of marriage migration. In this case, since some people younger than the state's non-consent age leave the state to marry; and since state residents who are at least as old as the state's non-consent age have no legal reason to leave the state, the distribution of age of marriage in the Vital Statistics should again show a discrete jump at the state's non-consent age relative to that in the Census. Finally, in states in which there is neither in- nor out-migration, Census and Vital Statistics the distribution of age of marriage should be the same across the Census and Vital Statistics.

One indicator of whether a state was, in 1950, a likely in-migration state; an out-migration state; or neither, is how its laws compare to those of neighboring states. In the available Vital Statistics data, there are seven states with non-consent ages of 18 which have at least one neighboring state with a higher non-consent age. These are likely in-migration states. There are eight states with female non-consent ages of 18, whose neighbors *all* also have non-consent ages of 18. These states, to a first approximation, are neither in- nor out-migration states. For men, we can identify 6 states with male non-consent ages of 21 which have at least one neighboring state with a lower non-consent age. These are likely out-migration states for men. And, we can identify 8 states with non-consent ages of 21 which are surrounded by only states whose non-consent ages are also 21. These states are likely neither in- nor out-migration states.

Figures 3.3 shows the distribution of age of marriage in these different states as estimated in the Census and Vital Statistics. The results are quite striking. Figures

3.3(a) and 3.3(b) show that women's ages of marriage in in-migration states differ sharply across Census and Vital Statistics in precisely the direction hypothesized, but there is little difference in the age of marriage distribution for non-migration states. Similarly, the two data sources reveal the sharp differences outlined above for men in out-migration states. These differences are totally absent in the non-migration states. Table 3.4 shows statistical tests that confirm the graphical evidence.

We think the foregoing evidence is strongly suggestive that there was some systematic marriage migration in 1950.[15] However, there are reasons to believe that migration is not the primary explanation for the differences in the data. In particular, observe that in Figure 3.1, the Census records, across all types of states and for both sexes, show significantly more marriages at young ages than the Vital Statistics, indicating an *aggregate* difference that cannot be explained by mobility. Additionally, we use disaggregated Vital Statistics records from the period 1968-1971, which pre-date the convergence of age of marriage laws in the mid-1970s, to measure the extent of migration. In this time period, cross-state variation in non-consent ages meant that the incentive to migrate should have been as high as in 1950, whereas higher incomes and greater mobility should have made it easier to migrate. We find that only a small share of marriages to young men and women occurred outside of their state of residence in this later period — a probable upper bound on marriage related migration in 1950.[16] In summary, some factor apart from migration must explain the 1950 disparity and 1970 convergence in Census and Vital Statistics data.

[15]The only two references to possible marriage migration that we could find in the literature are an early paper by Rosenwaike (1967) and Dahl (2005). Dahl indicates that in 1968-1969 data there is evidence of marriage migration, showing that women who marry out of state are more likely to marry in less-restrictive states. Unfortunately, the data to do this type of analysis in earlier years is not available.

[16]See the appendix for details.

Systematic Age Misrepresentation and Differences Across Data Sources

One important candidate explanation for the difference across the surveys is deliberate age misrepresentation. Specifically, we suspect that some young people, intent on not delaying their marriages to satisfy their state's minimum age rules, simply lied about their age when filling out marriage certificates. Thirty years later, facing no possible sanction from reporting the truth to the Census Bureau, they honestly reported their actual age at marriage.

The plausibility of an important role for systematic mis-representation of age depends on how easy it was to misrepresent age on a marriage certificate for this generation of young persons. Clearly, if a state required that an individual show a birth certificate, driver's license, or some other document, it would have been harder to evade the law. Misrepresentation should have been easiest when age was self-declared, with no external verification. Common experience today suggests that lying without forged documents is difficult because proof of age is routinely required for many things. But was this true in 1950? In earlier decades, a much larger share of the population did not hold a driver's license, either because they did not drive (especially younger women) or because states did not require people to carry licenses. The use of social security numbers for identification (other than employment) was relatively uncommon and many younger people did not have a social security number. Some in the population (especially black Americans in rural areas) did not even have birth certificates.

We do not have detailed information for 1950, but in 1929 the Russell Sage Foundation commissioned a document detailing marriage regulations in all states in the late 1920s (May, 1929). The first column of Table 3.5 summarizes the information in this document. In that year, 15 states indicated that information on the marriage certificate had to be certified by the oath of the parties involved, while another 13 states accepted an affidavit (essentially, a signature). Most of the remaining states did not specify that any testimony of age be offered, or indicated merely that such testimony could be requested at the clerk's discretion. In short, age was self-reported and certified by the signature or oath of the potential marriage partners.

We collected information on current requirements in all 50 states, and summarize this information in the second column of Table 3.5. By the mid-2000s, virtually all states required persons applying for licenses to provide some type of identification, usually in the form of social security numbers or birth certificates. Only a few states still have statutes that require only affidavits, and even these states appear to enforce standard practices that require marriage license applicants to show identification with proof of age.[17] We attempted to trace the statutory history of policies requiring documentary verification of age for a marriage license, but the complexity and thinness of the documentary record made this a prohibitively time intensive activity to conduct across all states. We therefore selected a set of fourteen geographically and demographically diverse states and tried to determine when these states started to require documentary proof of age to grant marriage licenses. In a few cases we could verify that such requirements were in place before or after a specified date, but could not find the year they were initially implemented. Column 3 of Table 3.5 provides this information for these 14 states. In all cases except Massachusetts, these "identifying documents" requirements appeared to have gone into effect sometime after 1960.

In short, our (admittedly fragmentary) evidence suggests that few states in 1950 appeared to require individuals to do more than swear to their stated age in order

[17]Our research assistant called a county clerk in each of these states and was told that he would need a driver's license or other identification in order to obtain a marriage certificate.

to receive a marriage license. By 1970, a growing number of states required that documentary proof of age be presented for a license. Lying about one's age to a county clerk almost surely became more difficult over the time period we study. This evidence is indirect, but when the incentives of young persons wishing to avoid state laws are taken into account it suggests that much of the difference we observe between Vital Statistics and Census is the product of young people in earlier cohorts having misrepresented their ages on official documents.

3.5 Discussion and Conclusions

The results in this paper show the massively different conclusions a researcher might be led to draw about the actual effect of marriage laws on marriage delay, depending on the data source used. In particular, we suspect that most researchers would have expected Vital Statistics data to more accurately reflect the effect of marriage laws on delayed marriage, but this expectation ignores the much greater susceptibility of administrative data to problems of systematic avoidance behavior.[18] Obviously, our results suggest important cautions for future work on questions about early marriage, especially if that work uses Vital Statistics data or exploits cross state variation in marriage rules.[19]

Our results suggest a broader lesson for *all* empirical scholars and especially those wishing to use information about laws as instrumental variables in their analyses. What we have found in the case of marriage laws is likely true for other behavior and laws: a law changes behavior among both compliers and non-compliers. Non-

[18]For instance, O'Connell (1980) benchmarks the accuracy of the retrospective Census data by comparing it to Vital Statistics data. He notes that the CPS reports a "more youthful distribution of women by age at first marriage" than the Vital Statistics in the 1940s and 1950s, while the two are more similar in later years, but he does not pursue this discrepancy.

[19]A growing literature studies various questions closely related to the issue of early marriage, including the effect of teenage childbearing on women's life outcomes, and the incidence of low education, higher levels of criminal activity, higher poverty and other negative outcomes among persons born to teens. See Ribar (1994), Hoffman (1998), Klepinger, Lundberg and Plotnick (1999), Hunt (2003), and Dahl (2005).

compliers — that is, persons whose actual behavior is not changed by the law — have an incentive to report information to administrative bodies in such a way as to suggest that it has. The researcher who naïvely assumes that the direct effects of a law can be readily estimated in administrative data, with little attention paid to the efforts to evade the law, may obtain deeply misleading estimates of the law's actual effect on the targeted behavioral change. Moreover, when avoidance itself is of interest, researchers can study it by comparing the administrative data to survey data.

Examples of the use of administrative data abound in empirical economics, and several previous authors have hinted at some of the concerns raised in this paper.[20] When, in general, is the quality of such data and the effects estimated using them likely to accurately reflect true changes in underlying behavior and when is it not? Of course, the answer will vary across contexts, but at a minimum it seems reasonable to argue that administrative data is surely better when it is directly and impartially observed (such as when height/weight information is directly measured rather than self-declared), or if information is verified as part of the data collection process (such as when birth certificates are required along with a declaration of age).

One final point illustrated by our results is that even when administrative data accurately record behavior or outcomes in a state, the likelihood that individuals may systematically move across states to avoid particular laws means that the population within which the behavior is being recorded may differ in important ways from the population researchers believe themselves to be studying. Administrative data may indeed be superior to general survey data for answering many questions, but

[20] For example, differences between reported income and earned income have long been a concern in the study of taxation. Recognizing some of the issues we discuss here, Slemrod (1992) proposes a hierarchy of responses to tax policy and emphasizes that both timing and reporting are likely to be more responsive than actual economic behavior.

whether this is actually so in a given instance is something the researcher should carefully assess by considering both types of data and by contemplating the likely contamination in each generated by agents' systematic efforts to avoid laws or rules aimed at preventing them from engaging in particular activities.

3.6 Appendix

Alternative Estimates of the Effect of Legal Restrictions on Age at Marriage Using the Census

This chapter focuses on the discrepancy between Census and Vital Statistics and the lessons to be learned for researchers about avoidance. Here, we note some evidence regarding the effect of legal restrictions on the age of marriage using the preferred data source, retrospective decennial Census data. We provide estimates of the effect of both non-consent and consent laws. The main text includes information about variation in non-consent laws. We provide analogous information here regarding consent laws in Table 3.6, which summarizes the variation across states and over time in consent laws for each sex.

Using data from the 1980 Census, we test for an effect of legal prohibition on actual age at marriage. We estimate the following equation:

$$Y^i_{gst}(a) = \beta_0 + \beta_1 P_{gst}(a) + \Gamma_b + \Gamma_s + \varepsilon^i_{gst} \tag{3.3}$$

where g indexes gender, s indexes state, and t denotes birth cohort within the Census. In equation 3.3, the vectors Γ_b and Γ_s are, respectively, birth cohort and state fixed effects; and ε^i_{gst} is a random error term. The binary outcome variable indicates whether an individual i of a given gender, state and birth cohort is ever married by age a; P_{gst} is a binary variable indicating whether, in a given year and state, the person was never able to legally marry before turning age a. So, for example, to assess the impact of non-consent laws on marriage before age 18, Y measures whether the individual was ever married by age 17, P_{gst} equals 1 if there was no time in the years before they turned 18 that the marriage laws allowed the person to legally marry. The coefficient β_1 measures how much a legal age constraint against marriage lowered the likelihood of marrying. The inclusion of state and cohort effects in equation 3.3 means that β_1 is identified from changes in marriage laws within states and across cohorts.

Table 3.7 presents the results from estimating equation 3.3. The top panel shows results for men and the bottom panel shows results for women. The first row of the top panel shows the estimate of the effect of laws that do not allow men to marry without parental consent before the age of 21 on the probability of being married by age 20. The results suggest that there is a significant negative effect of these laws on the cumulative probability of marriage at a younger age. The magnitude of the coefficient can be estimated by dividing it by the share of men married by age 20, which is 0.235 in 1970.[21] This suggests that the likelihood of being married by age 20 is reduced by 3.2 percent in a state that has a legal marriage age of 21 (without consent) versus a state with a lower legal marriage age.

We also look at the effect of changing the age of marriage with parental consent. For men, a significant number of states reduced this minimum age from 18 to 16 during the time period studied. Hence, we estimate the effect of not being able to marry without consent before age 18 on the probability of marriage by age 17. The point estimate is unexpectedly positive, but small and statistically insignificant.

The bottom panel shows similar estimates for women. When the legal age of marriage without parental consent was 20 or 21 among women, there is a statistically significant and negative effect on younger marriages. We also look at the effect of

[21] This is the average including data from all states.

age limitations on marriage with parental consent for women. The results indicate that imposing a 16 year old age of consent reduces marriage among women age 15 or younger. Relative to the mean number of marriages at age 15 or younger, our estimate suggests that the legal restriction is associated with a 15.3 percent decline in young marriages among women.[22] Our estimates suggest that age of marriage laws did impact the marriage choices of the young adult population, with larger effects on women than on men. Changes in the age of marriage without parental consent have a significant but not particularly large effect, with about a 2 to 3 percent change in the probability of marriage. This is much smaller than the effect of non-consent laws estimated in the Vital Statistics (with a different specification), as discussed in the main text.

Our Main Results are Robust to Concerns About Recall Bias

One potential reason for differences between contemporaneous Vital Statistics data and retrospective Census data is recall bias. We use the 1980 Census in our main results because a 5% sample is available in that year, whereas earlier samples are only one-fifth that size. If recall bias is a problem, it is because people's answers about their age at first marriage change depending on when you ask them. To demonstrate that recall bias is not generating our results, we show that data from the 1960 Census and the 1970 Census are statistically indistinguishable from the 1980 data we use in our main results.

Here, we replicate the measures used in our main analysis for marriages in 1950 in selected states and compare these across waves of the Census. The results, reported in Table 3.8, show that (with one exception) the age distributions are statistically

[22] In this respect, our results agree closely with Dahl (2005), who focuses only on age of marriage with consent laws for women.

indistinguishable at the 95% level.[23] There are a variety of ways to test the equivalence of two data sets. We chose to report these tests of the equivalence between waves of the Census because we think that they are most relevant to interpreting the potential importance of recall bias for the particular tests that we emphasize in the main text.

Our Main Results are Robust to Concerns About the Inclusion of Remarriages in 1950 Vital Statistics Data

The age disaggregated data in the Vital Statistics in 1950 includes both first marriages and remarriages, and we compare these, out of necessity, to first marriages from the Census in the main text of the paper. Since there are very few first marriages at the youngest ages, and proportionately more at older ages, the inclusion of remarriages will have a predictable effect on the distribution; including remarriages will exaggerate the older ages and undercount the younger ages. The inclusion of remarriages, therefore, could potentially explain part of the discrepancy between Vital Statistics and Census estimates of marriages in 1950.

In this appendix, we perform the best possible interpolation of remarriages using additional information from the 1950 Vital Statistics to demonstrate that the inclusion of remarriages, rather than explaining our findings, works against us. The 1950 Vital Statistics records indicate, in a separate table, the rates of first marriage and remarriage by state for several age categories (14-19, 20-24, etc.). We identify the set of states that match our sample as closely as possible (there is not perfect overlap in availability across the tables) and calculate average remarriage rates for the available age categories. Then, we take the same set of states and calculate the mean age within each category (e.g., the mean age of women 14-19 who marry in

[23] The pairwise comparison that we use here is the same that we describe in the text above.

1950 is not 16.5, it is closer to 19). This provides three pairs of data (each age and remarriage rate is a data point) for an appropriate set of states in 1950. We then use these three points to interpolate the estimated remarriage rate at each age.

Table 3.9 repeats the comparison emphasized in the main paper using 1950 Vital Statistics data, upon which this interpolation has been performed, to the Census estimates. On the left, the results reported in the main paper, using the original Vital Statistics data, are included for reference. On the right, we report interpolated estimates using a third degree polynomial to fit the remarriage data function exactly.[24] The interpolation has the expected effects. The adjustment closes the gap between data sources at younger ages, but only to a very small degree. The pronounced difference in Census estimates of the youngest marriages remains after the interpolation. The adjustment has a more significant effect on the spikes at middle ages in the Vital Statistics, which makes the discrepancy between data sources look much larger. The adjustment also causes Vital Statistics estimates of marriages at relatively old ages to shrink, bringing Vital Statistics closer to Census at these ages, which is consistent with our hypotheses about how the data sets should differ if individuals avoided restrictive laws.

In sum, the best available adjustment for remarriages in the 1950 data appears to strengthen the argument made in the main text. We therefore conclude that the conservative approach is to use the unadjusted data, which also has the added benefit of not requiring subjective decisions about how to perform an interpolation. We believe that the true discrepancy between first marriages reported in 1950 is probably larger than we estimate in the main text.

[24]We have experimented with other forms of interpolation and found consistent results.

Additional Information About the Quantity of Marriage Migration

The first type of legal avoidance that we analyze in the main text is the strategic movement of young people from their state of residence to nearby states with less restrictive laws, which we have called marriage migration. Here, we provide several pieces of evidence which suggest that marriage migration was relatively small in later years using disaggregated Vital Statistics data, which is available starting in 1968. Estimates of the extent of marriage migration in 1968 to 1971 provide a plausible upper bound on migration in 1950 because (a) rising incomes likely lowered the cost of traveling to marry in another state, (b) the pronounced convergence of non-consent laws had not yet taken place and (c) documentation requirements made age misrepresentation (a substitute for migration) more costly in 1968-1971, as compared to 1950. Here, we provide several pieces of evidence that migration was of modest proportions in this later period, which we interpret as suggesting that some other factor (age misrepresentation) must have driven the data discrepancy in 1950.

