Car notches: Strategic automaker responses to fuel economy policy[☆]

James M. Sallee [a,*], Joel Slemrod [b]

[a] The Harris School, University of Chicago, 1155 East 60th Street, Chicago, IL 60637, United States
[b] Ross School of Business and Department of Economics, University of Michigan, 701 Tappan Street, Ann Arbor, MI 48109, United States

ARTICLE INFO

ABSTRACT

Notches – where marginal changes in behavior lead to discrete changes in a tax or subsidy – figure prominently in many policies. In this paper, we analyze notches in fuel economy policies, which aim to reduce negative externalities associated with fuel consumption. We provide evidence that automakers respond to notches in the Gas Guzzler Tax and mandatory fuel economy labels by precisely manipulating fuel economy ratings so as to just qualify for more favorable treatment. We then describe the welfare consequences of this behavior and derive a welfare summary statistic applicable to many contexts. In brief, notches are an inefficient substitute for smooth policies because they create marginal incentives that vary among decision makers and induce some individual actions that have negative net social benefits.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Notches – where small changes in behavior lead to large changes in outcomes – figure prominently in many policies. Intuitively, these notches must have efficiency implications because they create large, capriciously varying local incentives to make small changes in behavior. Yet, despite their ubiquity, little economic research has tried to establish the consequences of using notched policies in place of smooth ones. The goal of this paper is to determine the consequences of using a notched schedule (a step function) of taxes or subsidies to approximate a smooth schedule, when the tax is intended to correct for an externality. In such cases, for those located near a notch, a marginal change in behavior leads to a large, discrete private gain but a small, marginal social reward. This disconnect between private and social returns implies that behavioral responses to notches erode the intended welfare benefits of corrective taxes.

Our specific objects of inquiry are notches in policies intended to encourage the production and use of fuel-efficient vehicles, which we call *car notches*. Key aspects of U.S. policy toward motor vehicle fuel economy feature such notches. The Gas Guzzler Tax, which penalizes vehicles with poor fuel economy, is a notched tax. For example, a car with a

14.5 miles-per-gallon (mpg) rating is subject to a $4500 tax, while a car with a 14.4 mpg rating is subject to a $5400 tax, so that a tax increase of $900 is triggered by a decrease of just 0.1 mpg. Likewise, mandatory fuel economy labels, which are intended to give consumers information about a vehicle's efficiency, are reported as integers. Thus, a small change in the underlying continuous measure of fuel economy corresponds to a discretely different published mpg for vehicles close to a rounding cutoff. Our aim is to document that automakers respond to these local incentives by fine tuning fuel economy and to estimate the welfare implications of their behavioral response.

In spite of a folk wisdom that they are sub-optimal,[1] policy notches are ubiquitous. Tax notches include the U.S. Saver's Credit, which provides a tax credit equal to a percentage of contributions to retirement savings accounts, where the credit rate is a notched function of adjusted gross income (Ramnath, 2009), and the U.K. Family Credit, which applies only to families that have one adult working 16 or more hours per week (Blundell, 2000). Energy efficiency labels on buildings and appliances are notches (Sallee, 2011a), and most countries have notched vehicle policies.[2] Many social programs have eligibility

---

[1] Only Blinder and Rosen (1985) have risen to the defense of notches. In the context of encouraging the consumption of a socially desirable good, they show via simulation that when general non-linear Pigouvian subsidies cannot be used in a model with multiple heterogeneous individuals, a single-notch program can improve efficiency because a notch is a non-linear subsidy that can, depending on the distribution of tastes, economize on the revenue loss from subsidizing inframarginal consumption. But, it is an inefficient tool when used to approximate a smooth schedule, as in our context.

[2] To list just a few, Sweden offers rebates of up to €4000 when purchasing a hybrid car with CO_2 emissions below 120 grams per kilometer. Malaysia has notched registration and fuel subsidies that depend on engine displacement. In the UK, the annual vehicle tax is a notched schedule determined by CO_2 emissions. Before 2001, it was a notched schedule in engine size, which may explain why there are so many cars with 1399 cc engines. Chinese fuel economy standards are a notched function of vehicle weight.

notches in the form of age requirements or means tests. Notches in time (a policy change takes effect on a specific date) and space (a policy changes at the border of a county, state or country) are present in most policies.[3]

In this paper, we document that automakers do indeed respond to notches in the Gas Guzzler Tax and fuel economy labels. We begin our empirical analysis by examining the distribution of fuel economy ratings for vehicles subject to the Gas Guzzler Tax. There, we find evidence of strategic bunching: there are a statistically significant number of "extra" vehicles with fuel economy ratings just on the tax-preferred side of notches. We show that vehicles with higher sales volume are more likely to end up on the tax-favorable side of a notch and that there is dramatic bunching around the mpg of the final notch in the tax, above which automobiles have a zero tax liability. We also present falsification tests that show that similar bunching does not exist in related distributions where tax implications are absent.

Next, we show that bunching occurs around the critical rounding values for mandatory fuel economy labels, the values of which must be rounded to integers. We call such mandatory rounding, which coarsens the information that producers provide to consumers, a *presentation notch*. Taken together, this evidence strongly suggests that automakers respond to local incentives created by notches by strategically altering fuel economy ratings by small amounts.

To understand the welfare implications of this behavior, we develop a framework that assumes that the notched fuel economy tax was meant to approximate a first-best linear Pigouvian subsidy for fuel economy. We show that, under certain conditions, there is a simple statistic, which we call the *average effective tax rate* around a notch, that summarizes how local manipulation distorts the intended effects of the tax. We demonstrate that this value, in conjunction with ex post aggregate data, determines a measure of the local welfare cost of using a notched policy, as compared to the first-best linear schedule.

Our conclusions regarding welfare can be summarized as follows. Notches cause automakers to re-engineer vehicles so as to just qualify for a more advantageous policy treatment. This has some social benefit, but because notches provide different marginal incentives to different decision makers, depending on their proximity to a notch, they are not efficient. A notched policy does not guarantee that behavioral changes are concentrated among those with the highest net benefits. Moreover, because automakers with vehicles located close to a notch face private benefits from marginal improvements that far exceed social gains, the *net* welfare implication of notch-induced manipulation may be negative. For the case of the Gas Guzzler Tax, we find that the tax is less efficient than a smooth tax would be, and we estimate that the local manipulation in response to notches has a net negative social impact.

Notched policies may also trigger the introduction of qualitatively new products, which Kleven and Slemrod (2009) call *tax-driven product innovation*. They note that in Indonesia, the preferential tax treatment of motorcycles relative to autos led to the creation of a new type of motorcycle with three wheels and long benches at the back seating up to eight passengers—car-like but not so car-like as to be taxed as a car. When Chile imposed much higher taxes on cars than on panel trucks, manufacturers soon offered a redesigned "panel truck" that featured glass windows instead of panels and upholstered seats in the back.[4] We provide suggestive evidence of something similar in the United States, where vehicles categorized as light trucks, as opposed to passenger cars, are exempt from the Gas Guzzler Tax.

The aim of this paper is to document the efficiency costs associated with notches, which appear to be under-appreciated by policy makers. As Slemrod (2010) discusses, notches might be justified by benefits including administrative simplicity or enhanced salience to consumers. Our conclusions suggest that there are real efficiency costs from notches in vehicle policies. Unless benefits that outweigh these costs can be documented, policy makers should consider removing notches from many programs.

We begin in Section 2 by describing fuel economy policies and the notches they create. In Section 3 we describe our data. Our empirical analysis begins in Section 4 with an analysis of the Gas Guzzler Tax. In Section 5 we turn to a welfare analysis of these tax notches. In Section 6 we provide evidence that automakers also manipulate vehicle fuel economy in response to rounding rules in fuel economy label regulations. Section 7 concludes.

2. Fuel economy policy and presentation notches

The Gas Guzzler Tax and fuel economy label ratings are both based on dynamometer tests, in which test vehicles are placed on a dynamometer (a treadmill for cars) and "driven" through a specified course. During the test, the vehicle's tailpipe emissions are captured, and the amount of fuel consumed is calculated based on the quantity of captured gases. Two different courses are used: the "city" test, which mimics driving in an urban environment, and the "highway" test, which approximates highway travel. Both ratings are reported as miles traveled per gallon of fuel consumed (mpg). The city and highway fuel economy ratings from these tests, which we will call C and H respectively, are used as inputs in different functions to get ratings for the Gas Guzzler Tax, fuel economy labels, and Corporate Average Fuel Economy (CAFE) ratings.

2.1. The U.S. Gas Guzzler Tax

The Gas Guzzler Tax was passed into law in 1978. The tax, which penalizes cars with low fuel economy, is a notched schedule (a step function) in fuel economy. Thus, vehicles with very small differences in ratings may be subject to discretely different taxes. The tax was phased in between 1980 and 1991, but the schedule has not changed since 1991. However, because the tax is not adjusted for inflation, the real value has eroded. Table 1 shows the schedule over time.

The Gas Guzzler Tax rating is a harmonic average of the city and highway test results, equal to $(.55/C+.45/H)^{-1}$. Harmonic averages are used to obtain the average fuel economy based on a division of miles traveled because miles per gallon is an inverse measure (distance/fuel) of fuel economy. The weights, .55 on city and .45 on highway, reflect government estimates of the fraction of driving that occurs on average in urban versus highway contexts.

Light trucks, a designation that includes pickups, sport-utility vehicles and vans, were exempted from the tax from its inception, originally with the stated intention of not penalizing vehicles used for farming and commercial purposes. Because the Gas Guzzler Tax applies only to vehicles with low fuel economy and does not apply to light trucks, only a small fraction of the vehicle market is subject to the tax. Affected vehicles tend to be high-priced, high-performance cars with relatively low sales volumes. In 2008, 77 (out of 1248) vehicle configurations – a unique engine (including cylinders and displacement) and transmission – were subject to the tax, which raised about $172 million in revenue.[5] The tax is remitted by

---

[3] Tax economists have recently taken interest in the study of *kinks*—points where a policy causes a discrete change in the *slope* of a tax, arguing that the extent of bunching can identify structural parameters of utility functions (e.g., see Saez (2010)). Often, empirical estimates have shown a more muted behavioral response to kinks than would be suggested by theory. The incentives surrounding notches are frequently much starker, suggesting that notches may prove more useful than kinks in uncovering behavioral parameters. Also note that notches play a central role in regression discontinuity research designs, but in cases where the regression discontinuity is appropriate, the key identifying assumption is that the running variable is *not* subject to choice or manipulation, as it is here.

[4] These examples are drawn from Harberger (1995).
[5] Source: Internal Revenue Service, Statistics of Income, Historical Table 20.

Table 1
Gas guzzler tax rates over time (dollars per car).

Vehicle fuel economy (mpg)	1980	1981	1982	1983	1984	1985	1986–90	1991–Present
Over 22.5	0	0	0	0	0	0	0	0
22.0–22.4	0	0	0	0	0	0	500	1000
21.5–21.9	0	0	0	0	0	0	500	1000
21.0–21.4	0	0	0	0	0	0	650	1300
20.5–20.9	0	0	0	0	0	500	650	1300
20.0–20.4	0	0	0	0	0	500	850	1700
19.5–19.9	0	0	0	0	0	600	850	1700
19.0–19.4	0	0	0	0	450	600	1050	2100
18.5–18.9	0	0	0	350	450	800	1050	2100
18.0–18.4	0	0	200	350	600	800	1300	2600
17.5–17.9	0	0	200	500	600	1000	1300	2600
17.0–17.4	0	0	350	500	750	1000	1500	3000
16.5–16.9	0	200	350	650	750	1200	1500	3000
16.0–16.4	0	200	450	650	950	1200	1850	3700
15.5–15.9	0	350	450	800	950	1500	1850	3700
15.0–15.4	0	350	600	800	1150	1500	2250	4500
14.5–14.9	200	450	600	1000	1150	1800	2250	4500
14.0–14.4	200	450	750	1000	1450	1800	2700	5400
13.5–13.9	300	550	750	1250	1450	2200	2700	5400
13.0–13.4	300	550	950	1250	1750	2200	3200	6400
12.5–12.9	550	650	950	1550	1750	2650	3200	6400
Under 12.4	550	650	1200	1550	2150	2650	3850	7700

Source: Internal Revenue Service, Form 6197. All values are nominal.

manufacturers, but it is visible to consumers because it appears as a separate item on the sticker price.

2.2. Fuel economy label ratings

Every new vehicle sold in the United States is required to display a label that details the vehicle's manufacturer's suggested retail price, and since 1978 this label has also included the vehicle's official EPA highway and city fuel economy ratings. The font size of each item is mandated by law, and the city and highway fuel economy ratings must be set in the largest font, making these by far the most prominent numbers. In a much smaller font, the label also displays the combined rating, a graphic that compares the vehicle to others in its class, and an estimate of the annual cost of gasoline. Fig. 1 is an example of the fuel economy label in the United States as of 2008, when our sample period ends.

The city and highway ratings are integers, which are determined by rounding off the underlying fuel economy estimate derived from the test procedure. This rounding creates what we call a *presentation notch* — where a marginal difference in an underlying characteristic creates a discrete change in the information transmitted in the marketplace. A vehicle with a highway fuel economy rating of 29.49 will be listed as 29 on the label, whereas a vehicle with a rating of 29.50 will be listed as 30. If consumers value fuel economy, and if they use the official EPA ratings as a source of information, then firms may undertake costly adjustment procedures to increase the fuel economy rating as displayed on the labels, just as they would respond to tax notches.[6]

The label ratings are based on the same test results, C and H, that are used to derive the Gas Guzzler Tax rating. Prior to 1986, the reported label rating was simply the integer nearest the value resulting from the test procedure.[7] Starting in 1986, the EPA modified the procedure
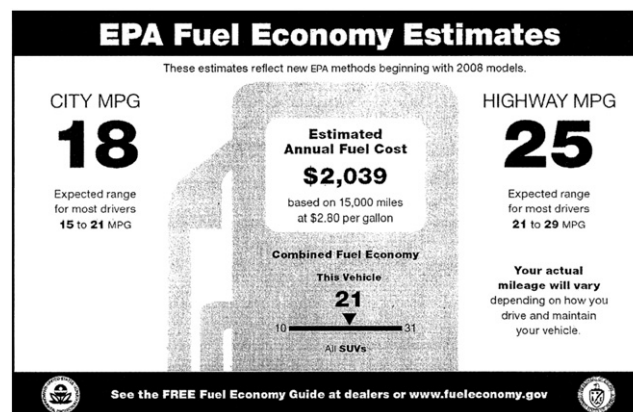

Fig. 1. A sample fuel economy label. Source: Environmental Protection Agency.

to adjust for "in-use shortfall" – the observation that experienced fuel economy was consistently below the ratings. After attempting to measure the discrepancy, the EPA decided to adjust the test numbers by multiplying the test output by a fixed factor—0.9 for the city and 0.78 for the highway test. Thus, the city label rating is $.9 * C$ rounded to the nearest integer, and the highway rating is $.78 * H$, rounded.[8] Automakers do have the right to adjust the label ratings downwards if they wish, and a very small percentage of ratings do reflect a downward adjustment, so that the test procedure indicates a higher value than appears on the label in practice. The automakers may do this to avoid consumer displeasure if actual fuel economy experiences fall short of their expectations.

