

Gillingham, Kenneth; Houde, Sébastien; van Benthem, Arthur A.

Working Paper

Consumer myopia in vehicle purchases: evidence from a natural experiment

CESifo Working Paper, No. 7656

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Gillingham, Kenneth; Houde, Sébastien; van Benthem, Arthur A. (2019) : Consumer myopia in vehicle purchases: evidence from a natural experiment, CESifo Working Paper, No. 7656, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/201882>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Consumer myopia in vehicle purchases: evidence from a natural experiment

Kenneth Gillingham, Sébastien Houde, Arthur A. van Benthem

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Consumer myopia in vehicle purchases: evidence from a natural experiment

Abstract

A central question in the analysis of fuel-economy policy is whether consumers are myopic with regards to future fuel costs. We provide the first evidence on consumer valuation of fuel economy from a natural experiment. We examine the short-run equilibrium effects of an exogenous restatement of fuel-economy ratings that affected 1.6 million vehicles. Using the implied changes in willingness-to-pay, we find that consumers act myopically: consumers are indifferent between \$1 in discounted fuel costs and 15-38 cents in the vehicle purchase price when discounting at 4%. This myopia persists under a wide range of assumptions.

JEL-Codes: D120, H250, L110, L620, L710, Q400.

Keywords: fuel economy, vehicles, myopia, undervaluation, regulation.

Kenneth Gillingham
Yale University
195 Prospect Street
USA – New Haven, CT 06511
kenneth.gillingham@yale.edu

Sébastien Houde
ETH Zurich
Zürichbergstrasse 18
Switzerland – 8092 Zurich
shoude@ethz.ch

Arthur A. van Benthem
The Wharton School
University of Pennsylvania
3620 Locust Walk
USA – Philadelphia, PA 19104
arthurv@wharton.upenn.edu

May 8, 2019

The authors are grateful to Hunt Allcott, Mark Jacobsen, Chris Knittel, Matt Kotchen, Josh Linn, Erica Myers, Mathias Reynaert, John Rust, Jim Saltee and Joe Shapiro for their comments and suggestions. We also thank seminar participants at Georgetown University, Simon Fraser University, Stanford University, University of Connecticut, University of Copenhagen, University of Maryland, University of Pennsylvania, and Yale University.

1 Introduction

The transportation sector is now the largest contributor of carbon dioxide emissions in the United States and emissions from petroleum constituted 45% of all energy-related carbon dioxide emissions in 2017.¹ Fuel-economy regulations are the dominant policy to reduce carbon dioxide emissions from the transportation sector in the United States and many other countries, despite economists long arguing for a Pigouvian gasoline tax to internalize climate change (and other) externalities (Parry and Small 2005).

Fuel-economy standards require automakers to meet average fuel-economy targets for new light-duty vehicles. A common argument for such standards is that they “save consumers money” due to buyers undervaluing fuel economy at the time of the vehicle purchase (Parry, Walls, and Harrington 2007). This argument suggests that consumers are buying lower fuel economy vehicles, with higher fuel costs, than is ex post privately optimal for them. Such myopia is a common explanation for what has become known as the “energy efficiency gap,” whereby consumers do not adopt seemingly high-return energy-efficiency investments (Hausman 1979; Gillingham, Newell, and Palmer 2009; Allcott and Greenstone 2012).² Indeed, there is a large and growing behavioral economics literature documenting cases where consumers appear inattentive or otherwise seem to misoptimize in many settings, such as health plans (Abaluck and Gruber 2011, 2016), sales taxes (Chetty, Looney, and Kroft 2009), and heuristics for large-number processing (Lacetera, Pope, and Sydnor 2012).³

This paper presents the first evidence on the consumer valuation of fuel economy

¹From https://www.eia.gov/energyexplained/index.php?page=environment_where_ghg_come_from.

²We follow the terminology in the existing literature (e.g., Hausman 1979; Busse, Knittel, and Zettelmeyer 2013) and use “myopia” to describe a range of behavioral phenomena that could cause undervaluation.