Table 3.10 provides evidence about the extent to which people married outside their state of residence during the years 1968-71. As the first row indicates, between 1968 and 1971, 15.7 percent of all men and 10.3 percent of all women who marry, marry outside their state of residence. Men under the age of 21 are those most likely to be affected by legal age limits. The results show that these men are somewhat less likely to marry outside their state of residence (13.6 percent), while younger women marry away from home at about the same rate as all women (10.6 percent). If we break this down by age, for younger teens we find relatively higher rates of marriage outside one's state of residence (at times exceeding 20 percent), and relatively lower rates among older teens.

Of course, people marry out of their state of residence for many reasons. How many of these young “marriage migrants” might have been seeking to avoid age of marriage laws? As Table 3.10 indicates, 66.8 percent of young men and 73.7 percent of young women who marry out of state did so in an adjacent state. Among these men, 25.8 percent of them were too young to marry without consent in their own state, but could marry legally in the adjacent state where their marriages actually took place. Among women, this rate is 19.4 percent. Since these persons were all too young to marry in their own state but could legally marry in an adjacent state, it can reasonably be argued that they were all migrating to avoid their home state’s minimum marriage age. These marriages constitute only 2.4 percent of all marriages among men under age 21 and only 1.5 percent of all marriages among women under age 21. While this is only an approximate estimate of marriage-related migration (some movers could have gone to non-adjacent states; some going to adjacent states may not have been consciously avoiding the laws, etc.), it suggests that a relatively small share of those under age 21 are likely to be migrating as a way to avoid age of marriage laws.

To further explore the importance of marriage migration, we compare migration in the period before and after age of non-consent laws converge across states. Figure 3.4 looks at these patterns. The solid dark line in Figure 3.4 shows the percentage of younger male migrants who move from more restrictive to less restrictive states, as classified by 1968 laws. The denominator is the number of men under age 21 who live in a state where the 1968 age of consent for marriage is 21 but who marry out of state; this is the number of marriage migrants who are too young to marry in historically restrictive states. The numerator is the number of these men who marry in a state where the 1968 age of consent law would have allowed them to marry legally. The ratio represents the share of younger marriage migrants who could plausibly be avoiding the law, if the 1968 laws were still in effect. We show this percentage for all years from 1968 to 1979, using the 1968 state laws to define restrictive and less restrictive states. If marriage migration is important, there should be more movement in the late 1960s between these states (when the restrictions were actually in place) than in the late 1970s (when almost all states had adopted age 18 as the legal age for marriage without parental consent). The dashed line shows the same data for women under age 21. Both of these lines decline during the period when marriage consent laws converge.

As one final check on the extent of marriage migration prior to convergence in age of marriage laws, we estimate difference-in-difference regressions, which are reported in Table 3.11. Our sample consists of all marriages among men (women) under age 25 in the periods 1968-71 and 1976-79. The dependent variable is a binary variable which denotes whether the man (woman) migrates to a state where the male (female) non-consent law is less than age 21 in 1968. We difference between the early and late period, and between men younger than age 21 and those ages 21-25. This implicitly compares changes over time (before and after the laws bind) in migration rates to states with historically lower non-consent laws among men who are of an age to be affected by these laws versus changes over time in migration rates among men who are too old to be affected.

We find that there was a statistically significant 1.8 percent higher incidence of marriages among younger men in less restrictive states in the early period than in the late period. This is quite consistent with our estimate of marriage migration in Table 3.10, suggesting a relatively small (but significant) marriage migration effect before the laws converge. Similar estimates among women find slightly larger effects. We estimate a statistically significant 2.9 percent greater rate of marriage among younger women in less restrictive states in the earlier period than in the later period.

In short, we find clear evidence of migration to states with less restrictive age of marriage laws among those who marry before age 21 in the period when there are significant cross-state differences in these laws. Marriage migration appears to be regularly used as a way to avoid state age of marriage laws. The magnitude of this effect is relatively small, however, and seems to have affected only somewhere between 1 and 3 percent of all younger marriages. Unfortunately, we can say nothing about the trend over time in legal avoidance through marriage migration before the late 1960s, but we strongly suspect the ability of teens to go out of state to avoid marriage laws would have been no greater and probably smaller in earlier years. If we take our estimate of marriage migration from the 1970s as a maximal estimate of this phenomenon in 1950, it will explain less than half of the discrepancy between Census and Vital Statistics data in 1950, suggesting that both migration and misrepresentation were occurring in this year.

3.7 Figures and Tables

Figure 3.1: Distributions of Age at Marriage from Contemporaneous vs Retrospective Reports for Individuals Marrying in 1950

(a) Women in States where Non-consent Age = 18

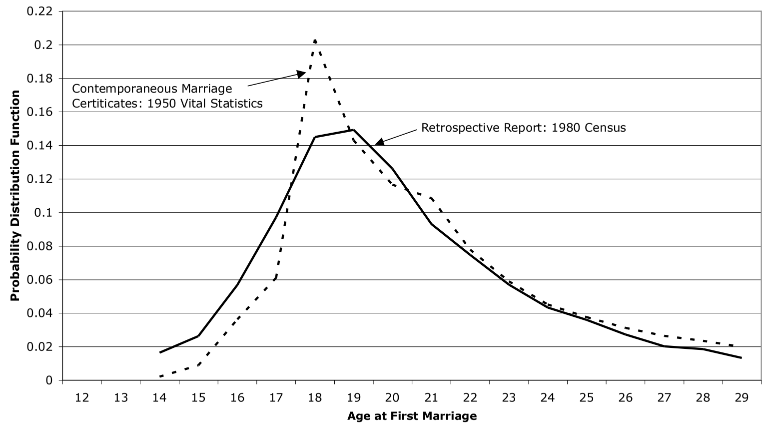

Based on the following 15 states with female non-consent age equal to 18: DE, ID, IA, KS, ME, MI, MS, MT, NH, ND, OR, SD, TN, VT and WY.

(b) Women in States where Non-consent Age = 21

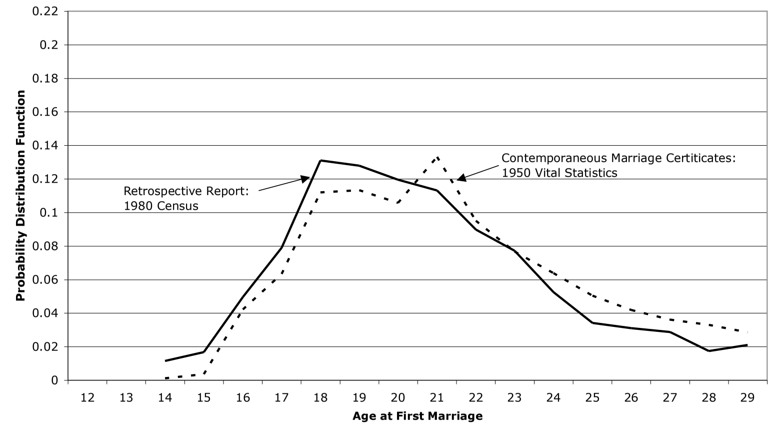

Based on the following 3 states with female non-consent age equal to 21: CT, FL and NE.

(c) Men in States where Non-consent Age = 18 or 20

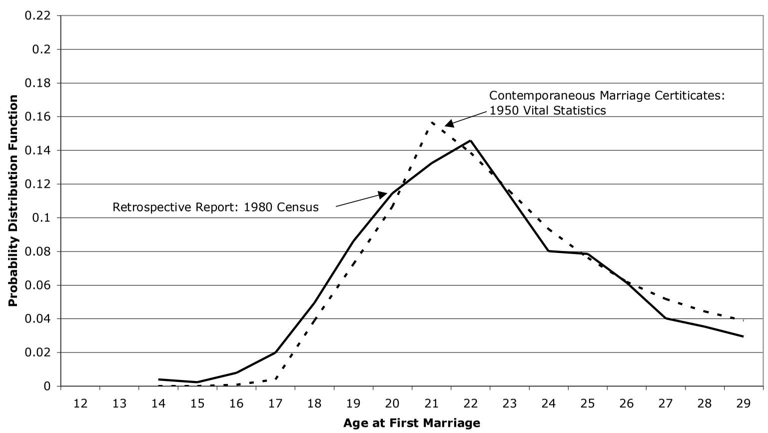

Based on the following 4 states with male non-consent age equal to 18 (ID, TN, MI) or 20 (NH).

(d) Men in States where Non-consent Age = 21

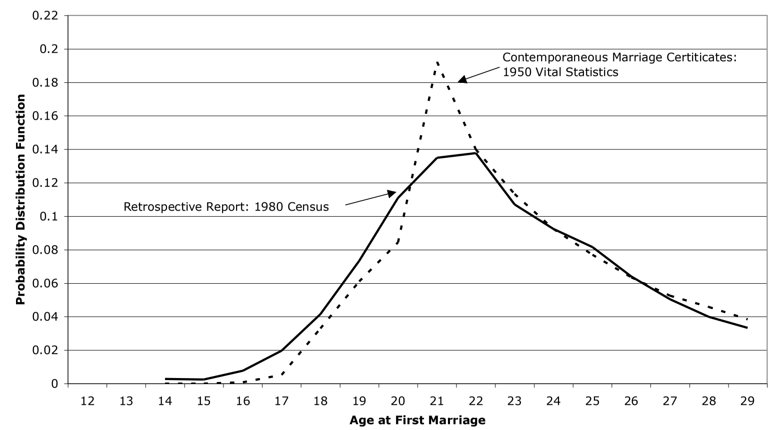

Based on the following 14 states with male non-consent age equal to 21: CT, DE, FL, IA, KS, ME, MS, MT, NE, ND, OR, SD, VT and WY.

Figure 3.2: Distributions of Age at Marriage from Contemporaneous vs Retrospective Reports for Individuals Marrying in 1970

(a) Women in States where Non-consent Age = 18

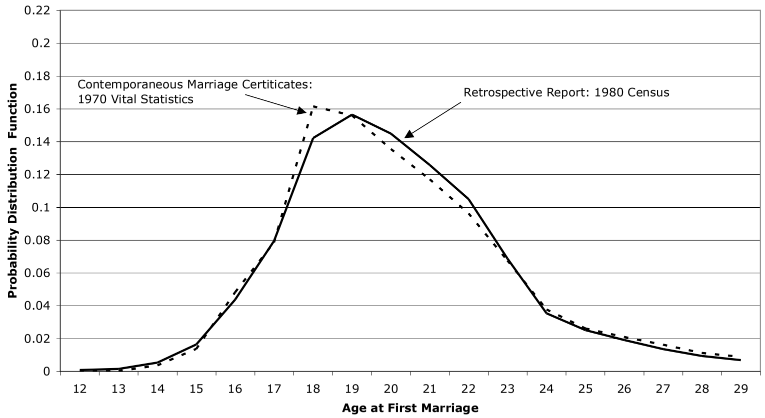

Based on the following 36 states with female non-consent age equal to 18: AK, AR, AL, CA, CO, DE, HI, IA, ID, IL, IN, KS, KY, MA, MD, ME, MI, MN, MO, MT, NC, ND, NH, NJ, NV, OH, OR, SC, SD, TN, TX, UT, VT, WA, WI and WY.

(b) Women in States where Non-consent Age = 19, 20 or 21

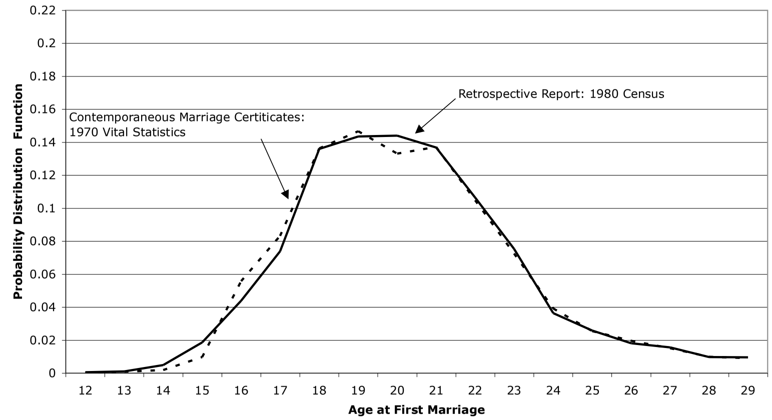

Based on the following 8 states with female non-consent age equal to 19 (GA), 20 (NE) and 21 (CT, FL, PA, RI, VA, WV).

(c) Men in States where Non-consent Age = 17, 18, 19 or 20

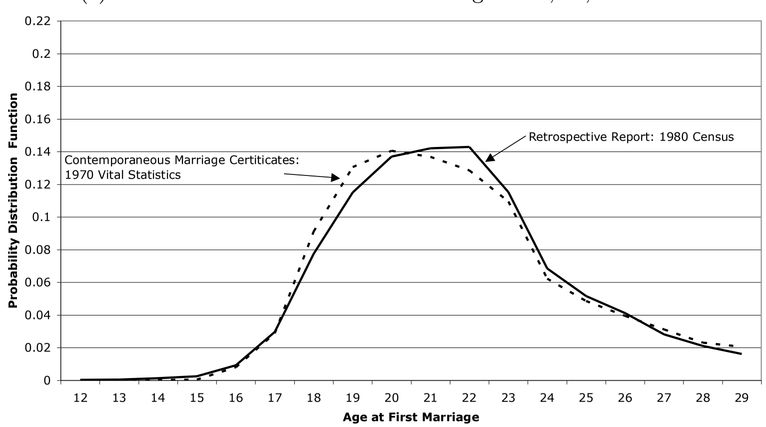

Based on the 14 states with male non-consent age equal to 17 (MS), 18 (KY, MI, NC, SC, TN, WA), 19 (AK, GA, TX) and 20 (HI, ME, NE, NH).

(d) Men in States where Non-consent Age = 21

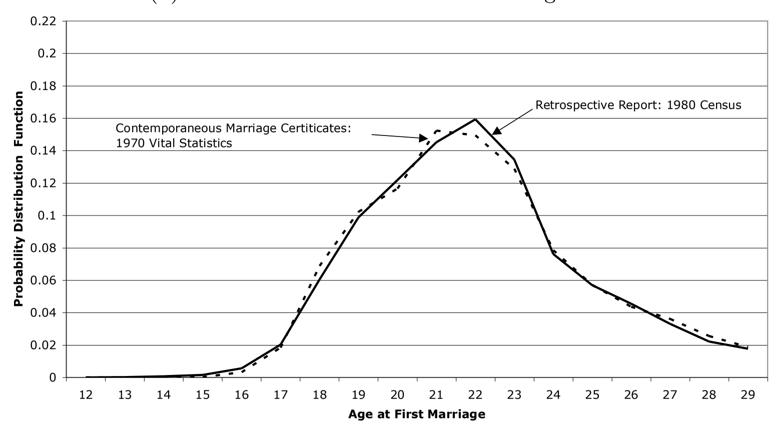

Based on the following 31 states with male non-consent age equal to 21: AL, AR, CA, CO, CT, DE, FL, IA, ID, IL, IN, MA, MD, MN, MO, MT, ND, NJ, NV, KS, OH, OR, PA, RI, SD, UT, VA, VT, WI, WV and WY.

Figure 3.3: Distributions of Age at Marriage from Contemporaneous vs Retrospective Reports for Individuals Marrying in 1950

(a) Women in States where Non-consent Age = 18 with Restrictive Law Neighbors

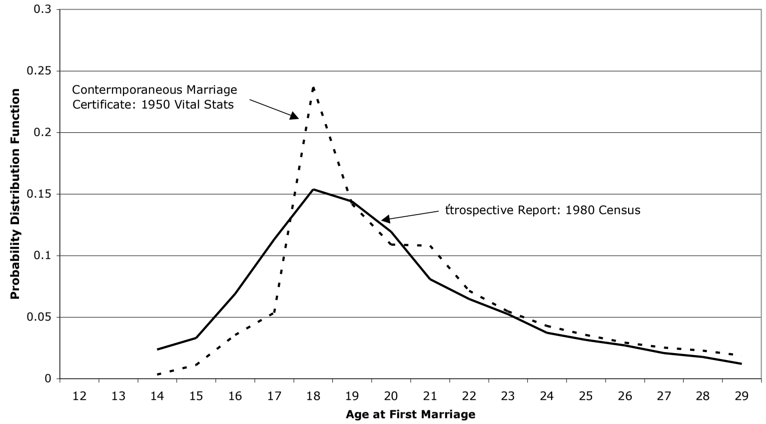

Based on the following 7 states with female non-consent age equal to 18 and with at least one adjacent state with a higher non-consent age: DE, IA, KS, MS, SD, TN and WY.