While both the label ratings and the Gas Guzzler Tax rating are based on the same two inputs, because of the difference in functions vehicles that are close to a notch in one policy will not be especially likely to be close to a discontinuity in the others. Even so, some vehicles will be located next to both the Gas Guzzler Tax notch and one or both label notches, so that a modification that improves fuel economy will generate multiple benefits. Nevertheless, we analyze responses to the Gas Guzzler Tax notches and the label rating notches separately because these multiple benefits are unlikely to be empirically important for either policy. For the labels, most vehicles facing a label notch will be well out of the Gas Guzzler Tax range. For the tax, which binds only for very expensive luxury and sports cars, the relevant consumers, who have a revealed preference for performance over fuel economy, are unlikely to have a high willingness to pay for fuel economy.[9]

2.3. How is fuel economy manipulated?

Our subsequent analysis is based on the premise that automakers perform local manipulation of fuel economy ratings in order to move over a tax or presentation notch. If an automaker wishes to boost fuel economy locally around a notch, how is this done?

First, automakers may simply repeat the underlying fuel economy test if there is sufficient variability across trials. U.S. regulation requires, however, that all valid tests must be reported and averaged, limiting the efficacy of this strategy. Second, each "model type" receives a single rating, but a single model type may involve several test vehicles of different weights, which are then averaged, weighted by vehicles sold, to

---

[6] The unrounded figures are public information (indeed, we make use of these data in subsequent analysis), and so in principle consumers could also obtain and consider these unrounded numbers. To do so, however, they would have to download the publicly available fuel economy data files; the unrounded numbers are not included in the fuel economy guide that is available at dealerships and on the EPA's website. Even with these files in hand, a car shopper would need to know how to adjust for a factor called "in-use shortfall" in order to convert the unrounded numbers. We think this is unlikely, but if consumers do consider the unrounded information, we would not expect automakers to respond to label notches, which we test directly below.

[7] The EPA uses ASTM International rounding, which rounds a value ending in exactly 0.50 to the nearest even integer.

[8] Starting in 2008, the EPA instituted a new testing procedure, which is designed to improve the accuracy of label ratings. In this paper, we investigate only models that were tested during the pre-2008 regime.

[9] Purchasers of high-end sports cars may even have a distaste for fuel economy. In online appendix section E, we provide evidence that supports both of these claims. We demonstrate that dropping vehicles affected by the tax does not change our results regarding labels, and we show that tax-paying vehicles that are near label notches are no more likely to respond to tax notches than those that are not.

create a single rating. An automaker could produce less of one configuration and more of the other to move the average rating, which determines the tax liability for all vehicles in a model type, but this is likely to be an expensive strategy.

Finally, and most importantly, an automaker may actually modify a vehicle to improve its fuel economy. Methods include "light-weighting" (substituting vehicle parts to reduce weight), engine recalibration (reprogramming the vehicle to operate in a different gear at certain speeds), use of low-friction lubricants, modifications to tires, or small aerodynamic changes such as the addition of a spoiler, side skirts, air dam reshaping, or the installation of "belly pans" that smooth air flow by covering parts underneath the vehicle.[10] It is important to distinguish these methods from what one might characterize as global design choices. The overall structure of the engine, weight, the shape of the car, the use of fuel saving technologies and the choice of transmission are all key determinants of fuel economy, but these decisions are made on a several year time horizon, long before vehicles are officially tested and the exact location of a vehicle vis-à-vis a notch is known. In contrast, the relatively minor adjustments we listed above could potentially be adopted late in the production cycle, in response to preliminary test results.

To capture this distinction, we model the vehicle design process as consisting of two stages. In the first stage, which lasts several years, major choices about a vehicle are made. In the second stage, which may be only a few months before production, small modifications can still be made, but most decisions are set. During the first stage, automakers have an expectation about what a vehicle's fuel economy will be, but they are not certain until prototype vehicles are available for testing. In the second stage, prototype vehicles are available, much uncertainty is resolved, and the automaker may choose to make small modifications to fuel economy – like the ones we list above.[11] These second-stage techniques allow automakers to fine tune fuel economy, but they do not necessarily allow automakers to choose an exact fuel economy rating because each of the tweaks involves a discrete change. Thus, we would expect strategic manipulation to result in bunching of fuel economy at, or slightly above, a notch cutoff. Also, because these modifications are costly, we would expect that some automobiles may remain at fuel economy ratings on the high-tax side of notches.[12]

Under the assumption that automakers have optimized the vehicle according to cost and consumer preferences in the first stage, the full cost of a second-stage design tweak for the purposes of moving over a notch includes both the direct cost and any drop in consumer value caused by a deviation from the pre-manipulation optimal design. For example, a spoiler might cost $100 per vehicle to install, but if the consumers of that model think a spoiler is unpleasant, it could cost much more by way of decreased consumer willingness to pay. Automobile construction involves careful planning about vehicle design, product testing, consumer marketing, parts sourcing, and assembly line production. Tweaks made weeks or months ahead of production may be infeasible or disruptive, or simply introduce uncertainty about production, such that what seems like a relatively straightforward material change to a car could be very costly.[13] As such, we believe it is not obvious ex ante whether automakers will be responsive to notches.

Automakers are reluctant to publicize information about how local fuel economy adjustment might occur or say whether they respond to notches. However, our conversations with experts who have worked for automakers and officials at the EPA indicate that this type of vehicle modification does indeed take place. Anecdotal evidence from the popular press also provides support. In 2009, when the recent "cash-for-clunkers" bill was passed, Nissan stated its intention to alter fuel economy for certain models to ensure that they met fuel economy eligibility requirements. Another oft-cited example concerns what is known as Computer Aided Gear Selection, sometimes called a "skip shift," which forces a manual transmission vehicle into a first-to-fourth gear shift at certain speeds. Popular consensus is that this feature is installed as a way of reducing the Gas Guzzler Tax, and after-market kits are available that claim to disable this feature.

3. Data sources

We gathered fuel economy data from several sources. For the Gas Guzzler Tax, we obtained from the Internal Revenue Service a complete list of all vehicles that were subject to the tax from the beginning of the program in 1980. These data include fuel economy ratings to a tenth of a mile-per-gallon, and are limited to the set of vehicles that were actually taxed, but are complete for all years between 1980 and 2009.

We complement this data with fuel economy ratings from the EPA, which provides unrounded city and highway test results from 1978 to 1983 and from 1999 to 2007. We use these underlying test statistics to reconstruct the Gas Guzzler Tax rating for all vehicles using the formulas published in federal regulations. The EPA data have the advantage of allowing us to calculate the Gas Guzzler Tax rating for vehicles that were not subject to the tax, but it has the disadvantage of a coverage gap. Between 1984 and 1998, the EPA data do not include the unrounded test results necessary for calculating the correct Gas Guzzler Tax rating.

These same EPA data are our source for fuel economy label ratings. We transform the unrounded city and highway test results according to the EPA's in-use shortfall adjustment factors in order to obtain the adjusted, unrounded fuel economy label ratings. In our analysis of CAFE ratings, we use official CAFE fuel economy ratings, which differ slightly from the EPA ratings used in the other sections, from the National Highway Transportation Safety Administration (NHTSA). CAFE data include sales volumes, and we match the CAFE data set to the IRS Gas Guzzler list in order to measure the sales volumes of taxed vehicles. We use this source of sales data instead of more conventional sources like Automotive News because conventional data sources do not divide the sales of a model line among the different engine configurations, so that the unit of observation is not the same across data sources.

4. The behavioral response to fuel economy notches

4.1. The Gas Guzzler Tax

If automakers respond to notches in the Gas Guzzler Tax by modifying vehicle fuel economy, then the distribution of fuel economy ratings should feature "extra" observations just on the tax-preferred side of notches. Fig. 2 is a histogram of the number of models by their one-decimal Gas Guzzler Tax rating between 1991 and 2009, during which time the Gas Guzzler Tax schedule was stable and notches were present at each rating ending in .5. Bars at a .4 decimal (the high-tax side of a notch) are shaded blue, and bars at a .5 rating (the low-tax side of a notch) are shaded red. (For those reading in black and white, in each pair of shaded values the blue bar is on the left and the red bar is on the right.) Of the ten different integer values, the number of vehicles at .5 exceeds the number at .4 in seven cases,

---

[10] These examples are drawn from Edmunds.com and National Research Council (2002). Aerodynamics influence fuel economy ratings because testers take vehicles on an actual driving course and "coast" the vehicle to test its aerodynamic efficiency. The results are used to make small modifications to the dynamometer output.

[11] Longer-horizon global choices are examined by Klier and Linn (2010) and Knittel (2011).

[12] These costs of manipulating ratings are analogous to optimization frictions that influence the shape of bunching around income tax notches, as in Kleven and Waseem (2012).

[13] For that reason, engineering cost estimates from federal regulatory documents are only a loose lower bound on the modification costs relevant for our analysis.
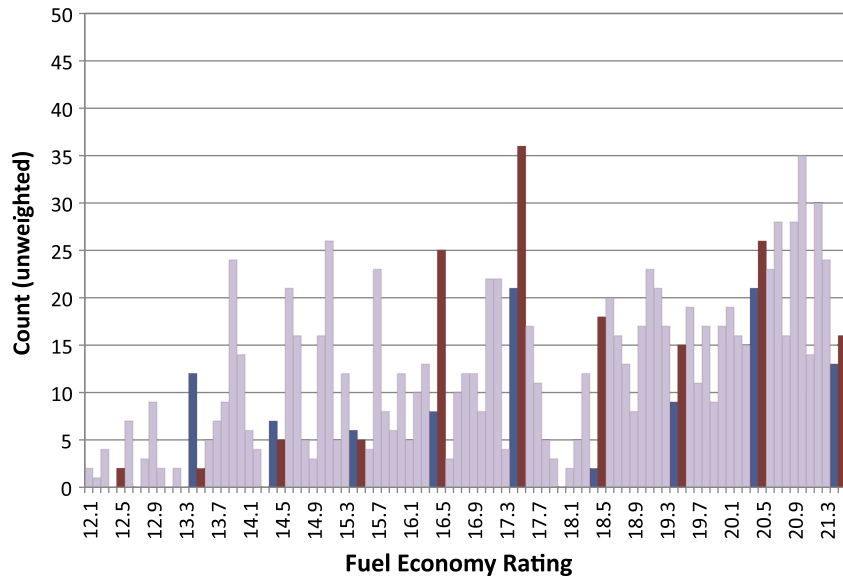
Fig. 2. Gas guzzler rating distribution, unweighted: 1991–2009. Note: IRS data, sample size is 1221. Higher fuel economy ratings correspond to lower taxes. Ratings ending in .4, all of which are just below a tax notch, are colored in blue, while ratings ending in .5 are colored in red. For those reading in black and white, in each pair of shaded values, the blue bar is on the left and the red bar is on the right.

sometimes by a large margin. Overall, there are 150 models at whose rating ends in .5 and 99 whose rating ends in .4. The probability that, of 249 draws, 150 or more would be drawn from a binomial distribution with equal probability is just 0.0007. If we widen our window and compare the number of models at .3 and .4 to the number at .5 and .6, the story does not change: 200 are just below the notch, and 295 are just above.

The counterexamples to the preponderance of .5 decimals over .4 decimals are high-performance, high-price ultra-luxury automobiles with very low fuel efficiency. Manufacturers of these cars may perceive that their prospective buyers care little about a few hundred dollars because it is a small fraction of the total cost, or even perceive

that a low mpg is a status symbol of high performance. These models also have relatively low sales volume, so that if modification involves a fixed cost, we would see less bunching among these vehicles. To capture this possibility, Fig. 3 replicates Fig. 2 but weights the distribution by sales volume. Note that the total economic impact of manipulation depends on the sales-weighted distributions. In this figure, the predominance of .5 decimals is even more pronounced, and the integers where .5 does not predominate feature very low sales.

Fig. 4 aggregates across integers to show a histogram of mpg *decimal* values for all vehicles subject to the Gas Guzzler Tax. For example, if a vehicle had a 20.5 fuel economy rating, we put that vehicle into the .5
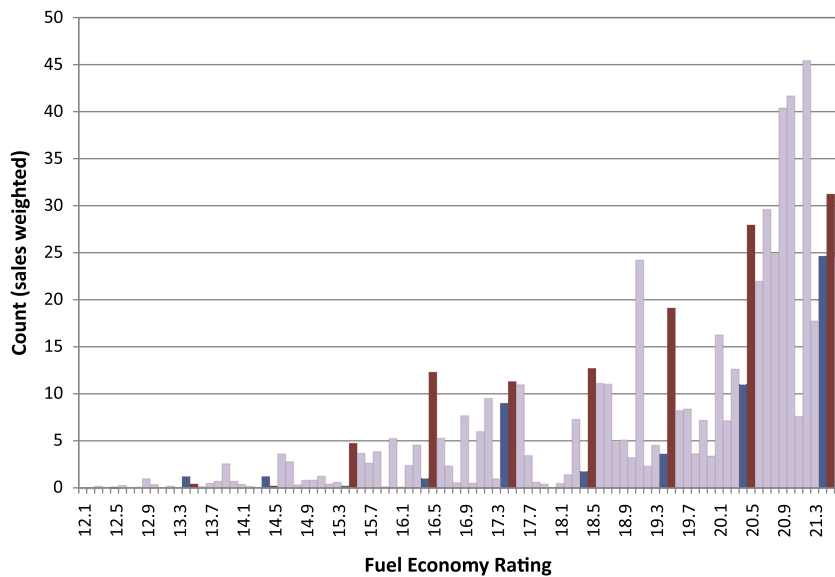