³Studies have examined how consumers and market performance respond to information disclosure in various contexts, including financial decisions (Duflo and Saez 2003; Bertrand and Morse 2011; Goda, Manchester, and Sojourner 2014), takeup of social programs (Bhargava and Manoli 2015), sexually risky behavior (Dupas 2011), vehicle choice (Tadelis and Zettelmeyer 2015), electricity consumption (Jesoe and Rapson 2014), and educational investment (Jensen 2010). In this paper, we test how consumers respond to information disclosure about fuel-economy ratings.

from a natural experiment providing exogenous variation in the fuel-economy rating that new vehicle buyers observe. In 2012, after an audit by the U.S. Environmental Protection Agency (EPA), the two major automakers Hyundai and Kia acknowledged that they had overstated the fuel economy for 13 important vehicle models from the 2011-2013 model-years by one to six miles-per-gallon. This overstatement—by far the largest in history—affected over 1.6 million vehicles sold, including several popular models such as the Hyundai Elantra and Kia Rio. Hyundai and Kia blamed a “procedural error” in the mileage testing and had to abruptly change the official fuel-economy ratings for these vehicles. Following the restatement, the automakers agreed to compensate buyers who had already purchased vehicles with misstated ratings, while new car buyers after the restatement did not receive compensation.⁴ The restatement was unexpected—even just prior to it, Hyundai and Kia often advertised the high fuel economy of their vehicles as a major selling feature.

We first examine the equilibrium price response by consumers and firms to this large unexpected restatement.⁵ Using detailed microdata on all new vehicle transactions in the United States over the period August 2011 - June 2014, we find a 1.2% decline in the equilibrium prices of the affected models (just under \$300). We then proceed by directly estimating the consumer valuation of fuel economy. Using our preferred set of valuation assumptions, our results indicate that consumers are indifferent between one dollar in future gasoline costs and 15-38 cents in the vehicle purchase price (a ‘valuation parameter’ of 0.15-0.38) depending on the affected model-year, and using a discount rate of 4%. We thus find that consumers systematically undervalue fuel economy in vehicle purchases to a larger degree than reported by much of the recent literature. This conclusion is robust to a wide range of valuation assumptions, including vehicle supply elasticities, as we illustrate in a bounding exercise.

Previous studies estimating the consumer valuation of fuel economy use several

⁴From <https://kiampginfo.com/>

⁵In focusing on the equilibrium effects of the restatement, our study relates to the literature estimating the equilibrium effects of boycotts on firms or products (e.g., Chavis and Leslie 2009; Hendel, Lach, and Spiegel 2017).

different identification strategies, but most leverage changes in gasoline prices to test whether vehicle prices fully adjust with the changes in the expected discounted present value of future fuel costs. This basic approach was used as early as the 1980s, with Kahn (1986) finding that used car prices adjust only one third to one half the amount that would be expected based on the changes induced by shocks to gasoline costs and argues that used car buyers must be myopic.

More recent studies have documented a wide range of valuation parameter estimates. Allcott and Wozny (2014) exploit variation in gasoline prices and estimate a valuation parameter of 0.72 for used vehicle purchasers in the United States. This result suggests more limited undervaluation of fuel economy. Allcott and Wozny also present a wide range around their preferred estimate (from 0.42 to 1.01) due to different assumptions going into the calculation of the discounted present value of future fuel savings. Several other recent studies present estimates centered around one, implying that consumers fully value future fuel savings. Busse, Knittel, and Zettelmeyer (2013) also rely on gasoline price variation and use both new and used vehicle data, while Saltee, West, and Fan (2016) estimate their model with used vehicle auction data and use variation in odometer readings. Grigolon, Reynaert, and Verboven (2018) use temporal variation in gasoline prices combined with cross-sectional variation in engine technology to find a central-case valuation parameter of 0.91 in Europe. Taken together, these studies suggest modest undervaluation at most.⁶ In contrast, Leard, Linn, and Zhou (2018) use data from new vehicles in the United States and exploit the timing of adoption of fuel-saving technologies. They find a substantially lower valuation parameter of 0.54. Leard, Linn, and Springel (2019) focus on using cross-sectional variation in engine technologies and find even lower values; most of their estimates are below 0.30.

We bring a new identification strategy to shed light on this unsettled question. An appealing feature of using our natural experiment to understand consumer myopia is that we can rely on a sudden and exogenous shifter of the official fuel-economy rating, yet be

⁶Some earlier studies that do not explicitly estimate a valuation parameter similarly suggest full valuation of fuel economy (Goldberg 1998; Verboven 2002).

assured that the vehicles themselves are identical before and after the change. The rating is the primary source of information provided by the government to help consumers compare fuel economy across different vehicles, and thus it serves as a policy-relevant exogenous shifter of expected future fuel costs. This rating is likely highly salient to consumers when shopping for cars, as it features prominently on dealer lots and all major automotive websites.

If consumers appear to act myopically and undervalue fuel economy in new vehicle purchases, this implies that it is possible for a policy that shifts consumers into more efficient vehicles to be welfare-improving, even if environmental externalities are fully internalized. Our analysis uses a novel approach to provide guidance to policymakers on this critical parameter for understanding the costs and benefits of fuel-economy standards and their performance relative to a tax on gasoline.