(b) Women in States where Non-consent Age = 18 and All Neighbors Have Same Law

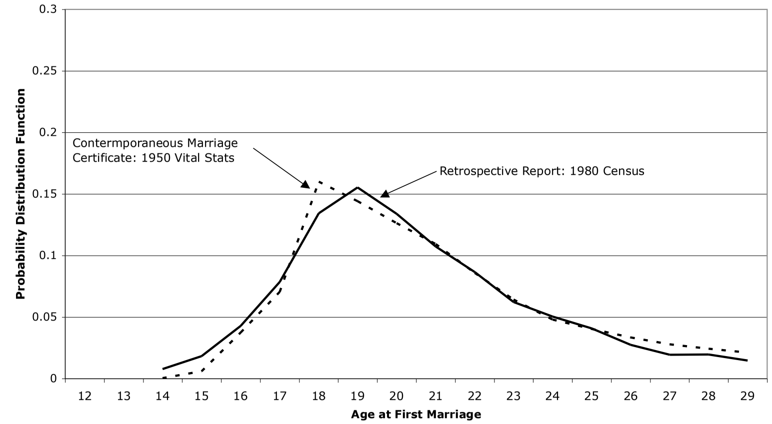

Based on the following 8 states with female non-consent age equal to 18, and for whom all adjacent states also have a non-consent age equal to 18: ID, ME, MI, MT, NH, ND, OR and VT.

(c) Men in States where Non-consent Age = 21 with Lax Law Neighbors

Based on the following 6 states with female non-consent age equal to 21 and with at least one adjacent state with a lower non-consent age: ME, MS, MT, OR, VT and WY.

(d) Men in States where Non-consent Age = 21 and All Neighbors Have Same Law

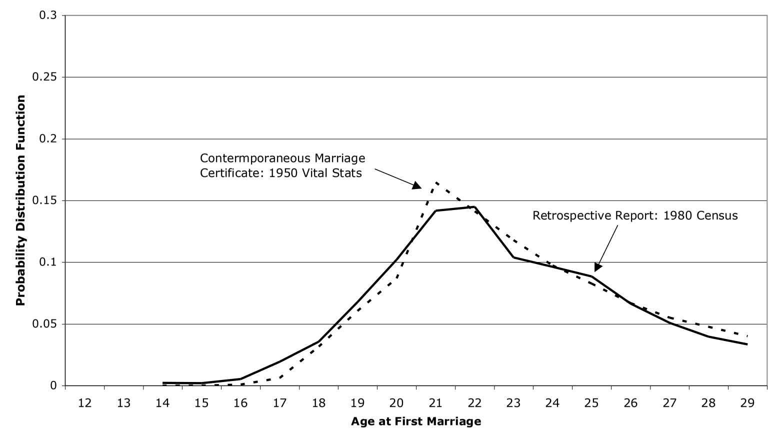

Based on the following 8 states with male non-consent age equal to 21, and for whom all adjacent states also have a non-consent age equal to 21: CT, DE, FL, IA, KS, NE, ND and SD.

Figure 3.4: Percentage of Young 'Marriage Migrants' From Restrictive States Who Move to Less Restrictive States, Classified by 1968 Laws on Age of Marriage without Consent

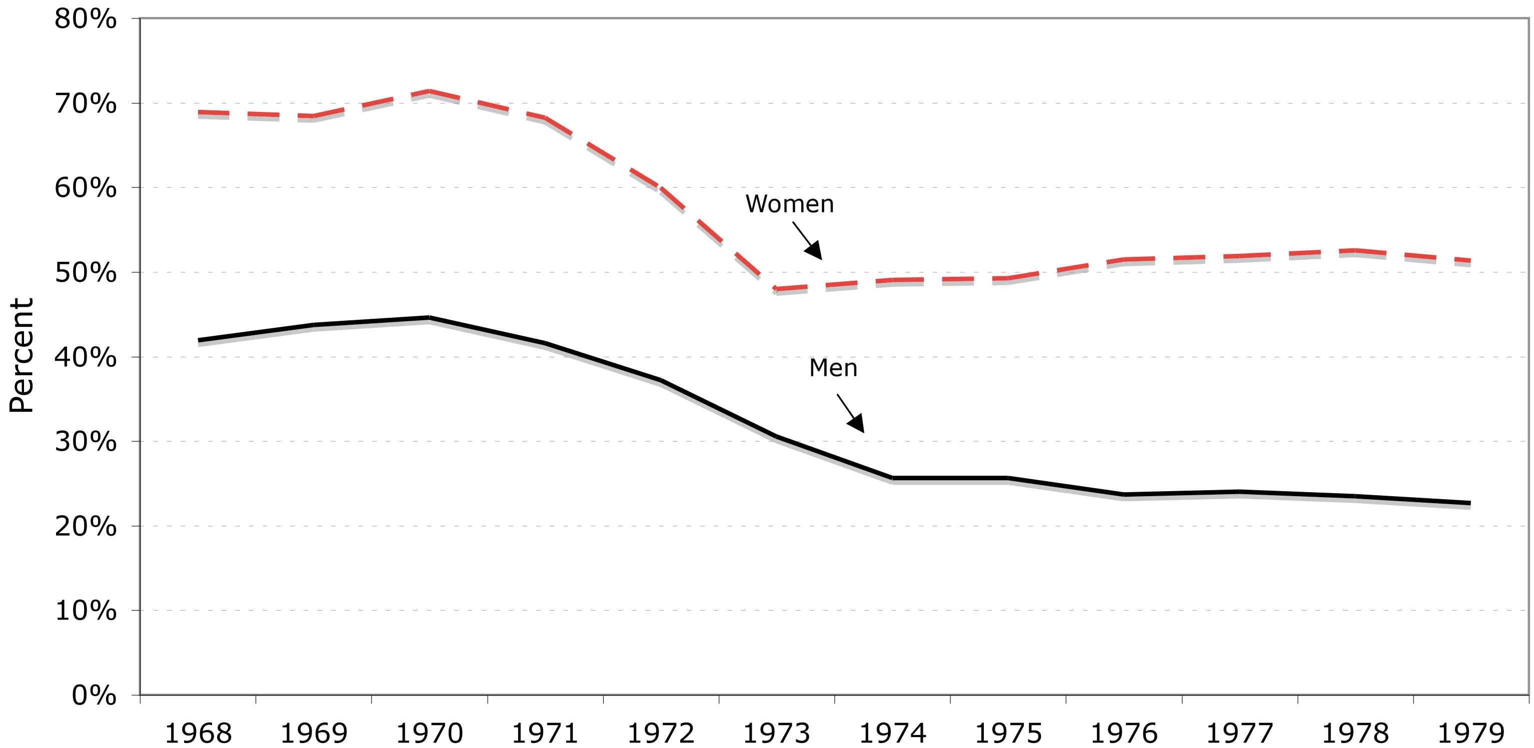

Note: The denominator for the men's (women's) line is men (women) under age 21 who live in a state where the 1968 age of consent law was 21 but who marry outside their home state; the numerator is the number of these 'marriage migrants' who marry in a state where they could legally marry based on 1968 age of consent law. The line thus shows the percent of young marriage migrants from historically more restrictive states who marry in historically less restrictive states. Data from Vital Statistics.

Table 3.1: Legal Age of Marriage Without Parental Consent in 1950, 1970 and 1980

State	1950		1970		1980	
	Women	Men	Women	Men	Women	Men
Alabama	18	21	18	21	18	18
Alaska	18	21	18	19	18	18
Arizona	18	21	18	21	18	18
Arkansas	18	21	18	21	18	21
California	18	21	18	21	18	18
Colorado	18	21	18	21	18	18
Connecticut	21	21	21	21	18	18
Delaware	18	21	18	21	18	18
Florida	21	21	21	21	18	18
Georgia	18	21	19	19	18	18
Hawaii	20	20	18	20	18	18
Idaho	18	18	18	21	18	18
Illinois	18	21	18	21	18	18
Indiana	18	21	18	21	18	18
Iowa	18	21	18	21	18	18
Kansas	18	21	18	21	18	18
Kentucky	21	21	18	18	18	18
Louisiana	21	21	21	21	18	18
Maine	18	21	18	20	18	18
Maryland	18	21	18	21	18	18
Massachusetts	18	21	18	21	18	18
Michigan	18	18	18	18	18	18
Minnesota	18	21	18	21	18	18
Mississippi	18	21	15	17	15	17
Missouri	18	21	18	21	18	18
Montana	18	21	18	21	18	18
Nebraska	21	21	20	20	19	19
Nevada	18	21	18	21	18	18
New Hampshire	18	20	18	20	18	18
New Jersey	18	21	18	21	18	18
New Mexico	18	21	18	21	18	18
New York	18	21	18	21	18	18
North Carolina	18	18	18	18	18	18
North Dakota	18	21	18	21	18	18
Ohio	18	21	18	21	18	18
Oklahoma	18	21	18	21	18	18
Oregon	18	21	18	21	18	18
Pennsylvania	21	21	21	21	18	18
Rhode Island	21	21	21	21	18	18
South Carolina	14	18	18	18	18	18
South Dakota	18	21	18	21	18	18
Tennesse	18	18	18	18	18	18
Texas	18	21	18	19	18	18
Utah	18	21	18	21	18	18
Vermont	18	21	18	21	18	18
Virginia	21	21	21	21	18	18
Washington	18	21	18	18	18	18
West Virginia	21	21	21	21	16	18
Wisconsin	18	21	18	21	18	18
Wyoming	18	21	18	21	19	19

Data on legal age requirements by state and year were collected by the authors from state statutes.

Table 3.2: Comparison of Estimated Effects of a Change in Non-Consent Age from 21 to 18 on Marriage Rates Across Data Sources

	Vital Statistics		Census	
	Women	Men	Women	Men
	Marriage Rate of 14-19 Year Olds	Marriage Rate of 14-19 Year Olds	Marriage Rate of 14-19 Year Olds	Marriage Rate of 14-19 Year Olds
Dummy for 3 years before change	-0.0040	-0.0007	-0.0005	-0.0008
	(0.0028)	(0.0010)	(0.0017)	(0.0007)
Dummy for 3 years after change	0.0065	0.0010	-0.000003	0.0002
	(0.0030)	(0.0010)	(0.0017)	(0.0007)
Dummy for 4 or more years before change	-0.0075	0.0012	-0.0004	-0.0008
	(0.0027)	(0.0012)	(0.0016)	(0.0007)
Dummy for 4 or more years after change	0.0064	0.0013	0.0002	0.0007
	(0.0027)	(0.0012)	(0.0018)	(0.0007)
State Fixed Effects	X	X	X	X
Year Fixed Effects	X	X	X	X
Years in Sample	1951-1979	1951-1979	1951-1979	1951-1979
Number of Observations	958	958	958	958
Difference (3 years after - 3 years before)	0.0105	0.0016	0.0004	0.0009
P-value	[.000]	[.071]	[.590]	[.082]
Percent Effect of Legal Change on Marriage Rate of 14-19 Year Olds	16.5%	6.8%	0.8%	4.2%

Standard errors in parentheses. We assign states based on state of birth in Census and state of marriage in Vital Statistics.
Vital Statistics estimates are based only on first marriages (unlike disaggregated Vital Statistics data studied below).
The regression also includes corresponding dummies for other non-consent legal changes. An identical set of states and years are used for both sexes and data sets.
Census estimates exclude a small percentage of observations that have imputed age at first marriage or birth place. Robustness checks indicate that this exclusion is not important for our results.

Table 3.3: Statistical Tests of the Equivalence of Marriage Proportions

Women: States with Non-Consent Age of 18 (Figures 1a and 2a)

	1950				1970			
Age Group	Vital Stats	Census	Difference	SE	Vital Stats	Census	Difference	SE
<=17	0.108	0.197	-0.088	0.004	0.145	0.148	-0.003	0.002
18	0.203	0.145	0.058	0.004	0.162	0.142	0.019	0.002
>=19	0.689	0.658	0.030	0.005	0.694	0.710	-0.016	0.002
N	201,564	8,051			121,687	46,480		

Women: States with Non-Consent Age of 21 (Figures 1b and 2b)

	1950				1970			
Age Group	Vital Stats	Census	Difference	SE	Vital Stats	Census	Difference	SE
<=20	0.441	0.535	-0.094	0.012	0.567	0.566	0.001	0.006
21	0.133	0.113	0.020	0.008	0.137	0.137	0.000	0.004
>=22	0.425	0.352	0.074	0.012	0.295	0.297	-0.002	0.005
N	45,623	1,672			24,365	11,691		

Men: States with Non-Consent Age of 18 to 20 (Figures 1c and 2c)

	1950				1970			
Age Group	Vital Stats	Census	Difference	SE	Vital Stats	Census	Difference	SE
<=17	0.005	0.034	-0.029	0.003	0.038	0.044	-0.006	0.002
18-20	0.218	0.250	-0.032	0.008	0.362	0.330	0.033	0.004
21	0.157	0.132	0.024	0.006	0.137	0.142	-0.005	0.003
>=22	0.621	0.584	0.037	0.009	0.463	0.485	-0.021	0.004
N	71,012	3,059			43,662	17,209		

Men: States with Non-Consent Age of 21 (Figures 1d and 2d)

	1950				1970			
Age Group	Vital Stats	Census	Difference	SE	Vital Stats	Census	Difference	SE
<=20	0.185	0.259	-0.074	0.006	0.310	0.309	0.000	0.003
21	0.192	0.135	0.057	0.005	0.152	0.145	0.007	0.002
>=22	0.623	0.606	0.017	0.007	0.538	0.545	-0.007	0.003
N	196,015	5,263			99,671	36,639		

Please see corresponding figures for a list of the states included in each sample.
The table reports the proportion of marriages that occur in several discrete age bins in each year for each sex.
Each proportion may be thought of as having a binomial distribution, so that a test of the equivalence of the distributions across data sets can be done by taking the difference of the proportions. This difference will have an approximate z-distribution.

Table 3.4: Statistical Tests of the Equivalence of Marriage Proportions in 1950 Vital Statistics Data Across States with Varying Opportunities for Migration

Women: States with Non-Consent Age of 18 (Figures 3a and 3b)

	States with High Migration Potential				States with Low Migration Potential			
Age Group	Vital Stats	Census	Difference	SE	Vital Stats	Census	Difference	SE
<=17	0.103	0.239	-0.135	0.007	0.114	0.147	-0.033	0.006
18-20	0.488	0.417	0.071	0.008	0.431	0.424	0.007	0.008
>=21	0.409	0.344	0.065	0.007	0.455	0.429	0.026	0.008
N	112,550	4,341			89,014	3,710		

Men: States with Non-Consent Age of 21 (Figures 3c and 3d)

	States with High Migration Potential				States with Low Migration Potential			
Age Group	Vital Stats	Census	Difference	SE	Vital Stats	Census	Difference	SE
<=17	0.009	0.042	-0.033	0.005	0.007	0.029	-0.022	0.003
18-20	0.150	0.271	-0.121	0.011	0.179	0.206	-0.026	0.007
>=21	0.841	0.687	0.155	0.012	0.814	0.766	0.048	0.007
N	65,068	1,638			87,110	3,625		

Please see corresponding figures for a list of the states included in each sample.
The table reports the proportion of marriages that occur in several discrete age bins in each year for each sex.
Each proportion may be thought of as having a binomial distribution, so that a test of the equivalence of the distributions across data sets can be done by taking the difference of the proportions. This difference will have an approximate z-distribution.
States with high migration potential are those states that have an abutting neighbor with a different legal regime. For women, high migration potential states have a neighboring state with a non-consent law of 21. For men, high migration potential states have a neighboring sate with a non-consent law of 18, 19 or 20. For both sexes, a state has a low migration potential state if all of its abutting neighbors have the same law as that state.