Fig. 3. Gas guzzler rating distribution, sales weighted: 1991–2007. Note: IRS fuel economy data and NHTSA sales data; sample size is 841. Higher fuel economy ratings correspond to lower taxes. Ratings ending in .4, all of which are just below a tax notch, are colored in blue, while ratings ending in .5 are colored in red. The sample differs from Fig. 2 because some vehicle types are missing sales information and sales data are unavailable for 2008 and 2009. For those reading in black and white, in each pair of shaded values, the blue bar is on the left and the red bar is on the right.
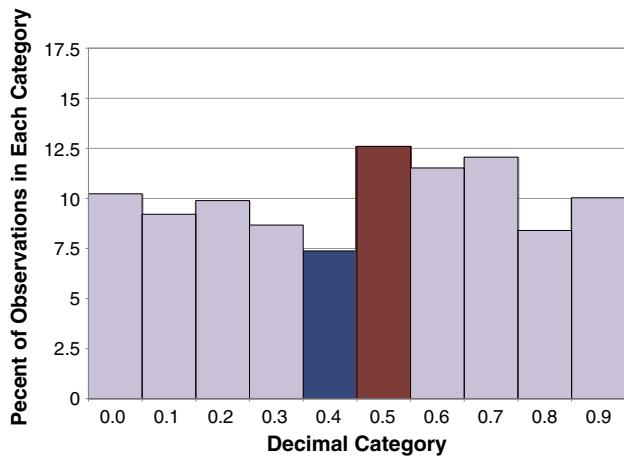
Fig. 4. Gas Guzzler Tax decimal distribution, unweighted: all vehicles, 1980–2009. Note: IRS data; sample size is 1476. In several years, the notch is at whole integers (the .0 bin). For those years, we shift decimals by .5, so the notch is always represented by the .5 bin.
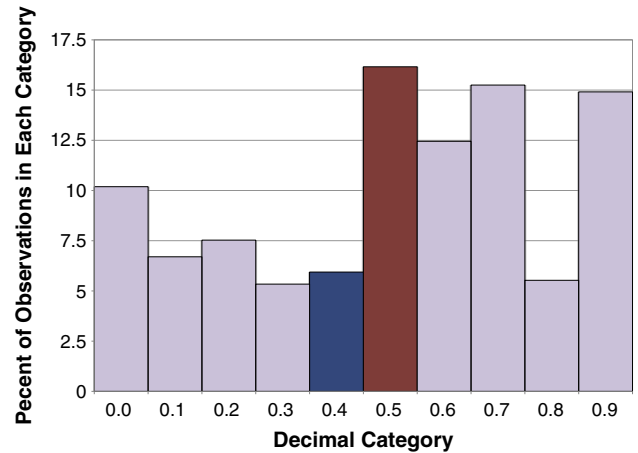

Fig. 5. Gas Guzzler Tax decimal distribution, sales weighted: all vehicles, 1988–2007. Note: IRS and NHTSA data; sample size is 945.

bin. Aggregation allows us to include more data by combining different tax regimes.[14]

In the absence of tax incentives, we might expect this decimal distribution to be uniform. The actual distribution shows a marked departure from uniformity, with far more observations just at, or just to the right of .5.[15] This difference is highly unlikely to be due to chance. Comparing either the number of vehicle configurations in the .4 bin to the .5 bin, or comparing the sum of the .3 and .4 bins to the sum of the .5 and .6 bins, yields a p-value less than .0001 on a test that they are drawn from a uniform distribution.[16] Fig. 5 provides a sales-weighted histogram of the ratings' decimals. Here, the gap between sales around the notch is even more pronounced, though the distribution shows greater variation overall.[17]

The statistical tests cited above are based on the assumption that, in the absence of notch responses, the occurrence of .4 and .5 decimals would be the same. This assumption may not be precisely correct if the overall fuel economy distribution has a positive slope, in which case there might be more .5 decimals for reasons unrelated to the notches. If this were driving our results, we would expect Figs. 4 and 5 to show a tilt across all decimals, i.e., there would be more .1 than .0, more .2 than .1, etc. We do not see this pattern.

To further dispel such concerns, we redid our statistical tests after accounting for the overall shape of the fuel economy distribution, the results of which we report in Table 2. First, we estimate a polynomial through the frequency distribution in Fig. 2, omitting observations at the .4 and .5 decimals. We then use the predicted values from these polynomials to predict the relative number of .4 and .5 decimals

that should occur, given the shape of the distribution. Combining these estimates yields the predicted probability that a vehicle would have a .5 decimal, conditional on the observation being either .4 or .5, under the null hypothesis that the polynomial predicts the relative prevalence correctly. We then use this new predicted probability to ask how likely it is that we would have observed 150 observations at .5 out of 249 that were either .4 or .5. Rather than simply doing a single t-test with the adjusted probabilities, however, we bootstrap this entire procedure (starting by resampling our microdata) so as to incorporate the variance that arises from the estimation of the polynomial.

Table 2 shows that this adjustment has very little impact on the estimated probabilities. The first row of the table shows what we label the binomial model, which is our original assumption that the counterfactual probabilities of ratings ending in .4 and .5 are equal. We observe 150 out of 249 above the notch. Under the binomial model assumption, the expected number of observations above the notch is 124.5, and the standard deviation is 7.89. The second row of Table 2 shows the expectation and standard deviation when only a linear control is used, and the third row shows a fifth-order polynomial. The expected number of observations above the notch and the standard deviation change only slightly. This is not surprising because the overall shape of the distribution in Fig. 2 does not exhibit a dramatic slope. The probability of observing 150 observations above

Table 2
Statistical tests of bunching above Gas Guzzler Tax notches.

	Observed number above notch (out of 249 within .1 mpg)	Expected number above notch under null hypothesis	Standard deviation of number above notch under null hypothesis	p-value
Binomial model	150	124.5	7.89	0.0007
Linear control function	150	125.4	7.89	0.0009
Quintic control function	150	124.8	7.97	0.0008

Note: The binomial model assumes that the probabilities that a vehicle has a fuel economy rating ending in .4 and .5 are the same. The linear model modifies this assumption by adjusting for the overall shape of the distribution using a linear fit, excluding observations within .1 mpg of a notch. The quintic model extends this by using a fifth order polynomial to estimate the distribution's shape. Statistics for the linear and quintic models are derived via bootstrap to account for sampling variation in the estimated polynomial.

---

[14] As shown in Table 1, the value of the Gas Guzzler Tax changes at each .5 in fuel economy ratings, except for 1980, 1981, 1983 and 1985, for which years we adjust the data to match the .5 notch point in the figure.

[15] Results are qualitatively similar if we restrict the sample to unique observations by dropping all vehicles with the same manufacturer, cylinders, displacement, transmission and fuel economy rating as some other vehicle, either within or across years. This restriction is intended to drop repeated observations of the same engine, which may be installed on several different models. All results in the paper are robust to dropping such cases.

[16] This significance test comes from treating the observations in the restricted distribution just around the notch as a binomial distribution, with points above the notch treated as a successful trial. We calculate the p-values reported here using the normal approximation to the binomial distribution.

[17] Results are even more dramatic for the Big Three automakers, for whom there is even more pronounced bunching, though the sample size is small – less than 7% of vehicles that pay the Gas Guzzler Tax were made by the Big Three.

Table 3
The correlation between bunching and sales volume and tax notch value.

Dependent variable = 1 if vehicle is on tax-preferred side of notch						
Sample includes vehicles with ratings ending in .3, .4, .5 and .6						
	IRS Data (1980–2009)				EPA Data (1998–2007)	
	1	2	3	4	5	6
Log sales		0.0439***		0.0435***		
		(0.0149)		(0.0161)		
Gas guzzler value ($100)			0.00179	−0.000548		0.0136***
			(0.00684)	(0.00679)		(0.00393)
Gas guzzler range dummy					0.126***	
					(0.0328)	
Constant	0.600***	0.340***	0.589***	0.345***	0.447***	0.450***
	(0.0201)	(0.0931)	(0.0463)	(0.0463)	(0.0113)	(0.0111)
Observations	593	371	593	371	2207	2207
R^2	0.00001	0.023	0.0001	0.023	0.007	0.005

Note: Columns 1 through 4 use IRS data, which include only vehicles that pay the Gas Guzzler Tax. Columns 5 and 6 use EPA data. This enables the inclusion of vehicles that do not pay the Gas Guzzler Tax, including both vehicles located just above the top notch and those further away who face a zero incentive for moving over .5. Sales data are not available in the EPA data, so not all specifications are repeated with both samples. All asterisks indicate statistical significance relative to a null hypothesis that the coefficient equals 0.

the notch is extremely unlikely under any of the modeling alternatives, which bolsters the conclusion that the data reveal strategic responses to notches.[18]

4.2. Determinants of bunching

If automakers respond to notches, we would expect to find more bunching around notches with greater tax values. And, if there are fixed costs to second-stage modifications, we would expect more bunching among vehicles with higher sales. We test these additional hypotheses using a linear probability model, where the dependent variable is coded as 1 if an observation falls on the tax-preferred side of a notch, in a sample restricted to cars within a window on either side of notches.

Table 3 shows results using a window of .2 mpg around notches, so the sample includes vehicles whose Gas Guzzler Tax rating ends in either .3, .4, .5 or .6, and the dependent variable is coded as 1 for the .5 and .6 observations. The first four columns use data from the IRS. This allows us to include all years of data, from 1980 to 2009, and to include sales volume, but it restricts the sample to observations that are subject to the Gas Guzzler Tax. This means that we cannot include the top notch, nor can we include vehicles that are not around a tax notch as a baseline. Columns 5 and 6 use EPA data instead, which allows us to include the top notch and vehicles not near a notch as a baseline, but it restricts the years of availability and does not allow us to include sales volume controls due to data limitations.

Column 1 includes only a constant, which is equal to the fraction of observations on the tax-preferred side of notches. This estimate is statistically different from .5, which would be the counterfactual value under the null hypothesis of no strategic bunching. Column 5 provides a similar test in the EPA data, but it includes both a constant and a dummy variable coded as 1 if the vehicle is in the range of the Gas Guzzler Tax because there are also vehicles in the sample that are not near a notch. The fact that the dummy is positive and statistically significant indicates that vehicles facing a Gas Guzzler Tax notch are more likely to have a decimal of .5 or .6 than vehicles with fuel economy outside the Gas Guzzler Tax range.

Column 2 shows that vehicles with higher sales volumes are significantly more likely to be on the tax-preferred side of a notch. If automakers were not bunching strategically and the "extra" observations on the tax-preferred side of notches were due to chance, there would be no reason to expect those observations on the tax-preferred side to be higher volume.[19] If, however, there are fixed costs in manipulation, we would expect higher sales volume vehicles to be more likely to bunch. The strong correlation between sales and bunching therefore provides additional support for our hypothesis that automakers respond to notches.

The IRS data (column 3) shows a small, statistically insignificant coefficient on notch value. In the EPA data (column 6), however, there is a positive coefficient, so there is more bunching surrounding notches with higher values. The difference is due to the importance of the top notch, which has many data points and a large notch value of $1000. Excluding data from the top notch in the EPA data produces a coefficient similar to the IRS data (not shown).

4.3. False experiments

One might be concerned that the preponderance of vehicle configurations with ratings at or just above .5 might be an artifact of some unknown engineering property or other anomaly. We test this by reproducing the Gas Guzzler Tax fuel economy measure for vehicles not subject to the tax and looking for the same pattern. This requires use of the EPA data, which limits the sample to 1999–2007. First, we reproduce the main result from Fig. 4, for vehicles subject to the tax, for this subsample of years in part (a) of Fig. 6, which shows evidence of bunching but with considerably greater noise than the main figure with the unconstrained sample.

Second, we look for bunching behavior around notch values when no tax incentive for bunching exists. Part (b) shows the distribution of fuel economy decimals for passenger cars that have ratings above the Gas Guzzler Tax threshold mpg and therefore have no incentive to bunch at or just above .5 decimals. Likewise, part (c) shows the distribution of rating decimals for light trucks in the same fuel economy range as the passenger cars in part (a). Because light trucks are not subject to the Gas Guzzler Tax, there is no incentive to bunch. The fact that neither of the untaxed classes of vehicles exhibits bunching

---

[18] We present several additional robustness results in the appendix, which is available online. An alternative methodology is to collapse the data and perform statistical tests on the aggregated cell counts, treating each fuel economy rating as the unit of observation rather than each vehicle. The appendix also presents placebo tests for bunching around decimal values where there are no notches. The appendix also repeats this analysis with standard errors clustered on manufacturer and on manufacturer interacted with cylinder. Generally, all of these results are supportive of the conclusions drawn here.

[19] Here we describe sales volume as an exogenous characteristic, but a lower Gas Guzzler Tax for a vehicle will lead to a lower price and thus increase sales, making sales an endogenous characteristic. This is true, but an extraordinarily high price elasticity would be required to explain the magnitude of the differences in sales volume for high and low tax vehicles observed in the data.
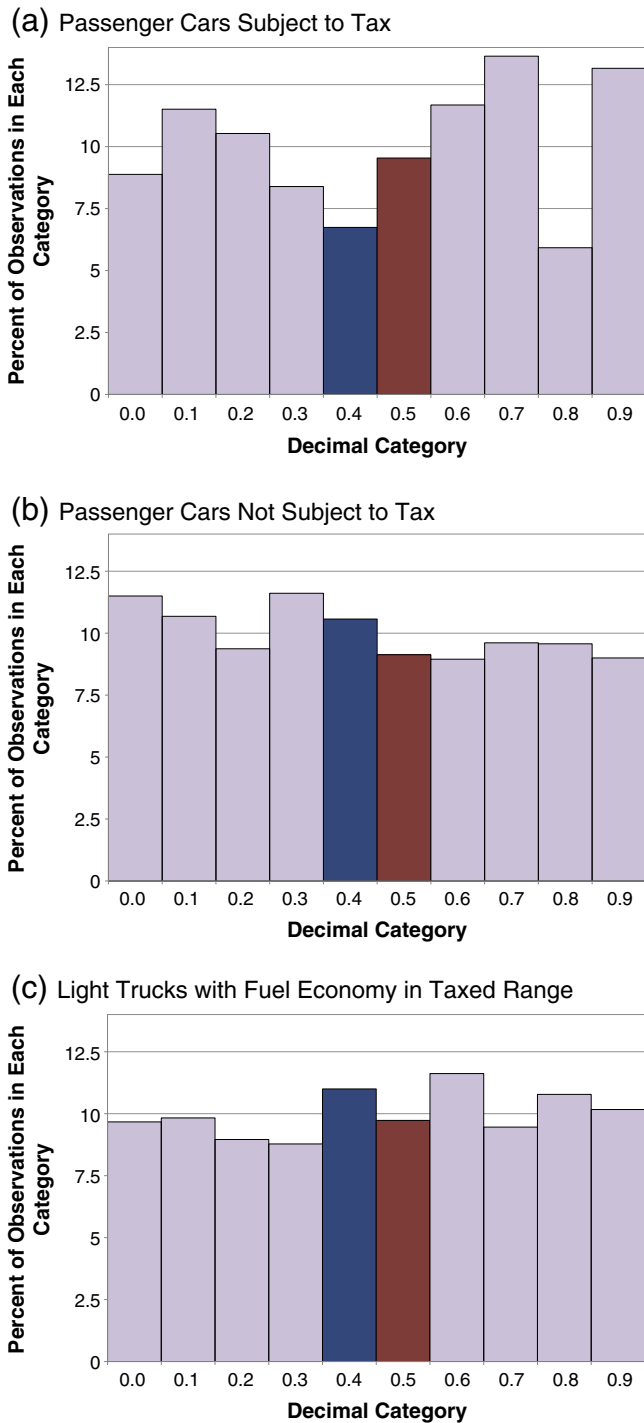
(a) Passenger Cars Subject to Tax

(b) Passenger Cars Not Subject to Tax

(c) Light Trucks with Fuel Economy in Taxed Range

Fig. 6. Gas Guzzler Tax rating decimals for several vehicle groups: 1999–2007. Note: Total sample sizes are (a) 608, (b) 5422 and (c) 3236. Vehicles in parts (b) and (c) are not subject to the tax and not expected to exhibit bunching.

is further evidence that the bunching in the vehicles subject to the tax is due to a strategic response to notches.