2 The 2012 Fuel-Economy Rating Restatement

The restatement was made public on November 2, 2012, when EPA stated in a press release that “in processing test data, Hyundai and Kia allegedly chose favorable results rather than average results from a large number of tests.”⁷ This was a result of a 2012 EPA audit of the model-year 2012 Hyundai Elantra, which revealed a large discrepancy between the test results and the self-reported fuel economy provided by Hyundai. Based on this finding, EPA expanded its investigation to other Hyundai and Kia vehicles, uncovering many more discrepancies, all of which overstated fuel economy. The two automakers claimed that “honest mistakes” had been made, such as a “data processing error related to the coastdown testing method.”⁸

Immediately after the EPA press release, the fuel-economy ratings for all affected vehi-

⁷The incident was widely discussed in the press, e.g., see <https://www.nytimes.com/2012/11/03/business/hyundai-and-kia-acknowledge-overstating-the-gas-mileage-of-vehicles.html>.

⁸See <https://www.autoblog.com/2014/11/03/hyundai-kia-300-million-mpg-penalties/>.

cles were updated on all new car comparison websites, at www.fueleconomy.gov, and on the EPA fuel-economy labels on all new vehicles on dealers' lots.⁹ Hyundai and Kia were also required to update all advertising that mentioned the incorrect fuel-economy ratings. At the time of the restatement, over 900,000 vehicles with incorrect fuel-economy labels had already been sold, which amounts to roughly 35% of all 2011-2013 models sold through October 2012 by the two automakers. Appendix A provides a list of the restated models and the change in miles-per-gallon for each.

Prior to the restatement, Hyundai and Kia often mentioned the high fuel economy of their vehicles as a selling point.¹⁰ This added to the unexpected and abrupt nature of the restatement. Following the restatement, the automakers offered compensation to buyers that had already purchased vehicles with misstated fuel economies (see Appendix A for details). New vehicles offered after the restatement—the focus of our analysis—were not subject to the compensation.

3 Data

Our first dataset contains all dealer-reported new vehicle transactions in the United States from August 2011 to June 2014 from R.L. Polk. These data include the vehicle identification number (VIN) prefix (often known as the “VIN10” because it includes the first 10 digits that provide information about vehicle characteristics), the transaction date, the transaction price, and the Nielsen Designated Market Area (DMA), which is a commonly used geographic delineation for media markets. There are 210 DMAs in the United States and each is a cluster of similar counties that are covered by a specific group of television stations. The transaction price is inclusive of all dealer and manufacturer incentives. The data do not allow us to observe movement on other dealer's margins, such as preferen-

⁹Appendix A shows an example label.

¹⁰Consider this quote from a November 2, 2012 article (<https://www.autoblog.com/2012/11/02/hyundai-kia-admit-exaggerated-mileage-claims-will-compensate-o/>): “Hyundai aggressively advertised the fact that the brand offers four models that boast 40 mpg, but that claim is no longer true.”

tial financing. The VIN10 uniquely identifies the vehicle trim, engine size, and further characteristics.

Table 1 presents means of key variables for the affected models, non-affected models by Hyundai and Kia, and all other models in market segments with at least one affected vehicle. Panel A presents total sales and average transaction prices. For Hyundai, sales of affected models were about half of total sales, while for Kia, they comprised about a third. Hyundai and Kia have similar pricing, with the affected models being priced slightly below the non-affected models. Both automakers specialize in smaller cars that are priced below the average for other automakers.

Panel B shows the composition of each of the fleets and some characteristics. 71% of the affected Hyundai vehicles are small cars, while 80% of the affected Kia vehicles are crossovers. We thus have identifying variation across different classes of vehicles. Both automakers have unaffected small cars and crossovers, providing variation within classes as well. On average, we see that the affected models tend to have slightly lower weight and cost slightly less than non-affected models or models from other automakers.

For our calculations of the valuation of fuel economy, we also bring in data on monthly nationwide gasoline prices from the U.S. Energy Information Administration and on average vehicle-miles-traveled from the National Household Travel Survey (NHTS). We provide estimates using the 2006 NHTS, following Busse, Knittel, and Zettelmeyer (2013), as well as the recent 2017 NHTS.