Table 3.5: Statutory Provisions for Proof of Age Among Those Applying for a Marriage License

State	Late 1920s	Current	Date of Statutory Change Requiring Documentation
Alabama	Not specified*	SSN	
Alaska	NA	SSN	1997
Arizona	Oath	Affidavit and SSN	
Arkansas	Affidavit*	BC	
California	Oath may be requested	Photo ID	Before 1988
Colorado	Affidavit	"Satisfactory proof" of age	
Connecticut	Oath	SSN	
Delaware	Oath	Affidavit and SSN	
Florida	Affidavit	Affidavit and SSN	
Georgia	Oath	BC, DL, or PP	1975
Hawaii	NA	SSN	
Idaho	Affidavit	BC	1967
Illinois	Affidavit	"Satisfactory proof" of age	
Indiana	Not specified	BC or DL	
Iowa	Affidavit or Certificate of age	SSN	1961
Kansas	Oath	Affidavit	
Kentucky	Not specified*	BC or DL	
Louisiana	Not specified	BC	
Maine	Not specified	Oath and SSN	
Maryland	Oath	Affidavit and SSN	
Massachusetts	Oath	BC, DL, or PP	1931
Michigan	Affidavit	BC	1968
Minnesota	Oath	Affidavit	
Mississippi	Affidavit	BC or DL	
Missouri	Not specified	SSN	
Montana	Not specified	BC	
Nebraska	Not specified	Photo ID	
Nevada	Oath may be requested	Affidavit	
New Hampshire	Not specified	BC, DL, or PP	
New Jersey	Oath	Oath by witness and SSN	
New Mexico	Not specified	Affidavit	
New York	Affidavit	BC, DL, or PP	Before 1974
North Carolina	Oath may be requested	BC	1957
North Dakota	Oath	BC	1981
Ohio	Oath	Affidavit and SSN	
Oklahoma	Evidence can be requested	BC	1961
Oregon	Affidavit	"Reasonable proof" of age	
Pennsylvania	Oath	Affidavit	
Rhode Island	Oath		
South Carolina	Affidavit	BC	1962
South Dakota	Testimony of witnesses	BC, DL, or PP	
Tennessee	Not specified*	Affidavit	
Texas	Not specified	BC, DL, or PP	Before 1997
Utah	Affidavit		
Vermont	Oath		
Virginia	Not specified		
Washington	Affidavit	Affidavit	
West Virginia	Not specified	BC or DL	
Wisconsin	Oath	BC	
Wyoming	Testimony of witnesses	Affidavit	

SSN: Social Security Number; BC: Birth Certificate; DL: Driver's License; PP: Passport
* Financial penalty specified for misinformation.
Late 1920 data from May (1929); current data and data on changes in statutes collected by authors.

Table 3.6: Legal Age of Marriage With Parental Consent in 1950, 1970 and 1980

State	1950 Women	1950 Men	1970 Women	1970 Men	1980 Women	1980 Men
Alabama	14	17	14	17	14	14
Alaska	16	18	16	18	16	16
Arizona	16	18	16	18	16	16
Arkansas	16	18	16	18	16	17
California	16	18	16	18	18	18
Colorado	16	16	16	16	16	16
Connecticut	16	16	16	16	16	16
Delaware	16	18	16	18	18	18
Florida	16	18	16	18	18	16
Georgia	14	17	16	18	17	16
Hawaii	16	18	16	18	18	16
Idaho	15	15	16	18	15	16
Illinois	16	18	16	18	18	16
Indiana	16	18	16	18	18	17
Iowa	14	16	16	18	18	18
Kansas	16	18	18	18	18	18
Kentucky	14	16	16	18	12	12
Louisiana	16	18	16	18	16	18
Maine	16	16	16	16	16	16
Maryland	16	18	16	18	16	16
Massachusetts	16	18	16	18	18	18
Michigan	16	18	16	18	16	18
Minnesota	16	18	16	18	16	16
Mississippi	12	12	15	17	15	17
Missouri	15	15	15	15	15	15
Montana	16	18	16	18	18	18
Nebraska	16	18	16	18	17	17
Nevada	16	18	16	18	16	16
New Hampshire	18	20	18	20	18	18
New Jersey	16	18	16	18	16	16
New Mexico	16	18	16	18	16	16
New York	16	16	16	16	16	16
North Carolina	16	16	16	16	16	16
North Dakota	15	18	15	18	16	16
Ohio	16	18	16	18	16	18
Oklahoma	15	18	15	18	16	16
Oregon	15	18	15	18	17	17
Pennsylvania	16	16	16	16	16	16
Rhode Island	16	18	16	18	16	18
South Carolina	14	18	14	16	14	16
South Dakota	15	18	16	18	16	16
Tennesse	12	12	12	12	16	16
Texas	14	16	14	16	14	14
Utah	14	16	14	16	14	14
Vermont	16	18	16	18	16	16
Virginia	16	18	16	18	16	16
Washington	15	12	17	17	17	17
West Virginia	16	18	16	18	16	18
Wisconsin	15	18	16	18	16	16
Wyoming	16	18	16	18	16	16

Data on legal age requirements by state and year collected by the authors from state statutes.

Table 3.7: The Effect of Marriage Laws on the Probability of Marriage before a Specified Age

Part a: Men

Marriage law	Dependent Variable: Probability of marriage by Age 20	Age 17
Never able to marry without consent before age 21	-0.0075	
	(0.0024)	
Never able to marry with consent before age 18		0.0004
		(0.0003)
Cohort fixed effects	X	X
State fixed effects	X	X
Birth cohorts included in regression	1930-1962	1930-1962
Number of observations	1,868,463	2,149,555

Share of men married by age 20 in 1970 = 0.235
Share of men married by age 17 in 1970 = 0.021

Percent effect of marriage laws limiting marriage without parental consent before age 21 (coefficient/1970 share) = -3.19%

Percent effect of marriage laws limiting marriage with parental consent before age 18 (coefficient/1970 share) = 2.03%

Part b: Women

Marriage law	Dependent Variable: Probability of marriage by Age 18	Age 15
Never able to marry without consent before age 19	-0.0088	
	(0.0030)	
Never able to marry with consent before age 16		-0.0036
		(0.0005)
Cohort fixed effects	X	X
State fixed effects	X	X
Birth cohorts included in regression	1930-1962	1930-1962
Number of observations	2,145,866	2,238,084

Share of women married by age 18 in 1970 = 0.245
Share of women married by age 15 in 1970 = 0.024

Percent effect of marriage laws limiting marriage without parental consent before age 19 (coefficient/1970 share) = -3.58%

Percent effect of marriage laws limiting marriage with parental consent before age 15 (coefficient/1970 share) = -15.32%

Standard errors in parentheses; standard errors are clustered by cohort.
Data are from the 1980 Census, including all 50 states (but not Washington, D.C.)

Table 3.8: Statistical Tests of the Equivalence of Marriage Proportions in 1950 Across Census Waves

Women: States with Non-Consent Age of 18

	1980 Census v. 1960 Census				1980 Census v. 1970 Census			
Age Group	1960 Census	1980 Census	Difference	SE	1970 Census	1980 Census	Difference	SE
<=17	0.205	0.197	0.008	0.011	0.210	0.197	0.014	0.011
18	0.136	0.145	-0.008	0.009	0.151	0.145	0.006	0.010
>=19	0.659	0.658	0.000	0.013	0.639	0.658	-0.020	0.013
N	1,708	8,051			1,697	8,051		

Women: States with Non-Consent Age of 21

	1980 Census v. 1960 Census				1980 Census v. 1970 Census			
Age Group	1960 Census	1980 Census	Difference	SE	1970 Census	1980 Census	Difference	SE
<=20	0.556	0.535	0.020	0.030	0.567	0.535	0.031	0.029
21	0.090	0.113	-0.023	0.017	0.128	0.113	0.015	0.019
>=22	0.354	0.352	0.002	0.028	0.306	0.352	-0.046	0.027
N	342	1,672			360	1,672		

Men: States with Non-Consent Age of 18 to 20

	1980 Census v. 1960 Census				1980 Census v. 1970 Census			
Age Group	1960 Census	1980 Census	Difference	SE	1970 Census	1980 Census	Difference	SE
<=17	0.005	0.034	-0.029	0.003	0.050	0.034	0.016	0.009
18-20	0.218	0.250	-0.032	0.008	0.253	0.250	0.003	0.019
21	0.157	0.132	0.024	0.006	0.146	0.132	0.013	0.015
>=22	0.621	0.584	0.037	0.009	0.552	0.584	-0.032	0.021
N	71,012	3,059			665	3,059		

Men: States with Non-Consent Age of 21

	1980 Census v. 1960 Census				1980 Census v. 1970 Census			
Age Group	1960 Census	1980 Census	Difference	SE	1970 Census	1980 Census	Difference	SE
<=20	0.185	0.259	-0.074	0.006	0.276	0.259	0.017	0.014
21	0.192	0.135	0.057	0.005	0.119	0.135	-0.016	0.011
>=22	0.623	0.606	0.017	0.007	0.605	0.606	-0.001	0.016
N	196,015	5,263			1,153	5,263		

The table reports the proportion of marriages that occur in several discrete age bins in each year for each sex.
Each proportion may be thought of as having a binomial distribution, so that a test of the equivalence of the distributions across data sets can be done by taking the difference of the proportions. This difference will have an approximate z-distribution.

Table 3.9: Statistical Tests of the Equivalence of Marriage Proportion in 1950: Comparison Between Unadjusted and Remarriage Adjusted Vital Statistics

Women: States with Non-Consent Age of 18

	Unadjusted Vital Stats				Remarriage Adjusted Vital Stats			
Age Group	Vital Stats	Census	Difference	SE	Vital Stats	Census	Difference	SE
<=17	0.108	0.197	-0.088	0.004	0.119	0.197	-0.078	0.004
18	0.203	0.145	0.058	0.004	0.221	0.145	0.077	0.004
>=19	0.689	0.658	0.030	0.005	0.660	0.658	0.001	0.005
N	201,564	8,051			201,564	8,051		

Women: States with Non-Consent Age of 21

	Unadjusted Vital Stats				Remarriage Adjusted Vital Stats			
Age Group	Vital Stats	Census	Difference	SE	Vital Stats	Census	Difference	SE
<=20	0.441	0.535	-0.094	0.012	0.493	0.535	-0.042	0.012
21	0.133	0.113	0.020	0.008	0.141	0.113	0.027	0.008
>=22	0.425	0.352	0.074	0.012	0.366	0.352	0.015	0.012
N	45,623	1,672			45,623	1,672		

Men: States with Non-Consent Age of 18 to 20

	Unadjusted Vital Stats				Remarriage Adjusted Vital Stats			
Age Group	Vital Stats	Census	Difference	SE	Vital Stats	Census	Difference	SE
<=17	0.005	0.034	-0.029	0.003	0.005	0.034	-0.029	0.003
18-20	0.218	0.250	-0.032	0.008	0.233	0.250	-0.017	0.008
21	0.157	0.132	0.024	0.006	0.165	0.132	0.032	0.006
>=22	0.621	0.584	0.037	0.009	0.597	0.584	0.013	0.009
N	71,012	3,059			71,012	3,059		

Men: States with Non-Consent Age of 21

	Unadjusted Vital Stats				Remarriage Adjusted Vital Stats			
Age Group	Vital Stats	Census	Difference	SE	Vital Stats	Census	Difference	SE
<=20	0.185	0.259	-0.074	0.006	0.198	0.259	-0.061	0.006
21	0.192	0.135	0.057	0.005	0.203	0.135	0.068	0.005
>=22	0.623	0.606	0.017	0.007	0.600	0.606	-0.007	0.007
N	196,015	5,263			196,015	5,263		

The unadjusted Vital Statistics numbers match the estimates reported in the main paper and reflect all marriages reported in Vital Statistics. The remarriage adjusted Vital Statistics numbers reflect the estimated number of first marriages in each category, based on a polynomial interpolation of remarriage rates.
The table reports the proportion of marriages that occur in several discrete age bins in each year for each sex.
Each proportion may be thought of as having a binomial distribution, so that a test of the equivalence of the distributions across data sets can be done by taking the difference of the proportions. This difference will have an approximate z-distribution.

Table 3.10: Incidence of First Marriage Outside State of Residence Among Youth

	Males	Females
Among those who marry, percent marrying outside state of residence:		
All Ages	15.7%	10.3%
Ages < 21	13.6%	10.6%
By age:		
Age 14	13.4%	21.6%
Age 15	17.4%	20.1%
Age 16	19.5%	13.2%
Age 17	20.1%	9.6%
Age 18	13.7%	11.9%
Age 19	12.8%	10.0%
Age 20	13.2%	8.7%
Among those < age 21 who marry outside state of residence:		
% marrying in an adjacent state	66.8%	73.7%
Among those < age 21 who marry in an adjacent state:		
% younger than own state's no-consent law but above no-consent law in marriage state	25.8%	19.4%
Assuming all marriages in the previous row are due to marriage avoidance, % of marriages among those < age 21 who avoid state law by marrying outside of home state	2.4%	1.5%

Note: All statistics are for first marriages. Data based on Vital Statistics records from 1968-1971, weighted by Vital Statistics sample weights. Data come from all 47 states with information reported during this time period.

Table 3.11: Difference-in-Difference Estimates of the Effect of Changing Laws on Selected Migration in Later Years

Women who resided in states whose 1968 non-consent law was 19, 20 or 21	Dummy for whether or not woman married in state whose 1968 non-consent law was under 19	Men who resided in states whose 1968 non-consent law was 21	Dummy for whether or not man married in state whose 1968 non-consent law under 21
Interaction (dummy for early year * dummy for bride under 19)	0.0293	Interaction (dummy for early year * dummy for groom under 21)	0.0185
	(0.0139)		(0.0072)
Dummy for early year	0.0332	Dummy for early year	0.0147
	(0.0164)		(0.0048)
Dummy for bride under 19	0.0096	Dummy for groom under 21	0.0008
	(0.0093)		(0.0020)
Constant	0.0383	Constant	0.0203
	(0.0166)		(0.0065)
Sample, Ages	Women (up to age 25)		Men (up to age 25)
Early Period	1968 - 1971		1968 - 1971
Late Period	1976 - 1979		1976 - 1979
Number of Observations	602,110		1,199,317

Data taken from Vital Statistics microsample; estimates are weighted by sample probabilities.
Standard errors in parentheses, clustered on source-destination pairs.
Sample includes a balanced set of source and destination states over the entire sample period.
Sample includes individuals from selected states who were married and were under age 25 at their marriage. Thus, individuals aged 20-25 (women) or 22-25 (men) form the control group.

CHAPTER IV

On the Optimal Allocation of Students and Resources in a System of Higher Education

4.1 Introduction

Approximately 77% of college students in the United States attend public institutions, where total annual expenditures now exceed $190 billion (NCES, 2005a,b). Despite the magnitude of public involvement in higher education, and despite the enormous body of research on the economics of education, economists have not established a normative model of how students and resources should be allocated in a system of public higher education. The aim of this paper is to provide a simple and tractable framework for the analysis of this issue and to derive several intuitive results regarding the optimal allocation of students and resources.

Given a distribution of student ability and a limited pool of resources for higher education, we model the social planner's decision to establish schools and populate them with students and resources. Our model is driven by two simple assumptions: (1) that there is complementarity between resources and student ability in producing educational outcomes; and (2) that there is a fixed cost to establishing a school. We show that these assumptions produce a tiered structure that sorts students by skill and results in discontinuous spending and educational output per student for essentially identical students at the margin between schools.

The existence of a fixed cost creates economies of scale, both for individual schools and for the whole system. Because they can tailor educational spending more closely to student quality, university systems serving a larger population produce more output per student, holding constant total resources per student. Improved tailoring raises aggregate social welfare, but it does not benefit all students (i.e., it is not Pareto improving). In particular, the lowest ability students at each school will lose when an additional school is introduced into the system, because they will be dropped into a lower tier. Larger systems with more schools will provide a college education to a larger fraction of the population, and they will feature a wealthier and more selective flagship school, at the optimum.

The principal contribution of this paper is that these results, which are broadly consistent with observed stylized facts, can be derived from a very simple model. Most states have a hierarchy of postsecondary institutions exhibiting markedly different levels of resources. The most obvious example is California, where the state's Master Plan for Higher Education clearly lays out a three tier structure comprised of the University of California system, the California State University system, and the California Community Colleges. Even without an explicit plan, most states have a flagship public university, some number of other four year institutions, and a system of local community colleges. Students are distributed among these schools largely on the basis of their measured academic ability, and discontinuous levels of public spending per student in each tier are strongly and positively associated with average ability. On average across the country, instructional expenditure per student in public universities that grant doctorates is more than twice that in community colleges. The difference in total expenditures directly relevant to education is higher.[1]

[1]Based on authors' calculations of 2004 data from the Integrated Postsecondary Education Data System.

Further, Winston (1999) shows that spending per student and subsidy per student are generally increasing in student quality in the U.S across all universities.

Our paper is closely related to existing work on systems of educational provision at the elementary and secondary level, as well as earlier work on competitive (as distinct from planned) systems of higher education. All of the earlier work that we are aware of contains explicit consideration of peer effects. Arnott and Rowse (1987) study an elementary and secondary system from a planning perspective similar to ours. They consider the allocation of students to various classes within a school, where students vary in ability and classrooms have a level of resources per student that applies to all students in the class. Their principal finding is that any type of partition is possible, depending on the strength of peer effects.[2]

Peer effects are a central feature of other related work. Rothschild and White (1995) analyze competitive outcomes in higher education with peer effects and demonstrate the potential for efficient private provision. Epple and Romano (1998) construct a model of private and public secondary schools in order to analyze the effects of voucher reforms. Epple, Romano and Sieg (2003) and Epple, Romano and Sieg (2006) consider a model of higher education in which universities compete on quality and university differentiation is driven by exogenous endowment differences.