For an additional false experiment, we examine a closely related fuel economy measure, individual vehicle CAFE ratings, which do not have a notch at .5. Each vehicle in a manufacturer's fleet is given a CAFE rating based on a weighted average of the vehicle's city and highway fuel economies. These combined ratings, which are calculated to the tenth of an mpg (e.g., 27.5 mpg), are used to calculate a sales-weighted average for all vehicles made by a given manufacturer. This sales-weighted average is then rounded to a tenth of a
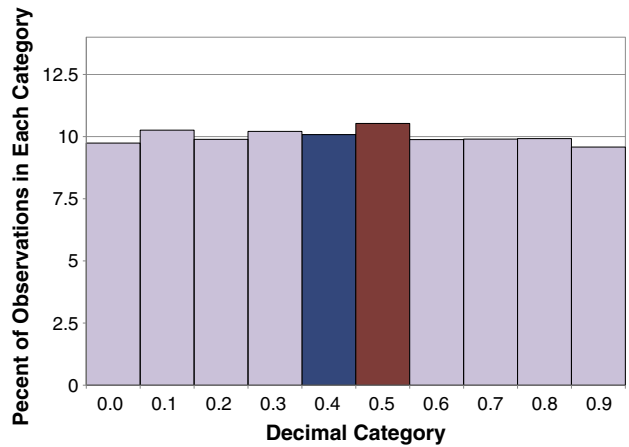

Fig. 7. Distribution of CAFE combined rating: 1978–2005. Note: NHTSA data; sample size is 18,045.

mile-per-gallon for use in determining compliance with CAFE. Because the individual fuel economy ratings are not rounded to integers prior to averaging, there is no incentive for manufacturers to push individual vehicle CAFE ratings above any particular decimal.[20] Fig. 7 shows the combined CAFE rating decimal distribution. The ratings are roughly uniform, as expected. There are slightly more observations in the .5 bin than the .4 bin, but this difference is not statistically significant.[21] This is further evidence of our main conclusion.

4.4. Bunching above the top gas guzzler notch

Looking only at the fuel economy ratings used in the IRS data does not reveal what is arguably the most striking evidence of bunching, which is found at mpgs just over the final Gas Guzzler Tax notch, above which no tax is due. Getting fuel economy above this final notch allows a vehicle to not only lower its tax burden, but also avoid being negatively branded as a "guzzler." The IRS does not publish ratings for vehicles that are not subject to tax. Thus, to examine bunching at the value that allows vehicles to escape the tax altogether, we are limited to our EPA data, available from 1978 to 1983 and 1998 to 2007.

Between 1978 and 1983, the Gas Guzzler Tax schedule was changing (see Table 1). In 1978 and 1979, there was no tax. The tax began in 1980, at which time it had a "top" notch of 15.0. This changed to 17.0 in 1981, and then to 18.5 in 1982, and finally to 19.0 in 1983. Fig. 8 shows the distribution of fuel economy ratings for passenger cars in each of these six years. In each diagram, the dashed blue vertical lines indicate the location of *future* top notches. The unbroken red vertical lines indicate the top notch effective for the year shown.

These six figures suggest a precise response to the top notch. Before the policy, a large fraction of vehicles lay to the left of the blue lines, so they would be subject to a tax in future years. When the tax is introduced for cars below 15.0 mpg in 1980, a majority of the vehicles that were previously below this level are gone. The same adjustments occur in 1981, 1982 and 1983; in each year, most of the vehicles that would have been just below the notch have moved. The entire distribution shifted rightward, not just vehicles

---

[20] Manipulation of the Gas Guzzler Tax rating will translate into manipulation of the CAFE rating, because the two numbers are identical in early years and extremely close to each other in later years. Thus, we omit passenger cars with combined fuel economy ratings below 23, which would be subject to the Gas Guzzler Tax.

[21] A test of the difference between the .4 and .5 bins yields a one-sided p-value of .092, and a test of the difference between the .5 and .6 bins from the .5 and .6 bins yields a one-sided p-value of .400. Overall, a chi-squared test statistic of the null hypothesis that the data are distributed multinomially with equal probability on each bin cannot be rejected (p-value of .994).
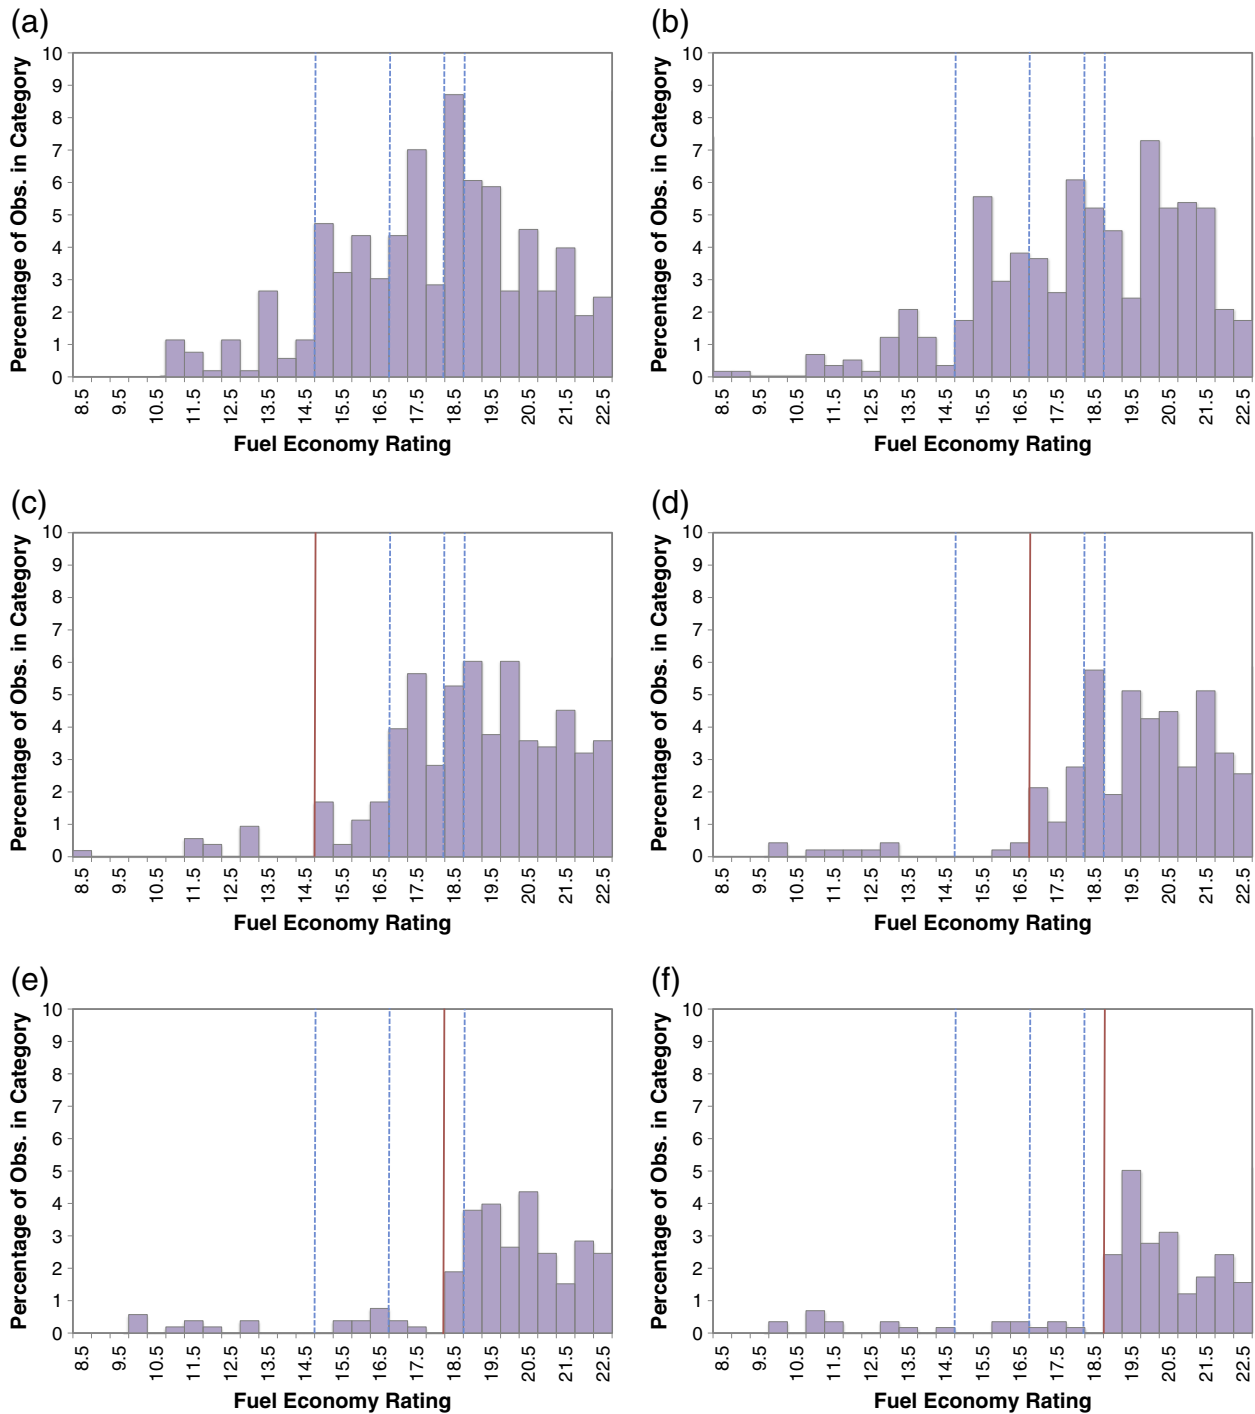
Fig. 8. Gas guzzler tax rating distribution during tax phase in (a) 1978 (no tax), (b) 1979 (no tax), (c) 1980 (top notch = 15.0), (d) 1981 (top notch = 17.0), (e) 1982 (top notch = 18.5), (f) 1983 (top notch = 19.0). Note: Data come from the EPA. The top notch in a year is the minimum value that a vehicle needed to avoid paying the Gas Guzzler Tax entirely. Distributions are truncated at 22.5 mpg, but percentages on the horizontal axis reflect the entire distribution, including the portion not shown.

near the notch.[22] The overall shift is likely due to CAFE, which was introduced at the same time, but it appears that the details of the distribution were driven by the location of the top Gas Guzzler Tax notch.

The EPA data are unavailable during the rest of the tax's phase-in period, but they are available for several years after the top notch of 22.5 mpg was established in 1991. Fig. 9 is a histogram of the number of models of all vehicles produced from 1991 to 2007 by their CAFE rating, whether they are subject to the Gas Guzzler Tax or not. The

dark vertical line is drawn at the tax threshold of 22.5 mpg. The vast majority of models are above the Gas Guzzler Tax threshold, and there is a clear asymmetry in the histogram. To do a rough quantification of the bunching above the top notch apparent in Fig. 10, we estimate an 8th-order polynomial on the distribution of passenger car Gas Guzzler Tax ratings over the available sample period from 1998 to 2007, omitting data within 1 mpg of the top notch at 22.5.[23] We then

---

[22] This is evident in Fig. 8 from the reduction in the fraction of vehicles getting under 22.5 mpg. The entire distribution is not shown to preserve visibility of the relevant portion, but the heights of the histograms are preserved to be comparable.

[23] Specifically, we run a regression of the form $count = \sum_k \beta_k mpg^k$ that includes vehicles with mpg between 10 and 21.5 or between 23.5 and 45. The predicted values of this regression will be a consistent estimate of the counterfactual distribution if no manipulation occurs outside of this 1 mpg range.
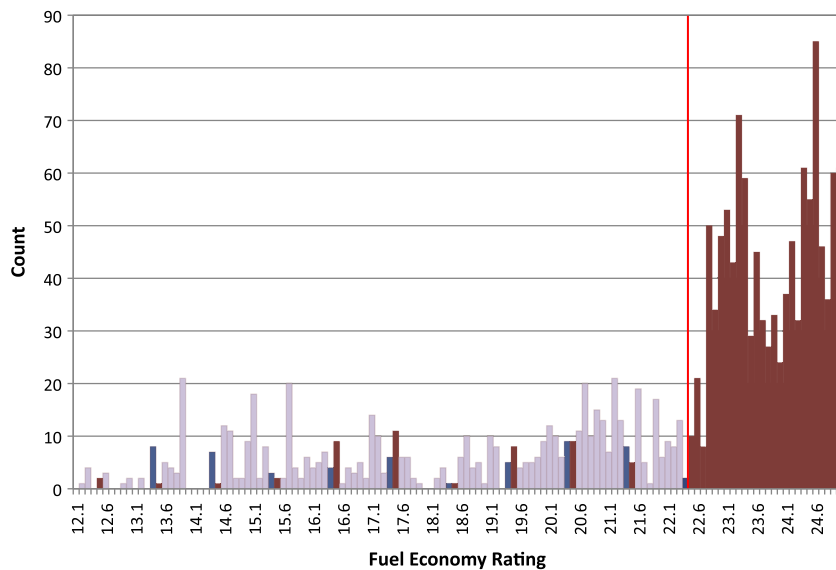
Fig. 9. Gas Guzzler Tax rating distribution: 1998–2007. Note: Data come from the EPA.

take the predicted number of observations near the Gas Guzzler Tax top notch and compare that to the actual number of observations. This regression, which has an R^2 of .81, predicts that 57% of the vehicles within 1 mpg of the top notch will be above the notch – this is more than half because of the upward slope in the distribution. In the actual data, 82% of the observations in this window are above the notch. This implies that over two-thirds of the vehicles within 1 mpg of the top notch moved in response to policy, which is equivalent to 3.4% of the vehicles in the entire car market (of whom only a modest fraction are near the notch).

There is nothing special about 22.5 mpg in terms of technology that could explain the apparent discontinuity in the distribution. To be convinced of that, consider Fig. 10, which is a scatter plot, separately for 1978 and 2004, of models by horsepower and fuel economy rating, where the size of the circles indicate the sales volume of each observation.[24] In each year, there is a negative relationship between the two attributes: other things equal, fuel economy suffers as a vehicle's power increases. The outward shift between 1978 and 2004 indicates technological progress in the intervening years — fuel economy for a given horsepower has increased markedly.[25] For our purposes, what is of interest is the bunching of observations just to the right at the Gas Guzzler Tax threshold of 22.5 mpg, shown by the vertical line. No such bunching is observed in the 1978 data, before the enactment of the Gas Guzzler Tax.