4 The Equilibrium Effects of the Restatement

4.1 Effects on Transaction Prices

We begin our empirical analysis by examining the equilibrium effects of the restatement on new vehicle transaction prices. Our empirical approach is a difference-in-differences

estimator:

$$Price_{jrt} = \beta 1(Post\ Restatement)_t \times 1(Affected\ Model)_j + \rho_{t \times Class_j} + \mu_{t \times Make_j} + \eta_r \times 1(Post\ Restatement)_t + \eta_r + \omega_j + \epsilon_{jrt}. \quad (1)$$

where *Price* is either the log or level of the transaction price for a VIN10 *j* sold in region *r* (DMA) in year-month *t*. $1(Post\ Restatement)_t$ is an indicator variable for after the restatement in November 2012 and $1(Affected\ Model)_j$ is an indicator variable for an affected model. We next include year-month indicators interacted with vehicle class indicators ($\rho_{t \times Class_j}$) to allow for flexible time controls specific to each vehicle class. We further add year-month indicators interacted with make indicators ($\mu_{t \times Make_j}$) for flexible time controls for each automaker. We include DMA indicators (η_r) and their interaction with the post-restatement indicator ($\eta_r \times 1(Post\ Restatement)_t$). Finally, ω_j are VIN10 fixed effects. We weight the regressions by monthly sales.

Our coefficient of interest, β , is identified using the temporal variation before versus after the restatement across the affected and non-affected vehicles within automakers and within vehicle classes. We restrict the sample to only include vehicle classes in which Hyundai and Kia have affected cars: subcompact, compact, midsize, fullsize, sport, compact crossover, and midsize crossover. Our fixed effects specification exploits the panel nature of our data along with its high level of disaggregation to address a variety of potential time-invariant and time-varying confounders.¹¹

We expect our coefficient of interest β to be negative if the market responds in equilibrium to the downward adjustment of fuel economy for the affected models. Panel A of Table 2 presents our primary results. Columns 1-3 estimate the model using the log of the transaction price as the dependent variable. Columns 4-6 use the level of the price. Columns 3 and 6 are our preferred specifications. The coefficients become slightly larger

¹¹In this respect, our identification follows recent studies. For example, Allcott and Wozny (2014) and Busse, Knittel, and Zettelmeyer (2013) use temporal variation in gasoline prices after conditioning on year or model-year fixed effects. Grigolon, Reynaert, and Verboven (2018) use a country-specific quadratic time trend rather than year fixed effects. Sallee, West, and Fan (2016) exploit variation in odometer readings within a model-year while controlling for VIN10-year-month.

as we add fixed effects (especially in levels), but are generally quite similar across specifications.

Our results indicate that the restatement led to a 1.2% decrease in equilibrium transaction prices, which amounts to a \$294 decline on average across all affected models. Figure 1 presents the average treatment effects by month. To create this figure, we interacted $1(Post\ Restatement)_t \times 1(Affected\ Model)_j$ with each year-month in our sample and plotted the coefficients over time. We see no discernable evidence of a treatment effect prior to the restatement, but afterwards we observe a decrease in transaction prices (that hovers around 1%) for the affected models until January 2014. After this there are only few treated vehicles left and the treatment effect reverts back towards zero. By the end of our sample, the 2014 model-year vehicles would have been selling for almost a year (note no 2014 model-year vehicles are affected) and very few 2013 model-years are left on dealers' lots.

4.1.1 Robustness Checks

A critical assumption underlying any difference-in-differences analysis is the Stable Unit Treatment Value Assumption (SUTVA), which requires that the treatment assignment does not affect the potential outcomes of the non-treated observations (non-interference).¹² SUTVA can be violated if there are spillovers between the treated and control (e.g., from strategic pricing in a market with differentiated products) or if there are broader general equilibrium effects due to the treatment.

We perform several robustness checks to confirm that SUTVA holds in our case. Panel B in Table 2 presents our first SUTVA robustness checks by showing the results after excluding close substitute vehicles, which are the most likely to be affected by strategic pricing. If excluding close substitutes does not affect our estimates, then we can be confident that SUTVA holds.

¹²The classic SUTVA assumptions also require stability in the treatment. In our context, the fuel-economy rating changes by different amounts, and thus our primary results should be interpreted as an average effect.

Columns 1 and 4 exclude the Hyundai and Kia vehicles that are the closest substitutes to the restated models, but were not subject to a restatement. Close substitute vehicles are defined as those offered by the same automaker in the same R.L. Polk vehicle class. Columns 2 and 5 provide an alternative test that excludes the five most popular close substitutes from other automakers, where we define substitutes across automakers using data from Edmunds.com and MotorTrend.com.¹³ Columns 3 and 6 exclude the Hyundai and Kia substitutes as well as the substitutes from other automakers. Removing close substitutes makes little difference to the estimated coefficients in Panel A. The coefficients excluding substitutes are all within the 95% confidence interval of our primary specification, indicating that the slight change in the competitive landscape from the restatement had little influence on the pricing of substitute models.¹⁴

We expect that the restatement for Hyundai and Kia had negligible general equilibrium effects on the much larger vehicle market. However, one might be concerned that the widely-publicized restatement had an effect on the overall Hyundai and Kia brands, so that equilibrium prices change due to a diminished perception of the brands rather than a change in the fuel-economy ratings. Note however that our primary specification uses variation across affected and non-affected models within automakers, so any change in brand equity affects both the control and treatment groups. In fact, when we estimate the model removing all other automakers besides Hyundai and Kia, we find very similar results. This result, along with further robustness checks on sample selection, can be found in Appendix Tables B.3 and B.4.