The driving force in our model is complementarity in production. The notion that complementarity leads to positive assortative matching (tiers, in this case) is hardly new. This is the underlying mechanism in the marriage market model of Becker (1973), for example. In the education literature, Arnott and Rowse (1987), Bénabou (1996), Epple and Romano (1998), Epple et al. (2003) and Epple et al. (2006) all

[2]Effinger and Polborn (1999) work with a model that, on the surface, appears similar to ours. They begin, however, by assuming that there are two different schools and that some students are innately better served at the "lower" school. They then solve an allocation problem, under the assumption that attendance at one school versus the other affects wages in a market with imperfect information.

derive at least some results that resemble our tiered structure.[3] In addition to assuming complementarity between resources and ability, this earlier work considers peer effects, and, in most cases, outcomes are influenced by the distribution of income.

Our model can be interpreted as a simplified case of much of this earlier work. As is often the case, simplification yields both benefits and costs. We show that it takes only two strong (but not unreasonable) assumptions to generate an optimal system that is broadly consistent with the stylized facts of state higher education systems. By eliminating peer effects, we demonstrate that complementarity, along with fixed costs, is sufficient to make a tiered system optimal. Our simple model also allows us to consider some issues that do not appear in the prior literature. In particular, our results on resource discontinuities, our analysis of the optimal number of universities, and our discussion of the selectivity of schools are new, even as they are latent in earlier work.

We cannot, however, claim to have captured everything that might be important in our simple model. As indicated by previous research, markets, peer effects and associations between income and ability are important in higher education systems. Transportation costs (and therefore spatial considerations) and the political economy of education finance are also surely influential. Abstracting from these concerns enables us to isolate the role of complementarity and to build intuition, but it also eliminates consideration of the ways in which these factors may reinforce or counteract our findings.

We also believe that our emphasis on a planner's perspective is of value. In Epple and Romano (1998), Epple et al. (2003) and Epple et al. (2006), schools compete with each other by maximizing quality. While this is a reasonable approach, we think

[3]A strand of the literature that studies continuous optimization problems has touched on the implications of complementarity between student ability and student quality. Fernandéz and Galí (1999) is an example. That paper differs significantly from our analysis by considering a continuum of pre-existing schools of exogenous quality.

that for the case of public universities in state systems, it is more natural to consider the planner's problem of maximizing total output across a set of schools. While both approaches lead to similar mathematical results on sorting, we think that framing the problem from a planner's perspective is of heuristic value.

We build up our model incrementally throughout the remainder of the paper. In section 2, we introduce the key elements of the model. Section 3 describes the solution when the optimal number of schools is fixed. Section 4 extends the analysis to consider the optimal number of schools. Section 5 concludes. All proofs are relegated to the appendix.

4.2 A Description of the Model

We model a social planner's problem. The planner takes as given the distribution of students, the amount of resources available, the education production function, and the fixed cost of establishing a school. The planner chooses the number of universities, selects which students attend each university, and decides how many resources to give each university.

We model educational output as a function of student ability and resources per student at the student's school. We assume that ability and resources are complementary. Our assumptions about the distribution of students will be innocuous, but our assumptions about the curvature of the education production function are key to our results.

We believe a planner's problem is attractive both because it is relatively simple and because it is a good approximation to the real world, where the vast majority of students attend public institutions. In our model, students are not explicitly decision-makers. If they were, however, they would all have unanimous preferences to attend schools with higher resources per student, because this is the only dimension along which universities vary. The planner can therefore use selective admissions to produce the desired allocation. In other words, a planner with control over admissions policy can satisfy all incentive constraints.

We consider a utilitarian social welfare function, so that the social planner seeks to maximize the aggregate level of educational output. Distributional considerations could be modeled by giving the social planner a preference for equality of outcomes (e.g., a concave social welfare function), equality of expenditure (e.g., a loss function for school quality disparities), or equality of access (e.g., a loss function for selective admissions). Courant, McPherson and Resch (2006) argue that distributional considerations may affect the design of higher education because higher education may be a useful instrument for smoothing preexisting differences in welfare. We acknowledge that distributional concerns are interesting, but we focus on the utilitarian case to maintain simplicity and because we believe it is a good characterization of higher education (as opposed to primary and secondary education), where selective admissions prevail and there is typically no presumption of education for all.[4]

We do not consider peer effects. They are not needed to obtain any of our results, and omitting them simplifies the model and makes the mechanics more transparent. Previous research has included both peer effects and expenditures per student as inputs to education (Arnott and Rowse, 1987; Epple and Romano, 1998; Effinger and Polborn, 1999; Epple et al., 2003), clouding the issue of what drives the model. Our results show that complementarity is sufficient for educational sorting.

The education produced by an individual student is denoted by $h(x, r)$, where x

[4]There is a class of concave social welfare functions that we could employ without changing any qualitative results. This leads to limited additional insight, at the cost of significant additional notation. A sufficient condition for our results to hold is that the transformation $U(h(\cdot))$ be supermodular and concave, where $h(\cdot)$ is the education production function and $U(\cdot)$ is the social welfare function. The conditions on $U(\cdot)$ and $h(\cdot)$ that ensure this have no obvious economic interpretation.

is that student's ability and r is the resources per student at the student's school. Ability follows a continuous, differentiable cumulative distribution, $F(x)$, with a probability density function denoted by $f(x)$ and a finite support bounded by $\underline{x}$ and $\overline{x}$.[5]

To establish a school, the planner pays a fixed cost, θ, and then purchases the variable input into education. We assume that all students at a school receive the same resources per student. In effect, educational resources at the school are a congestible public good. For any level of total resources that a school provides, the level of resources per student depends only on the number of students in the school.

We assume that the education production function, $h(x, r)$, is continuous, twice differentiable, and increasing and concave in each argument. Logically, output should be increasing in ability and resources. We also suppose that it is concave in each element. As a normalization, we assume that students with zero resources produce zero output, and we restrict the domain of h to weakly positive values of resources. Finally, we assume that the education output function exhibits complementarity. This may also be called supermodularity, and it is equivalent to a positive cross partial derivative.

$$h = h(x, r)$$

$$h_1 > 0 \qquad h_{11} < 0$$

$$h_2 > 0 \qquad h_{22} < 0$$

$$h_{12} > 0 \quad h(x, 0) = 0$$

Only complementarity should be a controversial assumption. Complementarity means that, at any given level of resources per student, higher ability students produce more when given a marginal increase in resources. While it is not obvious that this is true in all cases, we find it to be a plausible assumption. We note also that it is pervasive

[5] A finite support is not necessary generally, but it will be required when we later assume a uniform distribution.

in the literature (Arnott and Rowse, 1987; Epple et al., 2003).

The planner will choose to set up K universities, indexed by $k = 1, \ldots, K$. The planner must pay θK in fixed costs from the total available resources T. What remains, R, the resources net of the fixed costs, is partitioned among the schools. We denote the proportion of R allocated to school k as ρ_k. The planner must also partition the distribution $F(x)$ between schools. For each value of x, the planner allocates a proportion of the distribution to each school, denoted by $p_k(x)$. The total measure of students is denoted by S. The measure of students at a school is denoted by s_k and is equal to $S \int p_k(x) f(x) dx$. Thus, the resources per student at a school, r_k, may be written as $\frac{\rho_k R}{s_k}$.

The planner simultaneously chooses the number of universities and the partition of students and resources. It is useful, however, to write the planner's problem when the number of schools is fixed as a sub-problem. We denote the global value function as V, and the value function when K is fixed as W:

$$
\begin{aligned}
V(T, S, \theta) &= \max_K W(T, S, \theta, K), \text{ where} \\
W(T, S, \theta, K) &= \max_{\{\rho_\mathbf{k}\}, \{\mathbf{p_k(x)}\}} S \int_{\underline{x}}^{\overline{x}} h\left(x, \frac{\rho_1 R}{s_1}\right) p_1(x) f(x) dx + \\
&\quad ... + S \int_{\underline{x}}^{\overline{x}} h\left(x, \frac{\rho_K R}{s_K}\right) p_K(x) f(x) dx \\
\text{s.t.} \quad & \theta K + R \leq T \\
& \sum_{k=1}^{K} \rho_k \leq 1 \\
& \rho_k \geq 0 \ \forall \ k \\
& \sum_{k=1}^{K} p_k(x) \leq 1 \ \forall \ x \\
& p_k(x) \geq 0 \ \forall \ k, x
\end{aligned}
\tag{P1}
$$

Each integral of program P1 represents a school. The output of a school is the integral of individual student outputs with resources equal to r_k, integrated over $p_k(x)f(x)$, the distribution of students assigned to the school.

The first and second constraints are the planner's budget constraint. The third disallows "negatively funded" schools. The fourth and fifth restrict the planner's partition, disallowing negative assignments, while permitting the planner to not educate some students.

In choosing the optimal number of schools, the planner balances the burden of paying the additional fixed costs for more schools against the inefficiency of sending very different types of students to the same school. The planner allocates students and resources, which implicitly sets the resources per student at each school.

First order conditions for this problem can, at least in principle, be established using variational methods. In the interest of clarity, we shall instead demonstrate that the problem can be reduced to a more tractable form.

4.3 The Optimal Allocation When the Number of Universities is Fixed

We begin by isolating the allocation decision, taking the number of schools as fixed. With a fixed number of universities, the planner's solution is a mapping from the set of students and resources into universities. One class of partitions of the type space involves grouping the highest ability types together in one school, then grouping the next highest ability types in a second school, and so on. We call this a monotonic partition.

Definition 1. *A partition is monotonic if and only if, for least and greatest elements $\underline{x}_k$ and $\overline{x}_k$ in each school, a student x is assigned to school k if and only if $\underline{x}_k \leq x \leq \overline{x}_k$.*[6]

[6]Alternatively, this could be stated as, for least and greatest elements $\underline{x}_k$ and $\overline{x}_k$ in each school, $p_k(x) = 0$ if

Any partition that results in one school having both higher and lower ability students than another school cannot be monotonic. Any partition that puts two students of the same type into different schools cannot be monotonic. Supermodularity (complementarity) of the underlying education production function is a sufficient condition to make the optimal partition monotonic.

Proposition IV.1. *If $h(x, r)$ is complementary (supermodular), then the optimal partition of students is monotonic.*[7]

Supermodularity is sufficient to generate educational sorting, even when there are no peer effects. Imagine, instead, a non-monotonic partition between two schools. The allocation can be improved by replacing a lower ability student with a higher ability student in the school with more resources per student.

Corollary 1. *In any optimal monotonic partition, any school that has higher ability students than another school will also have higher resources per student.*

Resources and ability are complements. This immediately leads to the conclusion that universities with higher ability students should have more resources per student.

Proposition IV.1 tells us the shape of the optimal solution, allowing us to rewrite program P1. The planner sets an admissions policy by determining the lowest ability type admitted to each school.[8] We denote the highest type assigned to school k by

$x < \underline{x}_k$ or $x > \overline{x}_k$ and $p_k(x) = 1$ if $\underline{x}_k < x < \overline{x}_k$.

[7]All proofs are in the appendix.

[8]There will be no gaps between the lowest type in one school and the highest type in the next school; otherwise total output could be increased by giving a higher ability student the place of a student at the lower school.

a_k, with a_0 denoting the lowest type at the lowest school.[9]

$$
\begin{aligned}
V(T,S,\theta) &= \max_K W(T,S,\theta,K), \text{ where} \\
W(T,S,\theta,K) &= \max_{\{\rho_\mathbf{k},\mathbf{a_k}\}} S\int_{a_0}^{a_1} h\left(x, \frac{\rho_1 R}{s_1}\right) f(x)dx + \\
&\quad ... + S\int_{a_{K-1}}^{a_K} h\left(x, \frac{\rho_K R}{s_K}\right) f(x)dx \\
\text{s.t.} \quad & \theta K + R \le T \qquad \text{(P2)} \\
& 0 \le a_0 \le ... \le a_k \le a_{k+1} \le ... \le a_K \le 1 \\
& \sum_{k=1}^{K} \rho_k \le 1 \\
& \rho_k \ge 0 \ \forall \ k
\end{aligned}
$$

Program P2 has one fewer constraint than program P1, and the suboptimization problem for W is a standard static optimization problem. One can easily construct a Lagrangean and characterize the first-order conditions for any given K. In principle, the planner can find the optimal allocation for each value of K that is feasible, then choose the best among these.

We find significant heuristic value in further simplifying the problem. First, we assume that the distribution of student ability is uniform on $[0,1]$. This simplifies notation, but does not substantively affect any interpretations. Second, we normalize S to 1. Third, for the remainder of this section only, we assume that the number of schools is fixed at two. Again, this substantially clarifies the tension in the model, and all of the following results are easily translatable to other values of K.

Under these additional assumptions, the planner chooses one value of ρ, the proportion of resources to be allocated to the lower school, and two cut-off conditions, the lowest ability type admitted to the lower school, a, and the lowest ability type

[9] I.e., $p_k(x) = 1$ for $x \in [a_{k-1}, a_k]$ and $p_k(x) = 0$ otherwise.

admitted to the higher school, b.

$$\max_{a,b,\rho} \; H(a,b,\rho) = \int_a^b h\left(x, \frac{\rho R}{b-a}\right) dx + \int_b^1 h\left(x, \frac{(1-\rho)R}{1-b}\right) dx$$
$$\text{s.t.} \quad 0 \le a \le b \le 1 \tag{P3}$$
$$0 \le \rho \le 1$$

First-order necessary conditions for an interior solution follow from the unconstrained optimization problem. At an interior optimum, the Lagrange multipliers on the inequality constraints are all zero. Only one constraint, $a = 0$, can ever bind at the optimum.[10]

$$H_\rho = \frac{R}{b-a}\int_a^b h_2\left(x, \frac{\rho R}{b-a}\right) dx - \frac{R}{1-b}\int_b^1 h_2\left(x, \frac{(1-\rho)R}{1-b}\right) dx = 0 \tag{4.1}$$
$$H_a = -h\left(a, \frac{\rho R}{b-a}\right) + \frac{\rho R}{(b-a)^2}\int_a^b h_2\left(x, \frac{\rho R}{b-a}\right) dx = 0 \tag{4.2}$$
$$\begin{aligned} H_b =& h\left(b, \frac{\rho R}{b-a}\right) - h\left(b, \frac{(1-\rho)R}{1-b}\right) - \frac{\rho R}{(b-a)^2}\int_a^b h_2\left(x, \frac{\rho R}{b-a}\right) dx \\ &+ \frac{(1-\rho)R}{(1-b)^2}\int_b^1 h_2\left(x, \frac{(1-\rho)R}{1-b}\right) dx = 0 \end{aligned} \tag{4.3}$$

Another way to write these first-order conditions is to substitute s_k and r_k back into the equations.

$$H_\rho = \frac{R}{s_1}\int_a^b h_2(x, r_1)\, dx - \frac{R}{s_2}\int_b^1 h_2(x, r_2)\, dx = 0 \tag{4.1b}$$
$$H_a = h(a, r_1) - \frac{r_1}{s_1}\int_a^b h_2(x, r_1)\, dx = 0 \tag{4.2b}$$
$$H_b = -h(b, r_1) + h(b, r_2) + \frac{r_1}{s_1}\int_a^b h_2(x, r_1)\, dx - \frac{r_2}{s_2}\int_b^1 h_2(x, r_2)\, dx = 0 \tag{4.3b}$$

Equation 4.1 states that the full marginal output of a dollar spent at either school must be the same in equilibrium. Educational production per student depends not

[10] The other inequality constraints, $a = b$, $b = 1$, $\rho = 0$, and $\rho = 1$ all imply that one university is empty and unused. This cannot be optimal. Whenever the fixed cost has been paid, the optimal allocation uses all available schools to tailor resources per student. In the Cobb-Douglas case, which we explore in detail below, with $\underline{x} = 0$, $a = 0$ will not bind because the lowest student produces zero.

only on the total budget of a school, but also on the number of students over which this budget is spread; the key metric is resources per student. The price of an additional unit of resources per student in a school is equal to the size of the school. Rearranging 4.1 yields:

$$\frac{\text{Price of } r_2}{\text{Price of } r_1} = \frac{s_2}{s_1} = \frac{1-b}{b-a} = \frac{\int_b^1 h_2\left(x, \frac{1-\rho}{1-b}\right) dx}{\int_a^1 h_2\left(x, \frac{\rho}{b-a}\right) dx} = \frac{\text{Marginal Effect of } r_2}{\text{Marginal Effect of } r_1}$$

Equation 4.2 describes the condition for the lowest ability person who receives education. The first term represents the contribution to education made by the marginal person when he or she is admitted. The second term represents the reduction in education of those already at the school, due to congestion, when an additional student is added. When the marginal person is added to the school, holding the school's total resources fixed, the level of resources per student falls (at a rate of $\frac{\rho R}{(b-a)^2}$), and this causes a decrease (in the amount of $h_2(x, \rho R/(b-a))$) in production for each student in the school. Thus, at the optimum, the direct contribution of the marginal student just offsets the reduction that student causes by congesting resources.