Fig. 11, adapted from Fig. 8 in Sallee (2011b), plots the market share of vehicles near the 22.5 mpg threshold since 1978, separately for taxable cars, tax-exempt trucks, and all vehicles. The car series has a precipitous drop-off in that fuel economy class just as the Gas Guzzler Tax began to affect cars with mpg as high as 20, in 1985. In contrast, the market share of trucks in that mpg class jumped sharply at the same time. Sallee (2011b) interprets this as evidence that automakers responded to the tax by designing, producing and marketing vehicles in that part of the fuel economy distribution to qualify as light trucks, thereby avoiding the tax and relaxing CAFE constraints.

If correct, this implies that notches induce qualitative product design changes, as well as the second-stage modifications that are our primary focus.

One reason for the especially large bunching at the top tax threshold is that the tax saving at that notch is significantly larger than for most other notches. Moving over the top notch saves $1000 (since 1991), whereas other notches are as small as $300. We propose, though, that it may also be that car manufacturers placed a value on avoiding the stigma of a vehicle being officially labeled a gas guzzler, regardless of the tax liability that came with that designation.[26]

5. Welfare analysis

Having established that automakers manipulate fuel economy in response to notches in the Gas Guzzler Tax, we now develop a framework to assess the welfare implications of this behavior. We assume that a fuel economy subsidy exists because there is a constant per-mpg externality associated with it, and that automobiles are designed in two stages, the latter of which enables producers to fine tune fuel economy. Under the assumption that the per-mpg externality e is constant, a first-best response is achieved by a smooth linear tax $t = e$. Instead of a linear tax, however, the policy in place is a notched subsidy that jumps by e at every 1 mpg and is flat in between (a step function). For our calculations, we assume that $e = \$700$, which is the average size of a Gas Guzzler Tax notch. The tax is therefore assumed to have the correct global slope, but its notches create capriciously varying local incentives that distort second-stage choices, as we explicate below. Our aim is to understand the welfare impact of substituting a notched policy for a smooth linear tax in this setting.

Our interest is in how notches impact product design, and we make several assumptions that allow us to isolate this behavioral response. We assume that a continuum of consumers have unit demand, and that each is matched to a particular vehicle type (based on heterogeneity in tastes for attributes other than mpg) so that

---

[24] Fuel economy data in this plot are taken from CAFE data, where the fuel economy rating is very similar, but not identical, to the Gas Guzzler Tax rating. CAFE and Gas Guzzler Tax ratings differ by around .1 or .2 mpg, so the CAFE rating can be used to show a broad picture but not a precise decimal analysis.
[25] See Knittel (2011) for an analysis of this technological shift.

[26] It is often said that domestic automakers choose to comply with CAFE standards even though it would be cheaper to pay the regulatory fine for non-compliance because of a similar stigma attached to being labeled a dirty firm. The estimates in Anderson and Sallee (2011), however, suggest that compliance may actually be less costly than the fine so that stigma is not required to explain compliance.
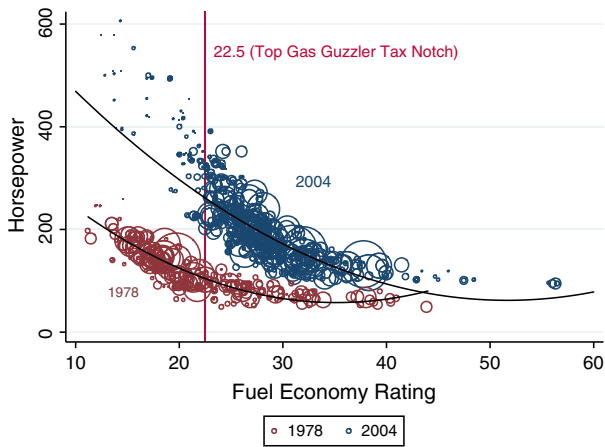
Fig. 10. Horsepower and fuel economy in 1978 and 2004: Passenger cars. Source: Authors' calculations of NHTSA data.

marginal changes in fuel economy do not change demand across vehicle types. This means that aggregate fuel economy can be changed by improving existing vehicles, but not by inducing substitution across vehicles. We also assume that firms price at consumer's marginal willingness to pay for all attributes, including fuel economy. This implies that firms capture any government subsidy, and, even in the absence of policy, all privately beneficial improvements to fuel economy are made. Finally, we assume that manipulation involves only an increase in marginal costs (and not a fixed cost) of production, which simplifies the analysis and accords with our abstraction from sales quantities in favor of a focus on product design.

As discussed in Section 2.3, we divide the automobile design process into two stages. In the first stage, which begins many months before production, the automaker makes choices regarding engine size, body style and vehicle features that cannot be changed quickly and have large impacts on fuel economy. Automakers have uncertainty regarding the fuel economy of each model. Specifically, we assume that the proximity to the next notch at the start of the second stage is uniformly distributed on 0 to 1. The automaker has no information about the distance each vehicle will be from the nearest notch at the beginning of the second stage, except that this distance is uniformly distributed. For concreteness, one example of this form of uncertainty would be if the automaker knows the integer value of each model's

fuel economy rating but not the decimal (which is uniformly distributed) during the first stage.[27] In the second stage, which may be only a few months before production (and at which point producers can run fuel economy tests on prototypes), automakers can make small design changes that influence fuel economy by modest amounts. Specifically, we assume that automakers can change fuel economy by less than 1 mpg in the second stage.

5.1. Producer choice for notched tax

In the first stage, producers know the integer of their fuel economy rating but not the decimal. At the beginning of the second stage, uncertainty is resolved and the automaker observes both the integer and the decimal of their fuel economy rating. If the subsidy is notched, the producer observes how far their vehicle is from the nearest notch in mpg units, which we denote as $z \in (0,1]$. Because second-stage modifications can improve fuel economy by at most 1 mpg, and because notches are 1 mpg apart, the automaker faces a discrete choice in the second stage. They can incur costly modifications that improve fuel economy enough to get over one notch, or they can make no changes (at zero additional cost).[28]

We denote by x the fuel economy improvement in the second stage, and its cost, which varies across j, is labeled $c_j(x_j)$. We denote the lowest cost fuel economy improvement that moves vehicle j at least z as x_j^*. We allow the cost function for a particular vehicle to be non-monotonic because actions are discrete, which explains why bunching may occur at values other than the exact notch (i.e., $x_j^* > z$). For example, if the cheapest way to improve fuel economy is to add a spoiler to a particular vehicle, and this improves fuel economy by .2 mpg, it will be more expensive to improve fuel economy by .1 than by .2, and a vehicle located .1 away from a notch may end up "overshooting." The cost of modifying a vehicle is net of any decrease (or increase) in consumer valuation.

We assume that there are analogous actions that can improve fuel economy at a given cost in the first stage, which are drawn from a different set of modifications. In online appendix A we show that, under the intuitive assumption that first-stage decisions do not impact the expected value of a vehicle's distance to the nearest notch at the start of the second stage, both a linear tax and a notched tax deliver identical first-stage behavior. This means that the only difference between a notched policy and a smooth one comes from differences in second-stage behavior. We thus focus on second-stage behavior for the remainder of the section.

A vehicle will be modified in the second stage if the subsidy increase from reaching a notch exceeds the cost. That is, a producer will improve a vehicle's fuel economy by x_j^* if $t > c_j(x_j^*)$, and otherwise they will make no improvement. It is useful to rewrite this condition in terms of t_{xz}, the *per-mpg* subsidy for a vehicle z units from the notch whose $x_j^* = x$. Written thusly, a producer will improve a vehicle's fuel economy by x_j^* if and only if $t_{xz} \cdot x_j^* > c_j(x_j^*)$.

Because the external benefit of fuel economy is a constant value e, social efficiency requires that fuel economy be improved if and only if $e \cdot x_j^* > c_j(x_j^*)$. This will occur if $t_{xz} = e$ for all x_j^* and z, which is what would occur under a linear tax. In contrast, a notched tax necessarily creates inefficiency in second-stage behavior because notches imply that t_{xz} will be different for different vehicles, depending on their values of x_j^* and z, which means that the privately optimal and socially efficient rules must differ for at least some vehicles.

We illustrate this in Table 4, which shows a complete listing of the per-mpg subsidies t_{xz} for a notched subsidy that, like the Gas Guzzler Tax, rises by $700 at every fuel economy rating ending in .5. Rows indicate the decimal of a vehicle's mpg rating before any fuel economy
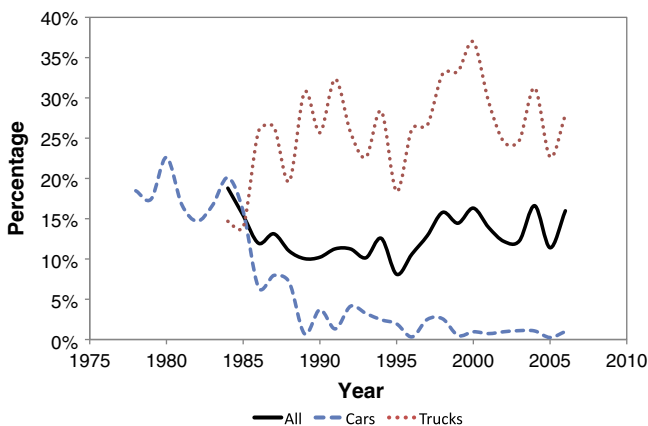

Fig. 11. Market share of vehicles rated between 20 and 22.5 mpg: Passenger cars, light trucks and all vehicles. Source: Authors' calculations of NHTSA data.

---

[27] We discuss this modeling assumption in more detail in appendix A available online.
[28] We assume that any negative cost (privately beneficial) design features are adopted in the first stage.

Table 4
Effective average subsidy per mpg, before and after second-stage manipulation.

Decimal at beginning of second stage ($= .5 - z + .5 \cdot 1(z \geq .5)$)	Decimal at end of second stage ($= z + x - 1 \cdot 1(z + x \geq 1)$)									
	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
0	X	0	0	0	0	1400	1167	1000	875	778
0.1	778	X	0	0	0	1750	1400	1167	1000	875
0.2	875	778	X	0	0	2333	1750	1400	1167	1000
0.3	1000	875	778	X	0	3500	2333	1750	1400	1167
0.4	1167	1000	875	778	X	7000	3500	2333	1750	1400
0.5	0	0	0	0	0	X	0	0	0	0
0.6	0	0	0	0	0	778	X	0	0	0
0.7	0	0	0	0	0	875	778	X	0	0
0.8	0	0	0	0	0	1000	875	778	X	0
0.9	0	0	0	0	0	1167	1000	875	778	X

Note: Table shows effective subsidy per mpg of a vehicle that starts at the fuel economy decimal in each row value and ends, after manipulation, in the column value. Table assumes a notch value of $700 (the Gas Guzzler Tax average) and includes "wrap-around" values, but assumes that all jumps are less than 1 mpg.

improvement that occurs in the second stage, and the column represents the decimal after the improvement. The table allows "wrap arounds" but assumes that all improvements are less than 1 full mpg. For example, a vehicle that begins at 14.8 mpg and improves to 15.3 mpg would face the subsidy listed in the ninth row and fourth column, which is $0 because such a change would not cross a notch. In contrast, a vehicle at 16.3 (fourth row) that improves to 16.7 (eighth column) would gain a subsidy equivalent to $1750 per-mpg ($700/.4 mpg). This latter subsidy is much larger than the social gain of $700 per-mpg. A smooth linear tax would have a single value for all t_{xz}, and every cell of Table 4 would be $70. The notched subsidy achieves the right "average" subsidy across all vehicles, but it does so by offering a too large subsidy to some vehicles and a too small one to others.

5.2. The welfare consequences of producer choice

A notched policy creates a schedule of second-stage per-mpg subsidies that vary across vehicles depending on their x and z. To assess the welfare impact of such a schedule, we need to characterize the welfare gain of second-stage manipulation conditional on values of x_j^* and z and characterize the distribution of vehicles across the various combinations of x_j^* and z. We assume there is heterogeneity in cost $c_j(x_j^*)$ across vehicles that face the same x_j^* and z. Specifically, we denote the cumulative distribution of these costs by $F_{xz}(c)$ and the overall mass of such vehicles as N_{xz}. The probability that a vehicle will exercise its option to improve mpg in the second stage is equal to the probability that the net cost is below the tax value, or $F_{xz}(t_{xz} \cdot x)$, and all vehicles with $c_j(x_j^*) \in (0, t_{xz} \cdot x_j^*)$ will undertake modifications.

The net second-stage social gain SG across all vehicles facing the same x^* and z can therefore be written as the integral of the individual social gains for net costs between 0 and $t_{xz} \cdot x^*$:

$$SG_{xz} = N_{xz} \int_0^{t_{xz} \cdot x^*} \left(e \cdot x_j^* - c_j\left(x_j^*\right) \right) dF_{xz}\left(c_j\right). \quad (1)$$

For tractability, we assume that F_{xz} is uniform with support on 0 to $\gamma > t_{xz} \cdot x^*$. This implies that the cost curve aggregated across vehicles that share values of x^* and z is linear. In that case, the net social gain of second-stage manipulation of all the models that are z distance from the nearest notch who have $x^*(z_j)$ is:

$$SG_{xz} = 1/2 \cdot (2e - t_{xz}) \cdot x^* \cdot \frac{t_{xz} x^*}{\gamma} \cdot N_{xz} \quad (2)$$

$$\equiv 1/2 \cdot (2e - t_{xz}) \cdot X, \quad (3)$$

where X is defined as the total change in fuel economy that is induced by tax rate t_{xz}.[29]

The social gain is maximized when $t_{xz} = e$, which is the familiar result from Pigou (1932) and which would occur under a smooth tax $t = e$. A subsidy below e will fail to induce all of the second-stage manipulations that have a net social benefit, whereas a subsidy above e will induce second-stage manipulations that have net social costs. This is shown graphically in Fig. 12, where the figure on the left shows the welfare gain from a per-mpg subsidy below e and the figure on the right shows the opposite case. Each figure represents a set of vehicles that share a common x^* and z. The horizontal axis is the total increase in fuel economy across vehicles with a given x^* and z, which aggregates over the discrete changes made by each vehicle.

The triangles in Fig. 12, which are directly analogous to Harberger triangles, have a net area of $1/2 \cdot (2e - t_{xz}) \cdot X$.[30] In the figure on the right, when $t_{xz} > e$, the lower blue (solid) triangle represents a welfare gain, while the upper red (striped) triangle is a welfare loss created by vehicles being modified that have a private cost larger than the total benefit, but an even larger private benefit via the subsidy (because the small fuel economy improvement moves the vehicle over a notch). When the red triangle is larger than the blue triangle (which will occur whenever $t_{xz} > 2e$), the net social benefit of second-stage manipulation for all vehicles that share an x^* and z will be negative.