4.1.2 Heterogeneous Effects on Transaction Prices

The restatement might be expected to influence the equilibrium pricing decisions of automakers differently based on the model-year of the vehicle and the magnitude of the

¹³Edmunds.com provides a list of other models that consumers considered for each model and model-year. MotorTrend.com explicitly provides a list of the closest competitors. We combined the two lists and then chose the five highest-selling vehicles from the combined list.

¹⁴In Appendix Tables B.1 and B.2, we use alternative vehicle class fixed effects and find the results are robust. These checks can be seen as changing the control group and the trends that the affected models are compared to.

change in the fuel-economy rating. In Table 3, we explore heterogeneous treatment effects with respect to these variables.¹⁵ Columns 1 and 2 replicate our preferred specification from Table 2. Columns 3 and 4 allow the treatment effect to vary by model-year. We see that the coefficients are generally similar, but the equilibrium price decline for the 2011-2012 model-years (1.7%) is somewhat greater than for the 2013 model-year (1.1%). In levels, the price reductions are \$544 and \$259, respectively. This difference could be due to differences in supply elasticities (see Section 4.2 for details) or automakers facing customers with different demand elasticities for the newest model-year vehicles.

Columns 5 and 6 allow the treatment effect to vary along with the change in the gallons-per-mile implied by the restatement. We use gallons-per-mile rather than miles-per-gallon because we anticipate consumers care about total expected fuel costs and fuel costs scale linearly with gallons-per-mile.¹⁶ The negative coefficient indicates that the price reductions are larger for models that faced a greater reduction in fuel economy (i.e., increase in fuel intensity). When evaluated at the mean change in gallons-per-mile (0.0019), the effects are smaller than in our preferred specification in columns 3 and 6 of Table 2 (-0.006 and -\$132 in logs and levels). This suggests that consumers do not respond to the magnitude of the restatement perfectly proportionately.

4.2 Effects on Other Outcomes?

In equilibrium, it is possible for there to be other adjustments as well. Busse, Knittel, and Zettelmeyer (2013) show that when gasoline prices change, sales of new vehicles tend to be affected even more than transaction prices. However, our setting is quite different. By November 2012, automakers had already ended production of model-year 2011 and 2012 vehicles and all remaining vehicles from those model-years were already on dealer lots. Model-year 2013 vehicles were still midway through their production cycle. Adjustments in production for these vehicle models are possible, but costly. Such adjustments would have required reallocating assembly lines or renegotiating contracts with suppliers, which

¹⁵Appendix Tables B.5 and B.6 explore heterogeneity by make and vehicle class.

¹⁶The results have nearly identical implications if we use miles-per-gallon.

may not be worth it for a one-time restatement. In fact, maintaining market share might well be the optimal long-run managerial strategy in response to a one-time negative shock (see Appendix C.1 for details). Therefore, supply was likely very inelastic for model-year 2011 and 2012 vehicles but possibly somewhat more elastic for model-year 2013 vehicles.

In Appendix C.1, we examine the equilibrium effects of the restatement on quantities using a specification similar to equation (1). Automobile sales tend to be highly idiosyncratic, however, with much difficult-to-explain variation occurring month to month. As a result, we obtain very noisy estimates: all coefficients are positive but imprecisely estimated. While we can only take this noisy evidence as suggestive, we certainly do not find clear evidence for a negative equilibrium quantity effect. We discuss the implications of positive or negative quantity effects for our eventual estimate of the fuel-economy valuation parameter in Section 5 and Appendix D.2, and find our conclusion about undervaluation to be robust.

Another possible adjustment could be to increase advertising expenditures. We examine this in Appendix C.2 and find no evidence of changes in either advertising expenditures or the number of advertisements after the restatement.