Equation 4.3 describes a similar condition for the marginal student between schools. Suppose the decision is made to send the best person from the lower school to the upper school. Their direct contribution rises by the amount $h(b, \frac{(1-\rho)R}{1-b}) - h(b, \frac{\rho R}{b-a})$, as a result of attending a school with higher resources per student (recall from corollary 1 that the higher school will have more resources per student at the optimum). This gain is exactly equal to the net crowding effect. The other two terms in 4.3 are the combined marginal benefit in the lower school of moving the student and the combined marginal loss to the students at the upper school from increased congestion.

This marginal student faces a discontinuity. He or she would produce discretely more at the upper school. Because of the complementarity between resources and student ability, the students at the better school enjoy more resources per student. The top person in the lower tier is almost exactly the same as the lowest person in the upper tier in terms of ability, but there is a discrete gap in their educational outcomes.

4.4 The Optimal Number of Universities

The above analysis characterized the optimal allocation, taking the number of schools as fixed. The planner must also choose the optimal number of schools. This analysis is less straightforward because the problem is discrete.[11] We can develop some intuition by looking at the case where $\theta = 0$, which again allows the use of standard calculus. When there is no fixed cost, the optimal solution is to tailor the resources per student to each ability type, with the resources per student rising in student ability.

Proposition IV.2. *If there are no fixed costs ($\theta = 0$), the optimal solution features a unique level of funding (a unique school) for each student ability that is funded at a positive level. The optimal amount of resources per student, $r(x)$, is an increasing function with $r'(x) = \frac{-h_{21}(x,r(x))}{h_{22}(x,r(x))} > 0$.*

When there are no fixed costs, the planner tailors education quality specifically for each ability level. The proof of proposition IV.2 solves a basic control problem. The solution demonstrates that resources will be rising in student ability, and that the rate of this increase will depend on the curvature of h. Greater complementarity increases the slope of resources as ability rises. Greater concavity in the value of resources will dampen the relationship.

[11] One may wish to appeal to discrete optimization tools such as integer programming to solve such a problem. Unfortunately, integer programming techniques, such as cutting plane methods, are not applicable because they require first solving the case where variables are not constrained to be integers. This will not work here because the objective function is not defined for non-integer values of K.

When there are fixed costs, the planner can provide only finite tailoring, which implies that almost all ability types will receive resources different from the infinite school optimum. This creates both winners and losers. Figure 4.1 shows a hypothetical resources per student function for the no fixed cost case and the same function when $K = 2$. For $K = \infty$, $r(x)$ must be increasing. The area trapped by $r(x)$ will represent the net resources available, R, when the distribution of ability is uniform and measure 1. It is possible, but not necessary, that all students receive some education in this system.

Suppose that $K = 2$ is the constrained optimum (which will be the case for some values of θ). Some measure of students at the bottom may receive no education in the constrained case. Students between a and b attend the lower tier school, and students above b attend the upper tier school. The lowest ability students at each school receive more funding than they would in the case with perfect tailoring. In general, an increase in the ability to tailor resources will increase total educational output, but it will not be Pareto improving. The lowest students at each school will lose if the number of schools increases.

When there are fixed costs, the planner must balance the benefits of tailoring against the costs of setting up new universities. A graphical representation of the planner's global choice provides further intuition. Figure 4.2 shows several hypothetical curves in total resources versus total educational production space. These curves are the $W(T, S, \theta, K)$ value functions from program P2, for several values of K, with S and θ held constant. Each curve shows how total output changes as total resources rises, holding fixed the number of schools, and the measure and distribution of students. These curves are increasing and concave. The global optimum, $V(T, S, \theta)$, is the upper envelope of these curves.

A number of comparative statics can be visualized as the expansion or compression of figure 4.2. Holding T constant, a fall in θ compresses the graph. Each W curve shifts horizontally to the left, and curves at a higher K shift more. As a consequence, systems with lower θ will be more productive.

Proposition IV.3. *The average product of resources,* $\frac{V}{T}$, *of the optimal system is rising in the measure of students when resources per student is held constant and falling in the fixed cost per school.*

The second part of proposition IV.3 is rather obvious. Average productivity rises when more resources are available for education and fewer are required for paying the fixed cost. The first part follows from the implied economies of scale. Larger systems will spread the fixed cost over more students, allowing more money to be used as an input.

Figure 4.2 also provides insight into how the optimal number of universities is chosen. For the given value of T, the planner will choose the highest curve. Each W curve begins on the T-axis at $T = \theta K$. For T above that point, W is increasing and concave. If each W satisfies the single-crossing property, then the optimal number of schools must be rising in T. Currently, we are unable to prove (or disprove) that the single-crossing property is satisfied without any additional assumptions, though our intuition is that the property will hold for a fairly broad class of functions. This property holds in the Cobb-Douglas case, and we can prove several further results with this functional form.[12]

Proposition IV.4. *If* $h(x, r) = x^{\alpha} r^{\beta}$, *with* $0 < \alpha, \beta < 1$, *then the optimal number of schools,* K^*, *is weakly rising in total resources,* T.

[12]The assumption of Cobb-Douglas can be slightly relaxed to an assumption that $h(x, r) = \alpha r^{\beta} g(x)$ without changing the proof used. We suspect that this property is true for a broader set of h functions, but the current proof uses the multiplicative separability of Cobb-Douglas, which is a relatively strong assumption.

Proposition IV.4 implies that richer systems should have more schools, thereby achieving better tailoring. Note, again, that better tailoring does not mean that all students benefit. Within each school, there are students who receive more resources than they would if perfect tailoring were feasible. Thus, there will be losers from an increase in total system resources if the addition of resources causes a rise in the number of schools. In particular, some portion of the lowest ability students at any given school will experience a decrease in educational quality when K rises. Increases in T are not, therefore, necessarily Pareto-improving, even if resources are dropped exogenously into the system.

Proposition IV.4 is closely tied to two additional comparative static results, which relate the optimal number of schools to the fixed cost of establishing a school and to the size of a system.

Proposition IV.5. *If $h(x,r) = x^{\alpha}r^{\beta}$, then the optimal number of schools, K^*, is rising in the measure of students when resources per student is held constant and falling in the fixed cost per school.*

The intuition behind proposition IV.5 is clear from figure 4.2. A reduction in θ shifts all the curves to the left. Curves with higher K values shift more. Thus, the diagram is contracted, and the cut-off points all move to the left. Holding T constant, K^* must weakly rise as a result. Raising the measure of students, while keeping total resources per student constant, has the same effect on the cut-off points.

University systems that serve a larger population, therefore, should be superior in several ways. Even if they are not richer per student, they should have more universities. They should do a better job of tailoring educational quality to students, and they should produce more per dollar of resources and more per student.

Larger university systems will also serve a greater fraction of the ability distribution, and they will feature more selective flagship universities.

Conjecture 1. *If $h(x,r) = x^{\alpha}r^{\beta}$, then the selectivity of the top university will be rising in K, holding R constant.*

This remains a conjecture, because we have been unable to prove this for the general case, but there are reasons to believe that the claim is true. First, it is clearly true in the limit. As the number of schools approaches infinity, the top school will become arbitrarily selective. Second, we have investigated this claim numerically, assuming that ability is distributed uniformly on $[0,1]$. Our assumptions require that $0 < \alpha < 1$ and $0 < \beta < 1$. We performed a grid search over these intervals with a .1 width, for $K = 1$ through $K = 5$.[13] For each α, β pair we numerically located the optimal cut points for each value of K and checked that the selectivity of the top school is rising in K. This procedure revealed no counterexamples. These examples, of course, do not prove the conjecture. Note, however, that even if there is some set of values for α, β and K that generate a counterexample to the claim, the predicted relationship will likely still emerge in the real world. A similar result about the low end of the distribution holds in simulations. Systems with more universities will serve a larger fraction of the distribution.[14]

We selected additional parameter values and extended the search up to $K = 10$. Two examples are provided in figure 4.3 for illustration. When $K = 1$, the lowest type admitted to the top school is the same as the lowest type admitted to the bottom school. As K rises, the lowest type admitted to the top school (A_K) also rises. The limiting argument suggests that as $K \to \infty, A_K \to 1$, giving a sense

[13] The numerically estimated solution is very sensitive to starting values when the parameters are near 0 or 1, which necessitates an extra layer of search. We checked many values close to 0 and 1, and we performed a finer grid search over the middle of the parameter space (from .2 to .8) where starting value sensitivity is reduced.

[14] An alternative approach is to calibrate the model. We prefer the grid search primarily because we do not believe there is a reliable way to calibrate α and β. Since we find no contradictions to our claim throughout the entire parameter space, we feel that the grid search is more comprehensive than a calibrated example, which would focus on a single pair of α and β.

of how these curves would project forward. A corresponding shape exists for the lowest student admitted to the bottom school, with this value approaching 0 in the limit. The corresponding curves have a similar shape for each of the large number of parameter value pairs that we have examined.

These results suggest that schools in states with a larger number of universities should have more selective flagships. Descriptive data from the Integrated Postsecondary Education Data System on university characteristics in 2001 support this hypothesis.[15] Figure 4.4 plots the 75th percentile of the combined SAT scores for students at each state's flagship university against the number of two- and four-year public universities in that state. It is clear from the graph that states which have more institutions (better tailoring) feature a more selective flagship university. Figure 4.5 plots the percentage of applicants admitted by the flagship university against the number of public universities in that state. The data again suggest that larger university systems have more selective flagships.

4.5 Extensions and Conclusions

The purpose of this paper is to provide a framework for analyzing the optimal allocation of students and resources within a system of higher education. Our hope is that future research will enrich the model and test its implications.

Our model does not include tuition.[16] At the optimum, the social cost of moving a student from their assigned university to a better one is the change in their educational output minus the net crowding effect. If individuals experience a private gain from educational output, there will be some measure of students at any university for whom the private gain from a university upgrade will outweigh the total social

[15] These data are available at http://nces.ed.gov/ipeds/.

[16] The existing literature primarily considers tuition policies that enable ability screening for schools maximizing quality (e.g., Epple and Romano (1998); Epple et al. (2003, 2006)).

cost. This suggests that there are gains to be made by allowing students to pay for an upgrade.

Similarly, students might be willing to pay a premium to attend a university out of state. If the social planner's objective function includes only the education of in-state residents, the optimal tuition policy will be to admit out-of-state students as long as their tuition, at the margin, exceeds the current resources per student at a school. In general, the introduction of tuition policy will make the total amount of system resources an endogenous variable.

Our model also makes empirical predictions about the relationship between the number of universities in a system and the selectivity of its flagship university and about the effects of introducing an additional university to a system. When new universities are introduced, our model suggests that some types of students will experience a reduction in educational quality, while others will experience an increase. At the same time, overall educational output, and the marginal value of additional revenue, should rise with the introduction of a new university. Our hope is that future research will utilize variation in the fixed cost (e.g., land grants, changes in federal support) and the size of the population (e.g., migration, the Baby-Boom) to test and further refine our findings.

In this paper, we focused on a deliberately simple model. Nevertheless, it captures a number of key features about the provision of public higher education. In particular, our model offers a normative explanation for a tiered university system, within which higher ability students receive more resources. It highlights the tradeoff inherent in tailoring education quality to student ability. It also provides a model for understanding the optimal number of universities in a system, and makes suggestions about how university systems should vary.

4.6 Appendix

Proposition IV.1. *If $h(x,r)$ is complementary (supermodular), then the optimal partition of students is monotonic.*

Proof: Fix the number of schools and the resources in each school.[17] Suppose the optimal partition is not monotonic. Call the two schools that violate monotonicity 1 and 2, and, without loss of generality, assume 1 has the higher resources per student. If monotonicity fails, then $\exists\ y \in 1 < z \in 2$. The proposed solution produces $h(y,r_1)+h(z,r_2)$. Switching the two students yields $h(y,r_2)+h(z,r_1)$. And,

$$h(y,r_2)+h(z,r_1) > h(y,r_1)+h(z,r_2) \Leftrightarrow$$

$$h(z,r_1)-h(z,r_2) > h(y,r_1)-h(y,r_2)$$

which is a definition of supermodularity, since $z > y$ and $r_1 > r_2$. **QED.**

Corollary 1. *In any optimal monotonic partition, any school that has higher ability students than another school will also have higher resources per student.*

Proof: Suppose that the optimal partition is monotonic, with students in 1 being higher ability than students in 2, but with $r_2 > r_1$. By supermodularity, swapping any two students between schools raises output. **QED.**

Proposition IV.2. *If there are no fixed costs ($\theta = 0$), the optimal solution features a unique level of funding (a unique school) for each student ability that is funded at a positive level. The optimal amount of resources per student, $r(x)$, is an increasing function with $r'(x) = \frac{-h_{21}(x,r(x))}{h_{22}(x,r(x))} > 0$.*

Proof: Part I: Suppose that there is a school with positive resources and two or more distinct ability types. Then there exists some school with both y and z with

[17] Clearly, if any two schools provide the same resources per student, there would be economies of scale gains to merging the schools. Thus, we can proceed as if the resources per student differs at each school.

$z > y$. Let s denote the total measure of students at the school. Without loss of generality, suppose that the measure of each type is the same, $s/2$. Then, the output of the proposed optimum can be written:

$$V = \frac{1}{s}h\left(y, \frac{\rho R - \varepsilon}{s}\right) + \frac{1}{s}h\left(z, \frac{\rho R + \varepsilon}{s}\right) \tag{4.4}$$

where $\varepsilon = 0$ and total funding at the school is ρR. We show that it is optimal to set $\varepsilon > 0$, which is equivalent to separating y and z into two different schools:

$$\frac{\partial V}{\partial \varepsilon} = -h_2\left(y, \frac{\rho R - \varepsilon}{s}\right) + h_2\left(z, \frac{\rho R + \varepsilon}{s}\right) > 0$$

The last inequality follows directly from supermodularity and contradicts the optimality of the proposed solution.

Part II: The infinite school problem may be written as a control problem:

$$\begin{aligned} \max_{r(x)} & \int_0^1 h(x, r(x))dx \\ \text{s.t.} & \int_0^1 r(x)dx = T \end{aligned}$$

It can easily be shown that the Hamiltonian leads to a degenerate solution with $h_2(x, r(x)) = \beta \in \mathbb{R}$. This is an implicit function, and the conditions of the implicit function theorem are satisfied because $h_{22}(\cdot) \neq 0$. The implicit function theorem yields the final result, which we sign from our assumption of concavity and complementarity:

$$r'(x) = \frac{-h_{21}(x, r(x))}{h_{22}(x, r(x))} > 0.$$

QED.

Proposition IV.3. *The average product of resources, $\frac{V}{T}$, of the optimal system is rising in the measure of students when resources per student is held constant and falling in the fixed cost per school.*

Proof: Fix T and S. Lowering θ relaxes the resource constraint for any value of K. $V(T,S,\theta)$ must therefore rise. Since T is fixed, $\frac{V}{T}$ must rise.