Given Eq. (3), we can characterize the second-stage net social gain from a notched incentive, for the vehicles facing a particular x^* and z. However, the data do not tell us the values of x^* and z. Instead, we only observe the ex post distribution of decimal ratings around the notches, which includes some vehicles whose fuel economy was strategically adjusted to the incentive and probably many whose fuel economy did not respond. Next, we show that it is possible to calculate the total net social gain from second-stage manipulation across all vehicles using only ex post data on the final distribution of vehicles vis-à-vis notches and an assumption regarding the counterfactual distribution of decimals that would occur if there were no second-stage manipulation (the distribution of decimals at the start of the second-stage).

The net social gain across all vehicles in a notched system is simply the sum of Eq. (3) across all of the pairs of x^* and z, weighted by the amount of manipulation in each of those markets:

$$\begin{aligned} SG &= 1/2 \sum_i (2e - t_i) \cdot X_i \\ &= e \sum_i X_i - 1/2 \sum_i t_i X_i \\ &= 1/2 \cdot (2e - \tau) X, \end{aligned} \quad (4)$$

---

[29] To see this, note the last three terms on the right hand side of Eq. (2) are the mpg gain per unit that makes a second-stage improvement (x^*), the proportion of vehicles who respond at the given tax level ($t_{xz} \cdot x^*/\gamma$), and the size of the market (N_{xz}).

[30] Here then, we see visually that our distributional assumption regarding F_{xz} produces a linear aggregate cost of improved fuel economy across a set of vehicles with a common x^* and z.
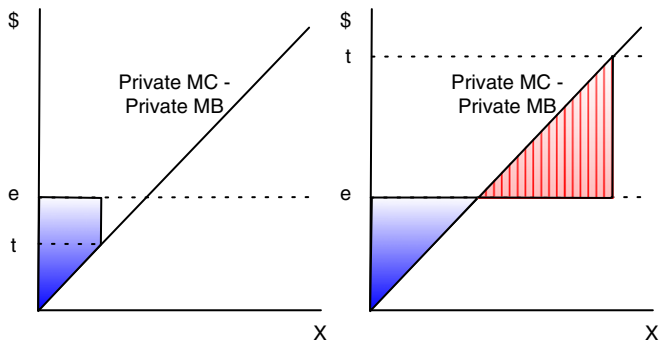
Fig. 12. Welfare gains at varied per-mpg subsidy rates. Note: The blue (solid) shaded area in the figure on the left is the welfare gain from a subsidy below the Pigouvian optimum. The figure on the right shows a subsidy in excess of the Pigouvian tax, with the blue (solid) area representing a social gain and the red (striped) area a social loss. Diagrams represent the entire market of vehicles facing a particular per-mpg subsidy, which share a common x^* and z. The smooth MC–MB curves come from averaging discrete second-stage modification choices across many vehicles.

Table 5
Distribution of Gas Guzzler Tax rating decimals.

Decimal	Observed distribution	Uniform counterfactual
.0	0.1023	0.1
.1	0.0921	0.1
.2	0.0989	0.1
.3	0.0867	0.1
.4	0.0738	0.1
.5	0.1260	0.1
.6	0.1152	0.1
.7	0.1206	0.1
.8	0.0840	0.1
.9	0.1003	0.1

Note: Sample matches Fig. 4, with data from 1980–2009. IRS data; sample size is 1476.

where i indexes a specific combination of x^* and z, $X \equiv \sum_i X_i$ is the total change in fuel economy, and $\tau \equiv \sum_i t_i \frac{X_i}{X}$ is the X_i-weighted *average effective tax rate*. Eq. (4) shows that the size of this parameter determines the total welfare from the second-stage responses to the notched incentives. If $\tau > 2e$ then the net benefit will be *negative*. Because vehicles that have a larger per-mpg tax subsidy will be more likely to perform second-stage modifications, there will be a positive correlation between X_i and t_i. As a result, τ will tend to be much larger than the simple average of t_i. The notched system sets the simple average of t_i equal to e, but τ will be larger. How large it is determines the net social gain.

To calculate τ, one needs to know how many vehicles faced different manipulation options x^*. It is reasonable to make assumptions about the distribution of z – the proximity of observations to the notch in the absence of manipulation – but we do not have an empirical basis for asserting precisely how many vehicles faced different manipulation options x^*. Consider Table 5, which shows the data underlying Fig. 4, the aggregated distribution of decimal ratings for vehicles subject to the Gas Guzzler Tax. We can assume a uniform counterfactual distribution, which is logical because, in the absence of second-stage manipulation in response to notch locations, there is no reason to suspect that certain decimals would appear more frequently than others. The data suggest that some vehicles moved from having a .3 or .4 decimal value to having a .5 or .6. But, we do not know exactly "how" that happened – whether some vehicles moved from .3 to .6 or all moved from .3 to .5, for example.[31]

When the aggregate cost curves for all vehicles facing a given x^* and z are linear (which is implied when $F_{xz}(c)$ is uniform), Eq. (4) is invariant to the mapping that leads to the final distribution. In this case, the ex post distribution is sufficient to calculate Eq. (4) because all mappings lead to the same social gain. Because our data show that almost exactly 30% of all observations fall between decimals .0 and .2 and almost exactly 30% fall between .7 and .9, we assume that all second-stage manipulation took place among vehicles starting at .3 or .4 and moving to .5 or .6.[32] We calculate the total change in average fuel economy, X, as the difference between the average fuel economy according to the actual fuel economy decimals in Table 5 and the average fuel economy under the uniform distribution, times the total market size N. Because τ is invariant to the mapping we pick to move between the counterfactual distribution and the observed data, we pick one mapping

to calculate that $\tau = \$4,170$.[33] Inserting this into Eq. (4) yields $SG = (700 - \frac{1}{2}4170)(.00663N) = -9.18N$.

This calculation suggests that the second-stage incentives of the notch generate a social *loss* of \$9.18 per car from second-stage manipulation. This is because the average effective subsidy rate of \$4170 per-mpg is much larger than the social benefit of \$700 per-mpg. As long as this weighted average subsidy is greater than two times the externality (\$1400), the net impact of the second-stage manipulation will be negative.[34]

5.3. Decomposition of social gain

Eq. (4) calculates the net social gain of all second-stage manipulation, some of which is efficient. Another insightful exercise compares this social gain to the social gain from a smooth subsidy, which would induce all of the efficient modifications and none of the inefficient ones. This requires first decomposing our social gain formula into the efficient (the blue triangle in Fig. 12) and the inefficient (red triangle in Fig. 12) portions. Elementary geometry provides the decomposition; for a given $t > e$, which induces both efficient and inefficient behavior:

$$ SG = \frac{1}{2}(2e-t)\Delta X = \underbrace{\frac{1}{2}\frac{e^2}{t}\Delta X}_{gain} - \underbrace{\frac{1}{2}\left(1-\frac{e}{t}\right)(t-e)\Delta X}_{loss}, \quad (5) $$

where the first term is the gain from efficient manipulation and the second term is the loss from inefficient manipulation. Just as with Eq. (3), each term can be summed across the various t to create a single statistic but, unlike the other measure, the decomposed summations are not invariant to the mapping between starting and ending decimal values.

In our case, the lack of invariance is easily solved through reasonable auxiliary assumptions. If we assume that mpg is manipulated no more than .2 in response to the notch (which is consistent with the fact that there is no shortage of observations at .2), then it must be that all of the vehicles that started at .3 ended at .5. This pins down the distribution of ending decimals for those starting at .4, and we can decompose our social gain into the efficient and inefficient portions, which reveals that the efficient gain is only \$.42N, but the inefficient loss is -\$9.60N.

---

[31] We do not need to know x^* for vehicles that did not make second-stage changes because their weight in Eq. (4) will be zero.

[32] We make a very small adjustment to the ex post value at .6 in order to make this assumption hold exactly.

[33] Specifically, we assume that 1.33% of all observations moved from a .3 ending decimal to a .5; 1.27% moved from .4 to .5; and 1.35% moved from .4 to .6. Thus $\tau = (.2*.0133/X)*3500 + (.1*.0127/X)*7000 + (.2*.0135/X)*3500 = \$4,170$, where $X = .2*.0133 + .1*.0127 + .2*.0135 = 0.00663$.

[34] We conjecture that, because there will be a positive correlation between effective subsidy rates and manipulation, τ will exceed the average statutory t, which implies that t, the notch policy parameter, should be set below the Pigouvian tax that would prevail in a system without notches. This conclusion, however, will depend on how the different notch size influences first-period behavior.

This decomposition puts the net social gain of -9.18N$ into perspective. The net social loss from the second-stage manipulation in response to the notched subsidies is not only negative, but it is also twenty times as large as the positive benefit from the affected vehicles. Because deadweight loss rises with the square of t and gains are zero when $t = 2e$, it is intuitive that effective subsidies of $7000 and $3500, which are 10 and 5 times the externality, yield very large inefficiency losses relative to gains. The inefficiency of the notched subsidy compared to the smooth subsidy is even worse if one considers the vehicles that face a zero subsidy under the notch, for which socially efficient improvements are not made.

5.4. Second-stage behavior for a smooth subsidy

In contrast to the notched case, when facing a smooth subsidy the producer's second-stage choice is to pick the fuel economy improvement that maximizes the difference between tax gain and loss. That is, a firm will choose x to maximize $t \cdot x - c_j(x)$ for each j. The socially optimal behavior would maximize the difference between social gain and private cost, which is $e \cdot x - c_j(x)$. Thus, the firm will choose the socially optimal x whenever $t = e$. When the policy is smooth, *all* vehicles will face the same subsidy per-mpg, so it is possible to give all j the correct incentive simultaneously. This contrasts with the notched policy, where the local incentives necessarily vary widely.

The decomposed social gain from the vehicles starting at the .4 decimal in the data is 0.23N$.[35] If all vehicles faced the smooth incentive, then the total benefit of second-stage manipulation would be ten times the gain for those at .4, or 2.30N$. Thus, while the smooth subsidy would generate a $2.30 net second-stage social benefit per unit sold, the notched subsidy creates a $9.60 net second-stage social loss per unit. If notches exist for administrative or salience reasons, these benefits must be large enough to offset costs of this magnitude.

5.5. Caveats and discussion

The welfare calculations presented here are second-stage calculations, which show how welfare is lost when agents react to the notches rather than a smooth Pigouvian subsidy. Our analysis shows the net social welfare consequences of second-stage behavior in response to a notched policy, on its own and compared to the net social gain in the second-stage from a smooth policy. We cannot calculate the social gain from first-stage manipulation, but we can say that the gain in the first stage is the same for a notched and smooth policy (see online appendix A). Without first-stage effects, we cannot say whether or not society would be better off with no policy at all, as opposed to a notched one. But, that is not our goal. Our goal to establish that there is a social cost incurred when policy makers use a notched schedule to approximate a linear one.

Our model assumed that a flat $700 per-mpg subsidy was the first-best policy, but this is not true because fuel economy policies are never first-best solutions on their own. The externalities related to fuel economy are direct functions of fuel consumed (which causes local air pollution, greenhouse gases, and energy insecurity) and miles driven (which causes traffic accidents and congestion). Improved fuel economy lowers the amount of fuel consumed per mile driven, but it also lowers the cost of driving, which induces extra miles driven. These considerations conflict, and if externalities related to extra miles driven are sufficiently large, the net social impact of increasing fuel economy can be negative. For these reasons, fuel economy taxation is not an efficient policy on its own (Sallee, 2011b), an intuition which closely parallels the literature on fuel economy regulation, which is surveyed in Harrington et al. (2007) and Anderson et al. (2011).

Even if we abstract from mileage responses (which are unlikely to play a major role in the ultra-luxury market), gasoline-related externalities are approximately a linear function of fuel consumption (gallons-per-mile), not fuel economy (miles-per-gallon). Thus, in the absence of a mileage response, if all vehicles were driven the same lifetime mileage, then the first-best policy would be a constant tax on gallons-per-mile (gpm), which implies a quadratic subsidy to fuel economy.[36] The shape of the Gas Guzzler Tax – i.e., how the notch size varies across mpg notches – does approximate a constant gpm shape, but it is imperfect. (Relative to the other notches, the notch size at 22.5 is too large, and the notch sizes at 17.5, 18.5 and 19.5 are too small.) Thus, it might be preferable to perform welfare calculations based on constant per gpm taxes, but the IRS data we use are measured in discrete increments of one-tenth of an mpg, not in gpm, and translating between them would necessarily force us to make fairly arbitrary judgment calls or extrapolations about where to draw gpm bins.[37]

The goal of our welfare exercise is to demonstrate how to model the welfare impacts of substituting a notched schedule for a linear one. The qualitative lessons from our exercise would remain if it was recast as a gpm analysis. The problem with the notched subsidy is that it creates uneven incentives for improvement given the proximate location of specific vehicles to notches. It is the heterogeneity in incentives across agents facing the same notch that creates inefficiency. Given that, our emphasis is on building a useful framework for other researchers to adapt, and we consider the heuristic benefit of abstracting from concerns about nonlinearity and the mileage responses to be worth the cost.

6. Bunching in fuel economy label ratings

Automakers face fuel economy policy notches not only in the form of tax incentives, but also in the form of fuel economy labels for consumers. Automakers are required by federal law to attach a fuel economy label to all new vehicles, and the contents of this label are strictly prescribed. The values must be reported as integers formed by rounding off test results. This results in a *presentation notch* at every .5 mpg, the rounding cutoff. If consumers value fuel economy and consider the rounded integers when shopping, then automakers have an incentive to manipulate fuel economy ratings around these presentation notches. Whereas only high-performance passenger cars are subject to the Gas Guzzler Tax, all vehicles have labels, which allow us to test for bunching throughout the entire market. The existence of bunching (which we find below) around fuel economy labels implies that vehicle-makers believe that consumers value fuel economy, which is a topic of some controversy in the environmental policy literature.[38]

To test for bunching, we generate histograms similar to the ones presented above for the Gas Guzzler Tax. Because the EPA's publicly available data files do not include the unrounded estimates from 1984 to 1997, we are limited to data from years before and after

---

[35] This follows from using the first half of Eq. (5) on vehicles estimated to have begun at .4 and moved to .5 and .6: $(1/2)*[700^2/3500]*(.2)*.0135 + (1/2)*[700^2/7000]*(.1)*.0127 = .23$.

[36] The total tax should be linear in gallons consumed g, as in $tax = e \cdot g = e \cdot M/mpg$, where e is the externality per gallon and M is miles driven. The derivative of the total tax with respect to mpg is then quadratic, $-e \cdot M/mpg^2$.

[37] Also, regardless of the form of the tax, one might ask whether or not it has the right overall level. As we discuss in online appendix B, the tax is probably too low for the majority of vehicles subject to the tax.