5 Implications for the Valuation of Fuel Economy

5.1 Valuing Fuel Economy

To understand how consumers value fuel economy, we are interested in how the discounted present value of future fuel costs influences vehicle purchase decisions. Going back to Hausman (1979), economists have examined how consumers trade off one dollar in upfront purchase costs against one dollar in the discounted present value of future energy costs. If consumers respond more to a change in upfront cost relative to future costs, this is taken as evidence of *undervaluation* of energy efficiency, or what is often described as myopia. It has become common to operationalize the valuation of energy efficiency through a valuation parameter, defined as the consumer response to the net present value

of future fuel costs over the response to the purchase price (e.g., Allcott and Wozny 2014; Sallee, West, and Fan 2016; Grigolon, Reynaert, and Verboven 2018; Leard, Linn, and Zhou 2018).¹⁷

Our approach to estimating undervaluation is inspired by Allcott and Wozny (2014). They start from a discrete choice model of vehicle choice with i.i.d extreme value idiosyncratic preferences, and invert the equation to arrive at a specification that regresses the vehicle purchase price on discounted lifetime fuel operating costs and controls. Our valuation specification is:

$$Price_{jrt} = \gamma \Delta G_{jt} + \rho_{t \times Class_j} + \mu_{t \times Make_j} + \eta_r \times 1(Post\ Restatement)_t + \eta_r + \omega_j + \epsilon_{jrt}. \quad (2)$$

where $Price_{jrt}$ is the vehicle transaction price and ΔG_{jt} is the change in the discounted lifetime fuel cost due to the restatement.¹⁸ In Appendix D.1, we motivate equation (2) from a random utility model and show that γ can be interpreted as the valuation parameter if sales do not adjust, which appears to be the case in our natural experiment (see Section 4.2). We thus interpret a value of -1 as full valuation—where an increase in expected future fuel costs is entirely reflected by a decrease in the purchase price—but discuss the implications of elastic supply in Appendix D.2.

There are two major empirical challenges to interpreting an estimate of γ in equation (2) as a causal estimate of undervaluation. First, ΔG_{jt} must be constructed based on assumptions about future driving, vehicle survival probabilities, expected future gasoline prices, and the car owner’s discount rate. We follow the existing literature in using an exhaustive set of assumptions to better understand the plausible range of γ . Second, ΔG_{jt} is

¹⁷Much of the early literature on energy efficiency valuation estimates an implicit discount rate that rationalizes full valuation, subject to assumptions about many other factors that could influence the valuation of fuel economy. We follow recent papers in presenting a valuation parameter subject to an assumed discount rate (and the same set of assumptions about other factors). This is mostly an expositional choice.

¹⁸Note $\Delta G_{jt} = 0$ for all non-affected models in this specification, so the variation in ΔG_{jt} is coming both from the differences between affected and non-affected models, as well as from the change in fuel economy due to the restatement. The only other source of variation in ΔG_{jt} could be from changing gasoline prices. Gasoline prices are similar before and after the restatement, but as a robustness check we replace the gasoline price with an average price over the entire period (shutting down this additional source of time series variation) and find similar results (Appendix Table D.2).

potentially endogenous due to a correlation between market shares (in ϵ_{jrt}) and expected future fuel costs, as well as potentially subject to measurement error (see Appendix D.1 for details). Our natural experiment overcomes these challenges because it provides a source of exogenous variation in ΔG_{jt} , and the restatement is perfectly observed.

We first estimate equation (2) using a baseline set of assumptions in constructing ΔG_{jt} : expected driving based on the 2017 NHTS, vehicle survival probabilities from Busse, Knittel, and Zettelmeyer (2013), and expected gasoline prices being held constant in real terms at the levels at time t (a martingale assumption, following evidence from Anderson, Kellogg, and Sallee (2015)). Panel A of Table 4 presents the results under these baseline assumptions. We show results for different discount rates, starting with a 1% rate in columns 1 and 2, and ending with a 12% rate in columns 7 and 8. For each discount rate, the first column presents the results using the pooled sample, while the second presents the results exploring heterogeneity in valuation across model-years.¹⁹

The results show that the equilibrium price changes induced by the restatement correspond to substantial undervaluation of fuel economy: the loss in the expected net present value of future fuel costs implied by the restatement far exceeds the equilibrium price changes, with the gap even larger for the affected 2013 model-years. The result in column 1 (1% discount rate) implies that consumers are indifferent between \$1 in expected future fuel costs and \$0.14 in the upfront purchase price (i.e., a valuation parameter of 0.14). The results in column 2 indicate substantial heterogeneity, with consumers buying the 2011-2012 model-years (35.4% of the affected vehicles) having a valuation parameter of 0.32, while for the 2013 model-year it is 0.13. Moving to a discount rate of 12%, the pooled sample shows a parameter of 0.25, where the 2011-2012 model-years have a valuation parameter of 0.56 and the 2013 model-year has a parameter of 0.22. It is difficult to arrive at a preferred specification when there are so many assumptions that could vary the parameter; we prefer a middle ground 4% discount rate (see Panel B of Table 4). This gives a valuation parameter of 0.15 for model-year 2013 and 0.38 for model-years 2011-2012.