A rise in the measure of students when resources per student is held constant means that S rises but $\frac{T}{S}$ is fixed. We can write this as a γ proportional change, with $\gamma > 1$. Denote the value function as $W(T,S,\theta,K)$. For any value of K, output per dollar of total resources can be written:

$$\begin{aligned}
\frac{W(\gamma T,\gamma S,\theta,K)}{\gamma T} &= \max_{\rho_k,a} \frac{\gamma S}{\gamma T}\sum_{k=1}^{K}\int_{a_k-1}^{a_k} h\left(x,\frac{\rho_k(\gamma T-\theta K)}{\gamma s_k}\right) f(x)dx \\
&= \max_{\rho_k,a} \frac{S}{T}\sum_{k=1}^{K}\int_{a_k-1}^{a_k} h\left(x,\frac{\rho_k(\gamma T-\theta K)}{\gamma s_k}\right) f(x)dx \\
&> \max_{\rho_k,a} \frac{S}{T}\sum_{k=1}^{K}\int_{a_k-1}^{a_k} h\left(x,\frac{\rho_k(T-\theta K)}{s_k}\right) f(x)dx \\
&= \frac{W(T,S,\theta,K)}{T}
\end{aligned}$$

The second equality uses the envelope theorem, which tells us that the optimal cut-points will not change when the parameters are varied in small amounts. The inequality uses the fact that $\frac{\partial}{\partial\gamma}\frac{\rho_k(\gamma T-\theta K)}{\gamma s_k} = \frac{\rho_k \theta K}{\gamma s_k} > 0$. Since this is true of any K, it must be true for the optimal K. **QED.**

Proposition IV.4. *If* $h(x,r) = x^{\alpha} r^{\beta}$*, with* $0 < \alpha, \beta < 1$*, then the optimal number of schools,* K^**, is weakly rising in total resources,* T*.*

Proof: Define $W(T,S,\theta,K)$ to be the constrained solution, when the number of universities is fixed at K, and denote the derivative of $W(\cdot)$ with respect to T by

W_T. In the Cobb-Douglas case, we can relate W and W_T:

$$\begin{aligned}
W(T,S,\theta,K) &= \max_{\rho_k,a} S \sum_{k=1}^{K} \int_{a_{k-1}}^{a_k} x^{\alpha} \left(\frac{\rho_k R}{(a_k - a_{k-1})S} \right)^{\beta} dx \\
&= \max_{\rho_k,a} S \left(\frac{R}{S} \right)^{\beta} \sum_{k=1}^{K} \int_{a_{k-1}}^{a_k} x^{\alpha} \left(\frac{\rho_k}{a_k - a_{k-1}} \right)^{\beta} dx \\
&= \max_{\rho_k,a} S^{1-\beta} (T - \theta K)^{\beta} \sum_{k=1}^{K} \int_{a_{k-1}}^{a_k} x^{\alpha} \left(\frac{\rho_k}{a_k - a_{k-1}} \right)^{\beta} dx \\
W_T(T,S,\theta,K) &= \max_{\rho_k,a} \beta S^{1-\beta} (T - \theta K)^{\beta - 1} \sum_{k=1}^{K} \int_{a_{k-1}}^{a_k} x^{\alpha} \left(\frac{\rho_k}{a_k - a_{k-1}} \right)^{\beta} dx
\end{aligned}$$

The vector of ρ and a values that maximize the objective function will also maximize the marginal value of resources, W_T, since W_T is an affine transformation of W: $W_T(T,S,\theta,K) = \frac{\beta}{T-\theta K} W(T,S,\theta,K)$. Now, consider any two numbers of universities, with $K > \hat{K}$:

$$\begin{aligned}
W_T(T,S,\theta,K) &= \frac{\beta}{T-\theta K} W(T,S,\theta,K) \\
&> \frac{\beta}{T-\theta K} W(T,S,\theta,\hat{K}) \\
&> \frac{\beta}{T-\theta \hat{K}} W(T,S,\theta,\hat{K}) \\
&= W_T(T,S,\theta,\hat{K})
\end{aligned}$$

Since the derivative of W with respect to T is higher the higher is K, the family of W functions will satisfy the single crossing property in the $T-W$ plane. For each $K > \hat{K}$ and $T > \hat{T}$, $W(\hat{T},S,\theta,K) > W(\hat{T},S,\theta,\hat{K}) \Rightarrow W(T,S,\theta,K) > W(T,S,\theta,\hat{K})$. As is illustrated in figure 4.2, W will be zero up until $T = \theta K$. So, W functions with higher K values start rising at a later point.

$V(T,S,\theta)$, the optimum when K is a choice variable, is the upper envelope of the family of W functions in figure 4.2. Because of the single crossing property, this upper envelope must lie on a W for a weakly higher K as T rises. Thus, K^* is rising in T. **QED.**

Proposition IV.5. *If $h(x,r) = x^\alpha r^\beta$, then the optimal number of schools, K^*, is rising in the measure of students when resources per student is held constant and falling in the fixed cost per school.*

Proof: Define $T^*(i,j)$ as the T that solves $W(T^*(i,j), S, \theta, i) = W(T^*(i,j), S, \theta, j)$, as in figure 4.2. Define $Q(K) = \max_{\rho_k, a} \sum_{k=1}^{K} \int_{a_{k-1}}^{a_k} x^\alpha \left(\frac{\rho_k}{a_k - a_{k-1}}\right)^\beta dx$. For the Cobb-Douglas case, as is shown in the proof of proposition IV.4, $W(T, S, \theta, K) = S^{1-\beta}(T - \theta K)^\beta Q(K)$. Therefore, $T^*(i,j)$, which will be unique if it exists by the single-crossing property from proposition IV.4, solves

$$S^{1-\beta}(T^*(i,j) - \theta i)^\beta Q(i) = S^{1-\beta}(T^*(i,j) - \theta j)^\beta Q(j) \tag{4.5}$$

Without loss of generality, assume $i < j$. Totally differentiate equation 4.5 with respect to T and θ:

$$(dT^*(i,j) - id\theta)\beta S^{1-\beta}(T^*(i,j) - \theta i)^{\beta-1} Q(i)$$
$$= (dT^*(i,j) - jd\theta) S^{1-\beta}(T^*(i,j) - \theta j)^\beta Q(j)$$

Rearrange:

$$\begin{aligned}
\frac{dT^*(i,j)}{d\theta} &= \frac{j\beta S^{1-\beta}(T^*(i,j) - \theta j)^\beta Q(j) - i\beta S^{1-\beta}(T^*(i,j) - \theta i)^{\beta-1} Q(i)}{\beta S^{1-\beta}(T^*(i,j) - \theta j)^\beta Q(j) - \beta S^{1-\beta}(T^*(i,j) - \theta i)^{\beta-1} Q(i)} \\
&= \frac{jW_T(T^*(i,j), S, \theta, j) - iW_T(T^*(i,j), S, \theta, i)}{W_T(T^*(i,j), S, \theta, j) - W_T(T^*(i,j), S, \theta, i)} \\
&> 0
\end{aligned}$$

The last two steps follow directly from the analysis in the proof of proposition IV.4. This shows that all cut-off points rise when θ rises. This implies that, when T is held constant, a rise in θ must weakly decrease the number of cut-off points passed with total resources T, which is equivalent to a weakly falling K^*.

For the size result, note that $W(\gamma T, \gamma S, \theta, K) = (\gamma S)^{1-\beta}(\gamma T - \gamma \frac{\theta}{\gamma} K)^\beta Q(K) = \gamma W(T, S, \frac{\theta}{\gamma}, K)$. Since γ does not depend on any of the parameters, the cut-off points for the $\gamma W(T, S, \frac{\theta}{\gamma}, K)$ system are equivalent to the cut-off points for the $W(T, S, \frac{\theta}{\gamma}, K)$ system. Thus, an increase in the measure of students, holding constant total resources per student, which is equivalent to choosing $\gamma > 1$, is equivalent in its effect on K^* to a reduction of θ to $\frac{\theta}{\gamma}$. Since K^* was proved above to be weakly falling in θ, it must be weakly rising in the measure of students. **QED.**

4.7 Figures

Figure 4.1: Resources per Student as a Function of Student Ability

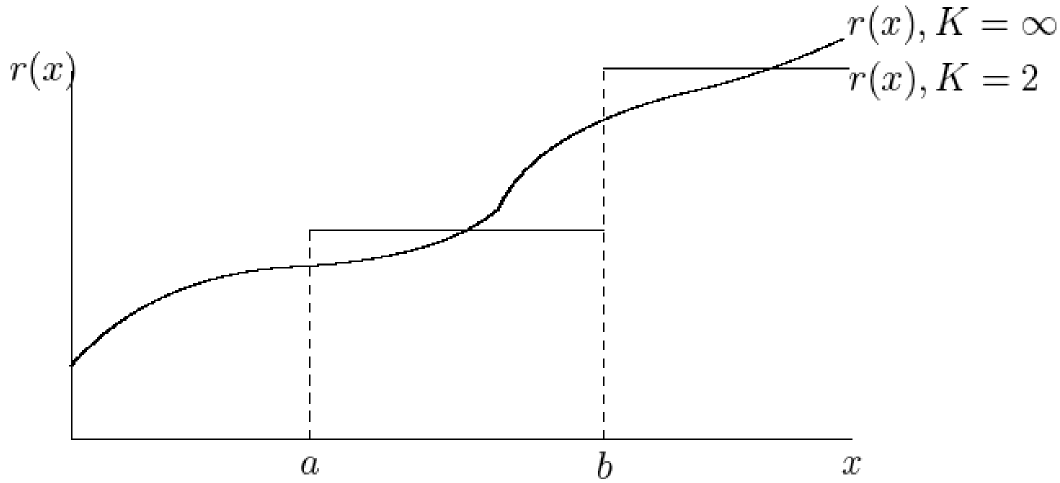

Figure 4.2: Total Output versus Total Resources for Several Values of K

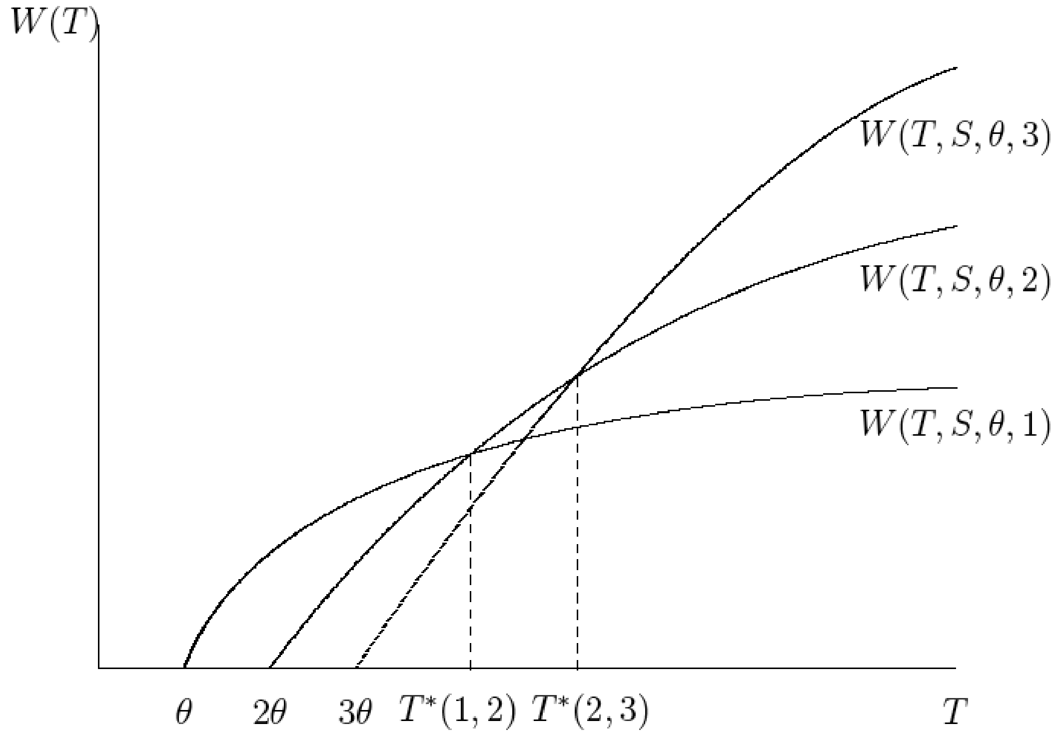

Figure 4.3: Numerically Estimated Optimal Values of A_0 and A_K for Selected α and β

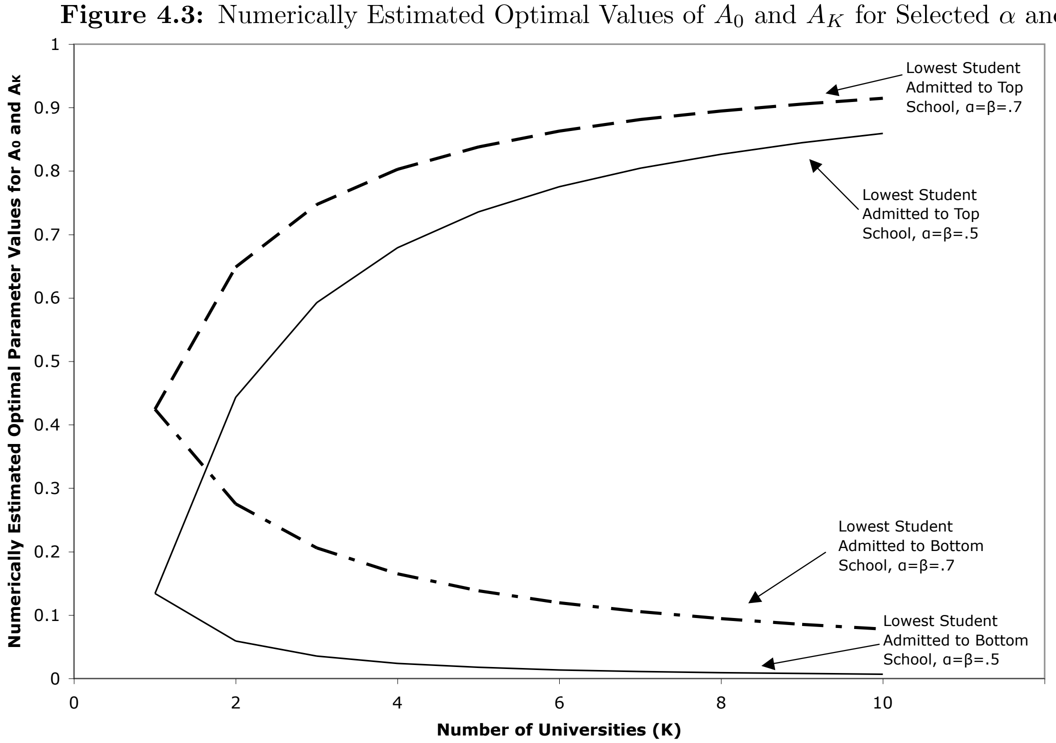

Figure 4.4: Flagship SAT Scores versus Number of Universities in State

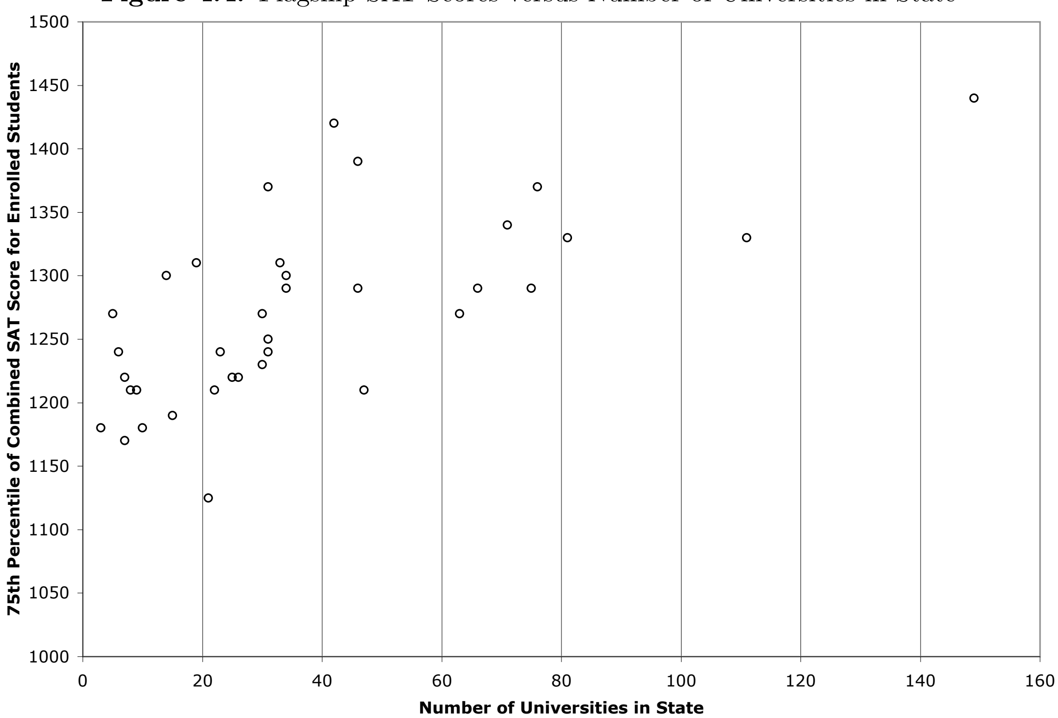

Data include all 36 states with available SAT scores from the 2001 Integrated Postsecondary Education Data System. A regression of 75th Percentile SAT scores at the state flagship on the number of colleges in a state yields the following estimates: SAT = 1213 (14) + 1.52 (.29) * Number of Schools + error, where standard errors are in parentheses.