[38] There is mixed evidence on whether or not consumers properly value fuel economy when purchasing a vehicle, which has important policy implications for the efficiency of market mechanisms relative to regulatory tools for promoting increased fuel economy. See Greene (2010) and Helfand and Wolverton (2009) for recent reviews. In an earlier version of this paper, we compared the amount of bunching around the Gas Guzzler Tax notches (where the economic value is known) to the amount around labels in order to infer how much the consumer willingness to pay for a vehicle increases when its fuel economy rating rises. Our estimates implied large valuations of fuel economy, with most exceeding a full valuation benchmark. We have omitted this analysis here because it rests on the strong assumption that the cost of manipulation is similar across vehicles facing the Gas Guzzler Tax and all other vehicles.
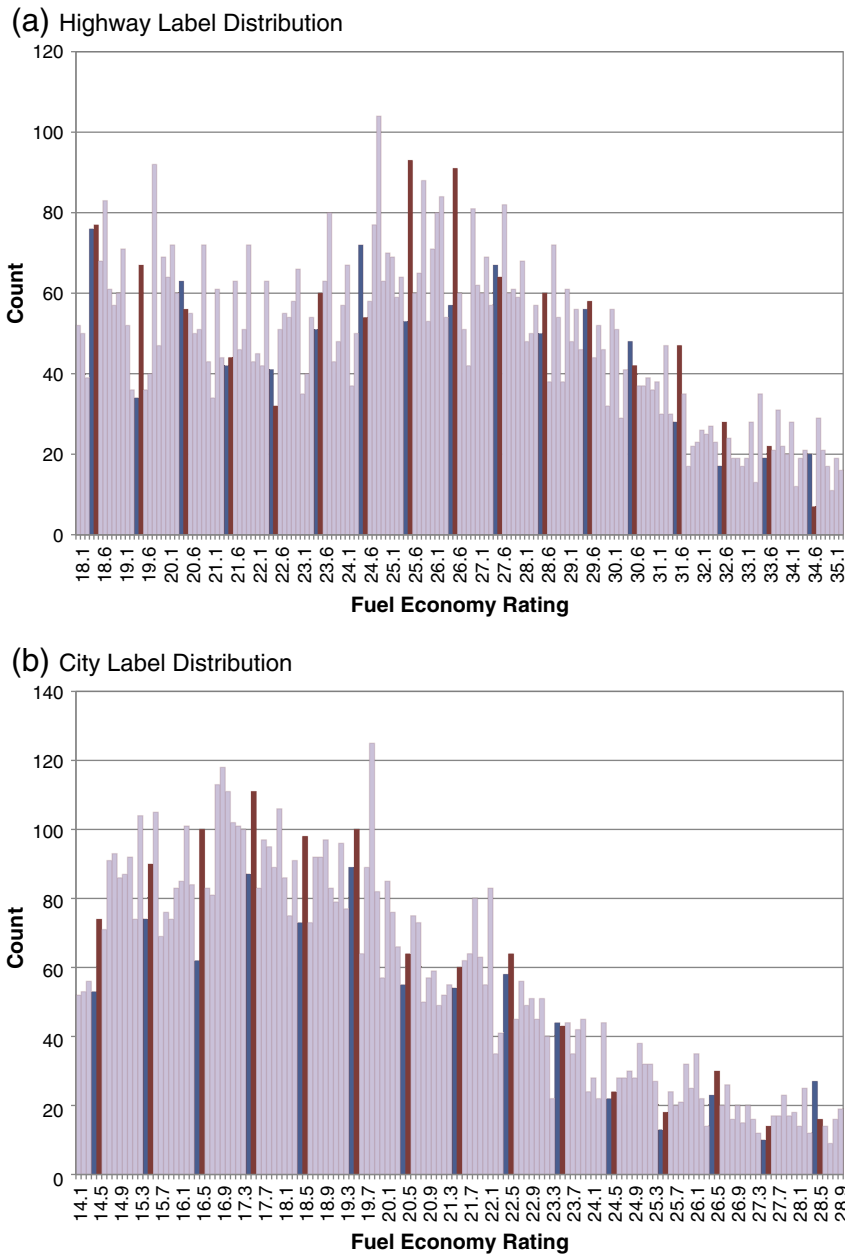
(a) Highway Label Distribution

(b) City Label Distribution

Fig. 13. Label distribution: 1978–1983 and 1998–2007. Note: Data come from the EPA. Sample excludes vehicles that run on alternative fuels and observations where the manufacturer, cylinders, displacement, drive type, fuel type, transmission and fuel economy ratings are identical within a model year.

this period. Recall that, although these mpgs are generated from the same underlying tests as the mpgs for the Gas Guzzler Tax, the test statistics are transformed in different ways for the two purposes, so a vehicle near a notch in label would not be more likely to be near a Gas Guzzler Tax notch. Fig. 13a and b show, respectively, histograms of highway and city fuel economy label ratings decimals for all vehicles.[39] The cutoff for rounding to the nearest integer is .5, so we

expect to see bunching at .5 (where the rating is rounded up) relative to .4 (where the rating is rounded down). Ratings ending in .4 are colored in blue, and those ending in .5 are colored in red. (For those reading in black and white, for each pair of shaded values, the one on the left is blue and the one on the right is red.)

Fig. 13a shows that there are more observations just above notches than just below in highway ratings, but there are several integer values where the reverse is true. Fig. 13b, though, shows consistent and large bunching above notches in city label ratings. Fig. 14a and b repeat this exercise for the "Big Three" domestic automakers – Chrysler, Ford and General Motors. Bunching just over presentation notches is amplified in this subsample.

Table 6 shows statistical significance tests for this bunching, repeating the analysis used in Table 2 for analysis of the Gas Guzzler Tax. All of the differences are statistically significant at any conventional level,

---

[39] The tabulations in these, and all subsequent figures in this section, exclude a very small number of vehicles with unusual fuel types (e.g., compressed natural gas) that are subject to a different rating procedure, and they drop any observation that appears to be a repeat in the sample such that the manufacturer, cylinders, displacement, drive type, fuel type, transmission and fuel economy ratings are identical within a model year, with the intention of restricting identical engines that are included in multiple models.
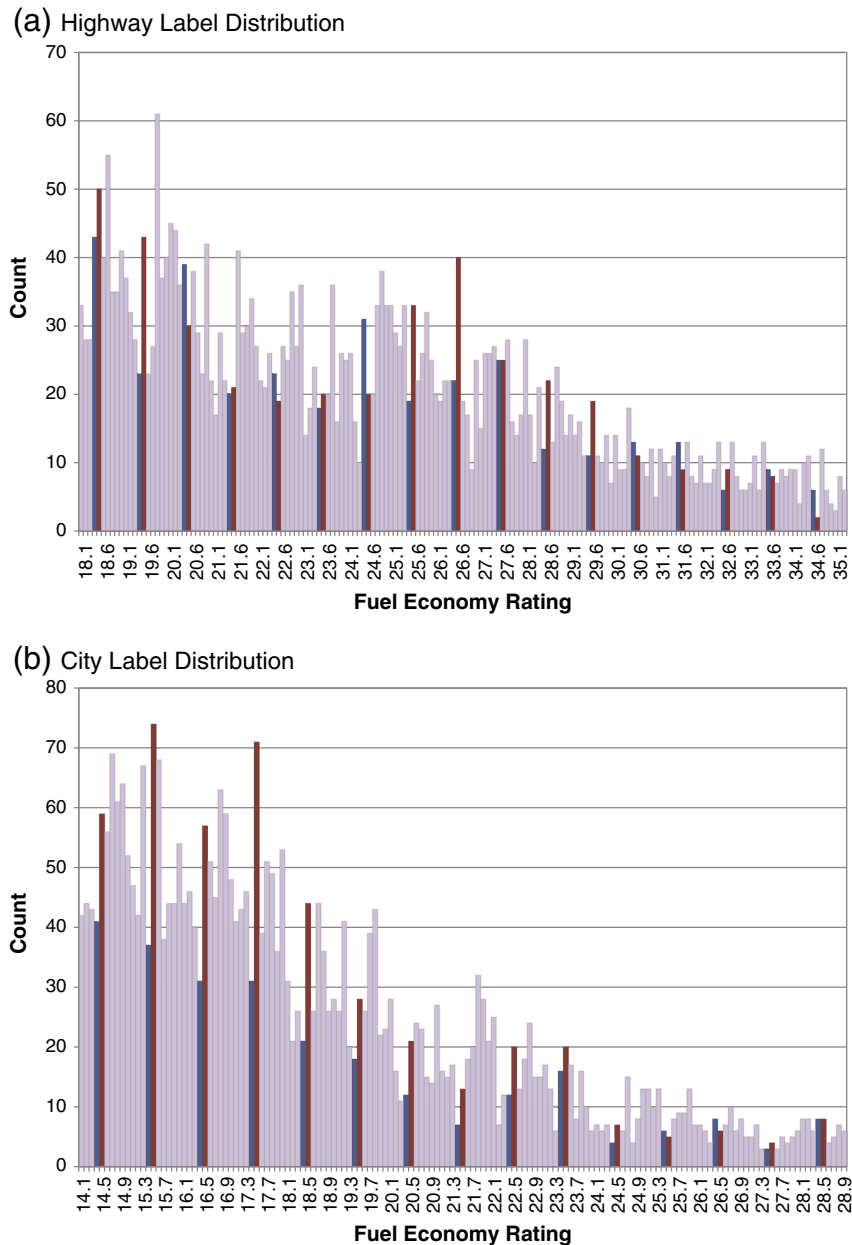
(a) Highway Label Distribution

(b) City Label Distribution

Fig. 14. Label distribution, Big Three: 1978–1983 and 1998–2007. Note: Data come from the EPA. Sample excludes vehicles that run on alternative fuels and observations where the manufacturer, cylinders, displacement, drive type, fuel type, transmission and fuel economy ratings are identical within a model year.

and adjustments for the overall shape of the fuel economy distribution make very little difference in the significance tests.[40]

Figs. 15 and 16 show the decimal distributions aggregated across integers, for the full sample and just the Big Three, in the late and early year samples separately. These figures show dramatic bunching in the city ratings in both time periods, and there is some evidence of bunching in the highway rating in recent years. Where there is bunching in the overall distribution, it is always greater when the sample is restricted to the Big Three.

There are several reasons why domestic automakers may be relatively more responsive to presentation notches. First, domestic automakers sell a much larger fraction of their vehicles to the U.S. market. Foreign producers may be reluctant to fine-tune vehicles for this purpose if only a modest fraction will be shipped to the U.S. Second, domestic automakers often face a binding CAFE constraint, which they meet through various strategic adjustments, including fuel economy tweaking. In contrast, European automakers are often out of compliance with CAFE and pay fines rather than adjust their fleet, and Asian automakers are generally well above the minimum mpg required by CAFE (Anderson and Sallee, 2011). As a result, the domestic automakers

---

[40] The appendix includes similar regressions based on statistical tests based on aggregated data –where the fuel economy rating is the unit of observation instead of the vehicle – which confirm the findings here. The appendix also presents placebo tests for bunching around decimal values where there are not notches. The appendix also repeats this analysis with standard errors clustered on manufacturer interacted with number of cylinders. Generally, all of these results are supportive of the conclusions drawn here.

Table 6
Statistical tests of bunching above label notches.

	Observed number above notch(total near notch)	Expected number above notch under null hypothesis	Standard deviation of number above notch under null hypothesis	p-value
Highway ratings (Fig. 13a)				
Binomial model	791 (1424)	712	18.87	0.00003
Linear control function	791 (1424)	709.6	18.64	0.00001
Quintic control function	791 (1424)	710.2	18.70	0.00002
City ratings (Fig. 13b)				
Binomial model	742 (1382)	691	18.59	0.00607
Linear control function	742 (1382)	690.5	19.20	0.00731
Quintic control function	742 (1382)	689.8	19.16	0.00644
Highway ratings, Big Three (Fig. 14a)				
Binomial model	290 (492)	246	11.09	0.00007
Linear control function	290 (492)	245.6	11.02	0.00006
Quintic control function	290 (492)	245.0	11.00	0.00004
City ratings, Big Three (Fig. 14b)				
Binomial model	296 (464)	232	10.77	<0.00001
Linear control function	296 (464)	232.2	10.17	<0.00001
Quintic control function	296 (464)	231.8	10.18	<0.00001

Note: The binomial model assumes that the probabilities that a vehicle has a fuel economy rating ending in .4 and .5 are the same. The linear model modifies this assumption by adjusting for the overall shape of the distribution using a linear fit, excluding observations within .1 mpg of a notch. The quintic model extends this by using a fifth order polynomial to estimate the distribution's shape. Statistics for the linear and quintic models are derived via bootstrapping to account for sampling variation in the estimated polynomial.
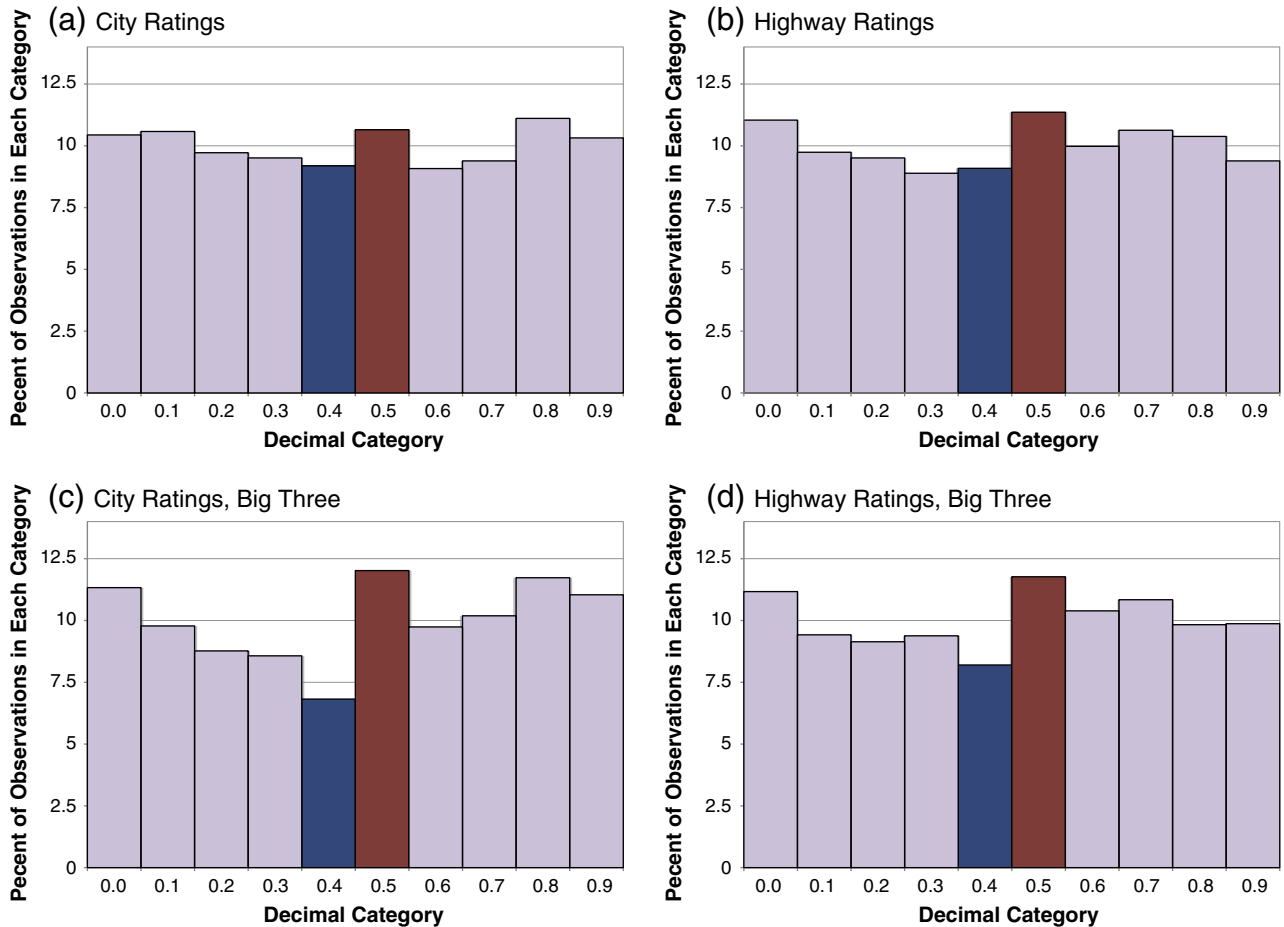