¹⁹For the pooled sample, an implicit discount rate of approximately 80% would be required to bring the valuation parameter to one.

We cannot emphasize enough that with different sets of assumptions, the undervaluation parameter would change. For a wide enough range of assumptions, the valuation parameter can be as low as zero or as high as one. However, at a 4% discount rate and using reasonable sets of assumptions for constructing ΔG_{jt} that closely follow the existing literature, we find a range for the valuation parameter almost entirely below 0.5 (Appendix D.3). Moreover, in Appendix D.2 we allow for negative or positive quantity effects and find that the valuation parameter stays below 0.5 even when quantity effects are large (+/-5%; pooled sample). Our finding of substantial undervaluation is therefore robust.

5.2 Comparison to Previous Literature

Panel B of Table 4 summarizes the range of our results along with several notable papers in the literature. The valuation parameters in Busse, Knittel, and Zettelmeyer (2013), Sallee, West, and Fan (2016), and Grigolon, Reynaert, and Verboven (2018) are all close to one, which implies near-full valuation. While Allcott and Wozny (2014) and Leard, Linn, and Zhou (2018) find parameters consistent with undervaluation, our estimates are even lower. Our estimates, however, align with the heterogeneous estimates of Leard, Linn, and Springel (2019), which range from 0.06 to 0.76 but are below 0.30 for most demographic groups. There are several possible explanations for why our estimates are lower than most others.

First, we use a different source of identification. Expected fuel expenditures are the product of a consumer's gas price expectations and her estimate of the car's fuel economy. Our experiment leverages a change in EPA miles-per-gallon ratings, which dealers are required to feature on vehicles in their lots and which play prominently on new vehicle comparison websites. Other studies leverage changes in gasoline prices and therefore price expectations. Fully-informed rational consumers should respond equivalently to changes in gasoline prices and fuel-economy ratings, but it is possible there is a difference, e.g., if consumers are not perfectly-informed about fuel economy. This could imply our

results would better capture consumer behavior around regulations that directly affect EPA ratings than previous work.

Another possible explanation is that we are focusing on new cars from Hyundai and Kia, while other studies provide estimates from different markets. Sallee, West, and Fan (2016) estimate their model on data from used car auctions. Busse, Knittel, and Zettelmeyer (2013) use estimates based on both the new and used vehicle markets. But our study is not the only one focusing on new cars (e.g., Grigolon, Reynaert, and Verboven 2018; Leard, Linn, and Zhou 2018). It is possible that buyers of new Hyundais and Kias are different. On the one hand, it seems likely that Hyundai and Kia, which are known for smaller, more fuel-efficient cars, draw a segment of new car buyers that are more attentive to fuel economy, and thus would be expected to value fuel economy more than average. On the other hand, these car buyers may also be lower-income households who are more prone to steeply discount future fuel costs (Leard, Linn, and Springel 2019).

Our sample period also differs somewhat from previous work. Our results are from 2012 when the economy was still in a slow climb out from the Great Recession. Interest rates were very low and gasoline prices were generally low. Fuel economy undervaluation may vary over time and economic conditions, but studying this issue in more detail would require a long time series of restatement events.

Another possibility is that consumers already knew that the Hyundai and Kia models had lower fuel economy than was stated by the EPA ratings. Given how much of a surprise the restatement was (as is evidenced by the media articles), we find this implausible. While one can find blog posts for automobile aficionados prior to the restatement that indicated they were having a hard time achieving the EPA fuel economy, this is also true for many other models that achieve lower on-road fuel economy than reported by the EPA for many drivers. In general, the EPA-rated fuel economy is considered reliable and is used in all car comparison articles, websites, and apps that we are aware of (Jacobsen et al. 2019). All things considered, it appears highly unlikely that consumers already knew about the restatement in advance.

Another potential explanation for why our estimates differ is the approach used to estimate the valuation parameter. Some papers, such as Sallee, West, and Fan (2016) and Allcott and Wozny (2014) estimate the parameter directly, just as in our equation (2). Others approximate the parameter by separately estimating the average change in equilibrium prices and the average change in future fuel costs, and then dividing the first by the second. In the closely-related context of appliances, Houde and Myers (2019) point out that this approximation is likely to provide a biased estimate of the true valuation parameter. The intuition is that the ratio of the means of two variables is usually not the same as the mean of their ratio if these variables are heterogeneous and correlated. Appendix D.4 illustrates the issue mathematically and provides a conceptual example.