Figure 4.5: Flagship Admissions Acceptance versus Number of Universities in State

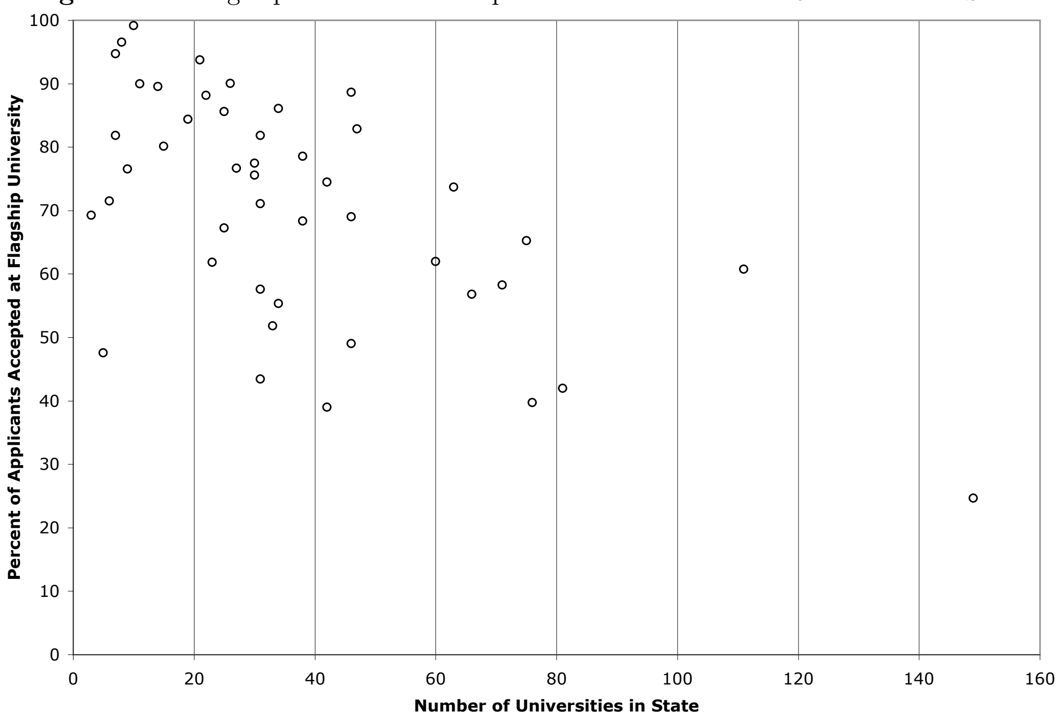

Data include all 45 states with available admissions data from the 2001 Integrated Postsecondary Education Data System. A regression of fraction admitted to the flagship on the number of colleges in a state yields the following estimates: Fraction Admitted = 107 (15) * -1.00 (.20) Number of Schools + error, where standard errors are in parentheses.

CHAPTER V

Conclusion

As stated in the introduction, the three essays in this dissertation are independent of one another. All provide an analysis of public policy by applying the tools of public economics. The substantive conclusions of each essay, however, are largely unrelated. The first essay demonstrates that tax incentives favoring hybrid vehicles have largely benefited consumers, not manufacturers. This finding suggests broader lessons for the economics of taxation, including situations where the economic incidence of a tax depends on statutory incidence. The second essay cautions future researchers against an instinctive faith in administrative data and demonstrates the importance of assessing how laws may be avoided by economic actors. The third essay provides a normative justification for a correlation between student ability and resources in higher education systems, which is consistent with what we observe in reality. It also provides basic intuition regarding the trade-offs in resource allocation in managing a system of schools. My hope is that all or some of this will aid in the future design and evaluation of public policy.

Bibliography

Alson, Jeff, Benjamin Ellies, and David Ganss, "Interim Report: New Powertrain Technologies and Their Projected Costs," October 2005. U.S. Environmental Protection Agency, EPA420-R-05-012.

American Council for an Energy-Efficient Economy, "New Light-Duty Hybrid and Diesel Tax Credits for Consumers Estimated by ACEEE," August 12, 2005. http://www.aceee.org/press/0508hybridtaxcr.htm.

Arnott, Richard and John Rowse, "Peer Group Effects and Educational Attainment," *Journal of Public Economics*, April 1987, *32* (3), 287–305.

Asplund, Marcus, Richard Friberg, and Fredrik Wilander, "Demand and Distance: Evidence on Cross-Border Shopping," *Journal of Public Economics*, February 2007, *91* (1-2), 141–157.

Associated Press, "Tax Credits for Some Hybrid Cars Halved," *USA Today*, September 21, 2006.

Automotive News, "Customer and Dealer Incentives," *Automotive News*, 2005 - 2007, *Various Issues* (December 12, 2005 to May 7, 2007).

Bailey, David, "Toyota Sees Big '07 U.S. Sales Jump for Prius," *Reuters*, February 7, 2007.

Becker, Gary S., "A Theory of Marriage: Part 1," *Journal of Political Economy*, July - August 1973, *91* (4), 813–846.

Bénabou, Roland, "Equity and Efficiency in Human Capital Investment: The Local Connection," *The Review of Economic Studies*, April 1996, *63* (2), 237–264.

Benton, Joe, "Prius Supplies Will Remain Tight Through 2006," *ConsumerAffairs.com*, December 16, 2005.

Blank, Rebecca M., Christine C. George, and Rebecca M. London, "State Abortion Rates: The Impact of Policies, Providers, Politics, Demographics, and Economic Environment," *Journal of Health Economics*, October 1996, *15* (5), 513–553.

Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer, "$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions," *American Economic Review*, September 2006, *96* (4), 1253–1270.

Busse, Meghan R., Duncan Simester, and Florian Zettelmeyer, "'The Best Price You'll Ever Get': The 2005 Employee Discount Pricing Promotions in the U.S. Automobile Industry," May 2007. NBER Working Paper no. 13140.

Campbell, Margaret C., "Perceptions of Price Unfairness: Antecedents and Consequences," *Journal of Marketing Research*, May 1999, *36* (2), 187–199.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," August 2007. NBER Working Paper no. 13330.

CNW Research, "Hybrid Vehicles – Private Personal Use," 2007. Document Number 481.

Copeland, Adam, Wendy Dunn, and George Hall, "Prices, Production and Inventories over the Automotive Model Year," March 2005. Federal Reserve Board: Finance and Economics Dicussion Series: Working Paper 2005-25.

Corrado, Carol, Wendy Dunn, and Maria Otoo, "Incentives and Prices for Motor Vehicles: What Has Been Happening in Recent Years?," January 2006. Federal Reserve Board: Finance and Economics Dicussion Series: Working Paper 2006-09.

Courant, Paul N., Michael McPherson, and Alexandra M. Resch, "The Public Role in Higher Education," *National Tax Journal*, June 2006, *59* (2), 291–318.

Dahl, Gordon B., "Early Teen Marriage and Future Poverty," April 2005. NBER Working Paper no. 11328.

Effinger, Matthias R. and Mattias K. Polborn, "A Model of Vertically Differentiated Education," *Journal of Economics*, February 1999, *69* (1), 53–69.

Environmental Protection Agency, *Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2005*, U.S. Environmental Protection Agency, 2007. EPA430-R-07-002.

Epple, Dennis and Richard E. Romano, "Competition between Private and Public Schools, Vouchers, and Peer-Group Effects," *American Economic Review*, March 1998, *88* (1), 33–62.

_ , Richard Romano, and Holger Sieg, "Peer Effects, Financial Aid and Selection of Students into Colleges and Universities: An Empirical Analysis," *Journal of Applied Econometrics*, September - October 2003, *18* (5), 501–525.

_ , _ , and _ , "Admission, Tuition, and Financial Aid Policies in the Market for Higher Education," *Econometrica*, July 2006, *74* (4), 885–928.

Fehr, Ernst and Simon Gächter, "Fairness and Retaliation: The Economics of Reciprocity," *The Journal of Economic Perspectives*, Summer 2000, *14* (3), 159–181.

Fernandéz, Raquel and Jordi Galí, "To Each According to...? Markets, Tournaments, and the Matching Problem with Borrowing Constraints," *The Review of Economic Studies*, October 1999, *66* (4), 799–824.

Frey, Bruno S. and Werner W. Pommerehne, "On the Fairness of Pricing – An Empirical Survey Among the General Population," *Journal of Economic Behavior and Organization*, April 1993, *20* (3), 295 – 307.

Fullerton, Don and Gilbert E. Metcalf, "Tax Incidence," in A.J. Auerbach and M. Feldstein, eds., *Handbook of Public Economics, Volume 4*, Elsevier Science B.V., 2002, pp. 1787 – 1872.

Gallagher, Kelly Sims and Erich J. Muehlegger, "Giving Green to Get Green? The Effect of Incentives and Ideology on Hybrid Vehicle Adoption," September 2007. Manuscript, Harvard University.

Healey, James R., "Toyota Goal: 1 Million Hybrids A Year," *USA Today*, June 11, 2007.

Hoffman, Saul B., "Teenage Childbearing Isn't So Bad After All...Or Is It? A Review of the New Literature," *Family Planning Perspectives*, July - August 1998, *30* (5), 236–239+243.

Horowitz, Joel L. and Charles F. Manski, "Identification and Robustness with Contaminated and Corrupted Data," *Econometrica*, March 1995, *63* (2), 281–302.

_ **and** _ , "What Can Be Learned about Population Parameters when the Data Are Contaminated?," in G.S. Maddala and C.R. Rao, eds., *Handbook of Statistics, Volume 15*, Elsevier Science B.V., 1997, pp. 439–466.

Hotz, V. Joseph, Charles H. Mullin, and Seth G. Sanders, "Bounding Causal Effects Using Data From a Contaminated Natural Experiment: Analysing the Effects of Teenage Childbearing," *Reveiw of Economic Studies*, October 1997, *64* (4), 575–603.

Hunt, Jennifer, "Teen Births Keep American Crime High," March 2003. NBER Working Paper no. 9632.

Inman, J. Jeffrey, Leigh McAlister, and Wayne D. Hoyer, "Promotion Signal: Proxy for a Price Cut?," *The Journal of Consumer Research*, June 1990, *17* (1), 74 – 81.

Internal Revenue Service, "Revenue Procedure 2002-42," 2002.

_ , "Some Toyota and Lexus Vehicles Certified for the New Energy Tax Credit," April 7, 2006. IR-2006-57.

_ , "Toyota Hybrid Begins Phaseout on October 1," September 20, 2006. IR-2006-145.

_ , *2005 Individual Income Tax Returns* Publication 1034, Washington, D.C.: U.S. Government Printing Office, 2007.

_ , "Alternative Minimum Tax by State," January 2007. Statistics of Income Division, Individual Master File System.

Jappelli, Tullio and Luigi Pistaferri, "Do People Respond to Tax Incentives? An Analysis of the Italian Reform of the Deductibility of Home Mortgage Interests," *European Economic Review*, February 2007, *51* (2), 247–271.

Kahn, James A., "Gasoline Prices and the Used Automobile Market: A Rational Expectations Asset Price Approach," *Quarterly Journal of Economics*, May 1986, *101* (2), 323–340.

Kahn, Matthew E., "Do Greens Drive Hummers or Hybrids? Environmental Ideology as a Determinant of Consumer Choice," *Journal of Environmental Economics and Management*, 2007, *Forthcoming.*

_ , "Who Are Early Adopters of New Energy Efficient Products? Evidence from Hybrid Vehicle Registrations," May 2007. Manuscript, UCLA.

Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, March 1979, *47* (2), 263–292.

_ , Jack L. Knetsch, and Richard Thaler, "Fairness as a Constraint on Profit Seeking: Entitlements in the Market," *American Economic Review*, September 1986, *76* (4), 728–741.

Kenkel, Donald S., "Drinking, Driving, and Deterrence: The Effectiveness and Social Costs of Alternative Policies," *Journal of Law and Economics*, October 1993, *36* (2), 877–913.

Klepinger, Daniel, Shelly Lundberg, and Robert Plotnick, "How Does Adolescent Fertility Affect the Human Capital and Wages of Young Women?," *Journal of Human Resources*, Summer 1999, *34* (3), 421–448.

Kopczuk, Wojciech and Joel Slemrod, "Putting Firms into Optimal Tax Theory," *American Economic Review Papers and Proceedings*, 2006, *2* (96), 130–134.

Lazzari, Salvatore, "Tax Credits for Hybrid Vehicles," December 20, 2006. Congressional Research Service: CRS Report RS22558.

Lee, David S., "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects," October 2005. NBER Working Paper no. 11721.

Leonhardt, David, "U.S. Hybrids Get More Miles Per Congress," *The New York Times*, June 21, 2006.

Linn, Joshua and Thomas Klier, "Gasoline Prices and the Demand for New Vehicles: Evidence from Monthly Sales Data," 2007. Manuscript, University of Illinois at Chicago.

Lovenheim, Michael F., "How Far to the Border?: The Extent and Impact of Cross-Border Casual Cigarette Smuggling," *National Tax Journal*, March 2008, *61* (1), 7–33.

McManus, Walter, "The Link Between Gasoline Prices and Vehicle Sales," *Business Economics*, January 2007, *42* (1), 53–60.

Moorthy, Sridhar, Brian T. Ratchford, and Debabrata Talukdar, "Consumer Information Search Revisited: Theory and Empirical Analysis," *The Journal of Consumer Research*, March 1997, *23* (4), 263 – 277.

NCES, *Digest of Education Statistics 2004*, Government Printing Office, U.S. Department of Education, National Center for Education Statistics, 2005.

_ , *Projections of Education Statistics to 2014*, Government Printing Office, U.S. Department of Education, National Center for Education Statistics, 2005.

O'Connell, Martin, "Comparative Estimates of Teenage Illegitimacy in the United States, 1940-44 ato 1970-74," *Demography*, February 1980, *17* (1), 13–23.

Ratchford, Brian T. and Narasimhan Srinivasan, "An Empirical Investigation of Returns to Search," *Marketing Science*, Winter 1993, *12* (1), 73–87.

_ , Myung-Soo Lee, and Debabrata Talukdar, "The Impact of the Internet on Information Search for Automobiles," *Journal of Marketing Research*, May 2003, *40* (2), 193–209.

Rechtin, Mark, "Spiffs Help Lift Mainstream Sales of Prius," *Automotive News*, April 9, 2007.

Ribar, David, "Teenage Fertility and High School Completion," *Review of Economics and Statistics*, August 1994, *76* (3), 413–424.

Romich, Jennifer L. and Thomas Weisner, "How Families View and Use the EITC: Advance Payments versus Lump Sum Delivery," *National Tax Journal*, December 2000, *53* (4), 1245–1265.

Rosenwaike, Ira, "Parental Consent Age as a Factor in State Variation in Bride's Age at Marriage," *Journal of Marriage and the Family*, August 1967, *29* (3), 452–455.

Rothschild, Michael and Lawrence J. White, "The Analytics of the Pricing of Higher Education and Other Services in Which the Customers Are Inputs," *Journal of Political Economy*, June 1995, *103* (3), 573–586.

Rowley, Ian, "Toyota Winning the Hybrid Race," *Business Week*, April 3, 2006.

Scott Morton, Fiona, Florian Zettelmeyer, and Jorge Silva-Risso, "Internet Car Retailing," *The Journal of Industrial Economics*, December 2001, *49* (4), 501–519.

Slemrod, Joel, "Do Taxes Matter? Lessons from the 1980's," *American Economic Review*, May 1992, *82* (2), 250–256.

_ , "A General Model of the Behavioral Response to Taxation," *International Tax and Public Finance*, March 2001, *8* (2), 119–128.

_ , “Show Me the Money: The Economics of Tax Remittance,” January 2007. Manuscript, University of Michigan.

Smeeding, Timothy M., Katherine Ross, and Michael O’Connor, “The EITC: Expectation, Knowledge, Use, and Economic and Social Mobility,” *National Tax Journal*, December 2000, *43* (4), 1187–1209.

Time Magazine, “Packing the Price,” *Time*, April 14, 1958.

Truett, Richard, “Suppliers Seek to Satisfy Hybrid Battery Demand,” *Automotive News*, January 24, 2005.

Union of Concerned Scientists, “State and Federal Hybrid Incentives,” September 2007. http://go.ucsusa.org/hybridcenter/incentives.cfm.

Weitzman, Martin L., “Optimal Search for the Best Alternative,” *Econometrica*, May 1979, *47* (3), 641–654.

West, Sarah, “The Effect of Gasoline Prices on the Demand for Sport Utility Vehicles,” 2007. Manuscript, Macalester College.

Winston, Gordon C., “Subsidies, Hierarchy and Peers: The Awkward Economics of Higher Education,” *Journal of Economic Perspectives*, Winter 1999, *13* (1), 13–36.

Woodyard, Chris, “Prius Finally Available Without a Wait,” *USA Today*, November 11, 2006.