Fig. 15. Label rating decimals: 1999–2007. Note: Data come from the EPA. Sample excludes vehicles that run on alternative fuels and observations where the manufacturer, cylinders, displacement, drive type, fuel type, transmission and fuel economy ratings are identical within a model year.
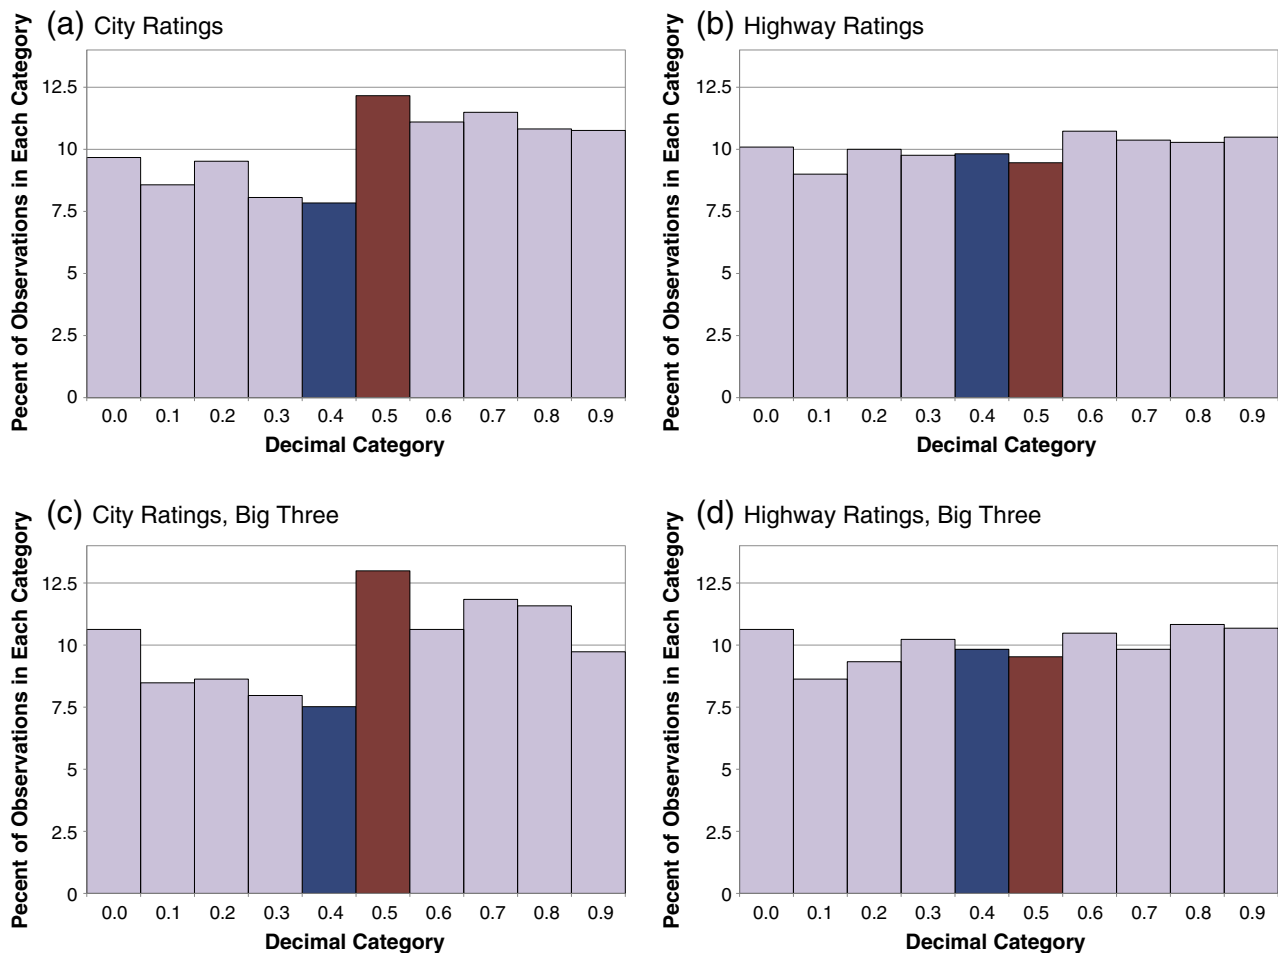
Fig. 16. Label rating decimals: 1978–1983. Note: Data come from the EPA. Sample excludes vehicles that run on alternative fuels and observations where the manufacturer, cylinders, displacement, drive type, fuel type, transmission and fuel economy ratings are identical within a model year.

may have developed greater expertise in finely tuning fuel economy to meet CAFE standards. For example, until very recently, only domestic firms made use of a CAFE loophole for flexible-fuel vehicles (Anderson and Sallee, 2011). Third, relative to American consumers of European cars, American consumers of domestic cars may be more concerned with fuel economy and, relative to the Asian automakers, domestic firms may be more concerned with boosting their fuel economy image.

Finally, Fig. 17 separates the sample into passenger cars and light trucks, restricting the data to the Big Three. This figure suggests that the city rating shows more bunching for light trucks than cars, whereas the highway rating shows more bunching for cars than trucks. This could be due to the different uses of trucks and cars on average — truck consumers may be more concerned about city fuel economy ratings that reflect the typical usage patterns of larger vehicles, whereas car buyers are more concerned about highway driving and commuting.

7. Conclusion

Key aspects of vehicle fuel economy policy are designed with notches, so that for many vehicles there is no incentive to incrementally improve fuel economy, but for others there are large and varying incentives for improvement. In this paper we show that the policy notches have real consequences, as there are significantly more vehicles produced (and purchased) just on the policy-beneficial side of

the notches than otherwise would be expected. We observe this behavior not only in response to explicit notches in tax and subsidy policies, but also in response to implicit "presentation notches," where government policy dictates what (coarse) information a firm must provide to consumers. We develop a simple framework within which the negative welfare effects of local manipulation can be calculated, a framework which may prove useful in a variety of contexts as it can be utilized with only ex post data.

Future fuel economy policies are likely to increase the importance of notches. The state of California recently explored a comprehensive feebate program with notches, and similar legislation has been introduced in the U.S. Senate. The EPA has considered the adoption of a grading system that would assign notched letter grades to vehicles by fuel economy and emissions. Recent CAFE reforms have dramatically tightened its standards in a way that increases the value of moving vehicles over the light-truck classification notch. Policy notches may have administrative or salience benefits, but for notches to be warranted, these benefits must outweigh the substantial inefficiency costs that we document here.

Appendix A. Supplementary data

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.jpubeco.2012.06.005.
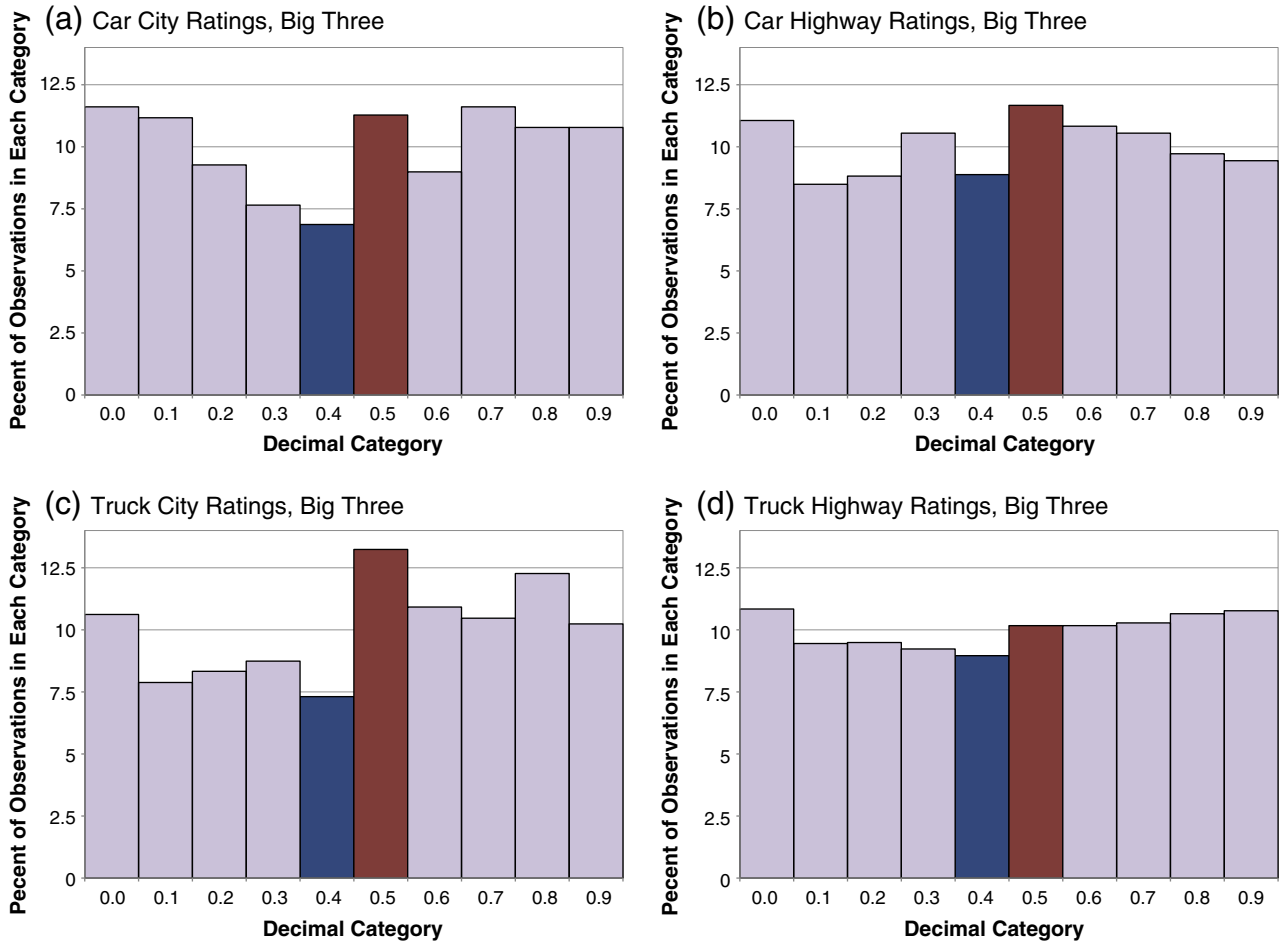
Fig. 17. Label rating decimals: 1978–1983 and 1998–2007. Note: Data come from the EPA. Sample excludes vehicles that run on alternative fuels and observations where the manufacturer, cylinders, displacement, drive type, fuel type, transmission and fuel economy ratings are identical within a model year.

References

Anderson, S.T., Sallee, J.M., 2011. Using loopholes to reveal the marginal cost of regulation: the case of fuel-economy standards. The American Economic Review 101, 1375–1409.

Anderson, S.T., Parry, I., Sallee, J.M., Fischer, C., 2011. Automobile fuel economy standards. Review of Environmental Economics and Policy 5, 89–108.

Blinder, A.S., Rosen, H.S., 1985. Notches. The American Economic Review 75, 736–747.

Blundell, R., 2000. Work incentives and 'in-work' benefit reforms: a review. Oxford Review of Economic Policy 16, 27–44.

Greene, D.L., 2010. How consumers value fuel economy: a literature review. Technical Report EPA-420-R-10-008. U.S. Environmental Protection Agency.

Harberger, A.C., 1995. Tax lore for budding reformers. In: Dornbusch, R., Edwards, S. (Eds.), Reform, Recovery and Growth: Latin America and the Middle East. University of Chicago Press, Chicago.

Harrington, W., Parry, I., Walls, M., 2007. Automobile externalities and policies. Journal of Economic Literature 45, 373–399.

Helfand, G., Wolverton, A., 2009. Evaluating the consumer response to fuel economy: a review of the literature. National Center for Environmental Economics Working Paper No. 09-04.

Kleven, H.J., Slemrod, J., 2009. A characteristics approach to optimal taxation and tax-driven product innovation. Working Paper: London School of Economics.

Kleven, H.J., Waseem, M., 2012. Behavioral responses to notches: Evidence from Pakistani tax records. Manuscript: London School of Economics.

Klier, T., Linn, J., 2010. New vehicle characteristics and the cost of the corporate average fuel economy standard. Resources for the Future Discussion Paper 10-50.

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

National Research Council, 2002. Effectiveness and Impact of Corporate Average Fuel Economy (CAFE) Standards. The National Academies Press, Washington, D.C.

Pigou, A.C., 1932. The Economics of Welfare, 4th edition. Macmillan and Co., London.

Ramnath, S., 2009. Taxpayer responses to notches: Evidence from the saver's credit. Manuscript, University of Michigan.

Saez, E., 2010. Do taxpayers bunch at kink points? American Economic Journal: Economic Policy 2, 180–212.

Sallee, J.M., 2011a. Rational inattention and energy efficiency. Manuscript: University of Chicago.

Sallee, J.M., 2011b. The taxation of fuel economy. Tax Policy and the Economy 25, 1–37.

Slemrod, J., 2010. Buenas notches: lines and notches in tax system design. Working Paper, University of Michigan.