Our results suggest that this approximation bias may be large in the context of fuel-economy valuation. In Panel B of Table 4, we divide up the recent studies based on the approach taken, first showing studies estimating an exact valuation parameter and then showing studies using the approximation. We also provide our own estimates based on the same discount rates used in the previous studies, and present a range of valuation parameters allowing for heterogeneity between affected model-years 2011-2012 versus 2013. For comparison purposes, we also calculate the approximated valuation parameter. In our setting, we divide the estimated change in the equilibrium vehicle price in levels (Table 2) by the sales-weighted change in discounted future fuel costs implied by the restatement.

Our estimates show a wide range, but tend to be below 0.5 when the exact valuation parameter is estimated, suggesting much more substantial undervaluation than previous work. When we use the approximation, we find much greater valuation of fuel economy, with upper bound estimates as high as one, as in several previous papers. In fact, our estimate with a 1.3% discount rate is in line with Leard, Linn, and Zhou (2018). Altogether, these results suggest that some of the findings of nearly-full valuation of fuel economy in the literature may suffer from upward bias due to this approximation.

6 Conclusions

This paper exploits an unexpected restatement in the EPA-rated fuel economy for thousands of vehicles. A highly desirable feature of this natural experiment is that the vehicles themselves are identical before and after the restatement, providing us with a clean source of variation in expected future fuel costs by consumers. This restatement reduces equilibrium prices by 1.2%, or just under \$300. This variation allows us to estimate the valuation of future fuel costs, through a valuation parameter that captures how consumers weigh future fuel costs against the upfront purchase price. We find a wide range of valuation parameters that depend on several assumptions about consumer expectations, discounting, supply elasticities, and other factors, but even the upper end of our range suggests substantial undervaluation of fuel economy. For the 2011-2012 model-year vehicles, we find that consumers are indifferent between a \$1 increase in discounted future fuel costs and a \$0.38 increase in the upfront vehicle purchase price. The estimate drops to \$0.15 for the 2013 model-year vehicles.

This finding of substantial undervaluation differs from some—but not all—of the recent literature, but it differs much less after accounting for whether the study estimates the exact valuation parameter or an approximation. Other factors may also make a difference, including the empirical setting and the variation being exploited. We emphasize that our results are the first in the literature to use a natural experiment that actually changes EPA-rated fuel economy, and thus we believe that they provide valuable guidance to policymakers who are attempting to better understand the costs and benefits of fuel-economy standards.

References

- Abaluck, Jason and Jonathan Gruber. 2011. "Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program." *American Economic Review* 101:1180–1210.
- . 2016. "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance." *American Economic Review* 106:2145–2184.
- Alberini, Anna, Markus Bareit, and Massimo Filippini. 2016. "What Is the Effect of Fuel Efficiency Information on Car Prices? Evidence from Switzerland." *Energy Journal* 37 (3):315–342.
- Allcott, Hunt and Michael Greenstone. 2012. "Is There an Energy Efficiency Gap?" *Journal of Economic Perspectives* 26 (1):3–28.
- Allcott, Hunt and Christopher Knittel. 2019. "Are Consumers Poorly Informed about Fuel Economy? Evidence from Two Experiments." *American Economic Journal: Economic Policy* 11 (1).
- Allcott, Hunt and Nathan Wozny. 2014. "Gasoline Prices, Fuel Economy, and the Energy Paradox." *Review of Economics and Statistics* 96 (5):779–795.
- Anderson, Soren, Ryan Kellogg, and James Sallee. 2015. "What Do Consumers Believe about Future Gasoline Prices." *Journal of Environmental Economics and Management* 66 (3):383–403.
- Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4):841–890.
- Bertrand, Marianne and Adair Morse. 2011. "Information Disclosure, Cognitive Biases, and Payday Borrowing." *Journal of Finance* 66 (6):1865–1893.

- Bhargava, Saurabh and Dayanand Manoli. 2015. "Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment." *American Economic Review* 105 (11):3489–3529.
- Borenstein, Severin and Nancy Rose. 1995. "Bankruptcy and Pricing Behavior in U.S. Airline Markets." *American Economic Review* 85(2):397–402.
- Busse, Meghan, Christopher Knittel, and Florian Zettelmeyer. 2013. "Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review* 103 (1):220–256.
- Chavis, Larry and Phillip Leslie. 2009. "Consumer Boycotts: The Impact of the Iraq War on French Wine Sales in the U.S." *Quantitative Marketing and Economics* 7:37–67.
- Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review* 99:1145–1177.
- Davis, Lucas and Gilbert Metcalfe. 2015. "Does Better Information Lead to Better Choices? Evidence from Energy-Efficiency Labels." *Journal of the Association of Environmental and Resource Economists* 3 (3):589–625.
- Duflo, Esther and Emmanuel Saez. 2003. "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment." *Quarterly Journal of Economics* 118 (3):815–842.
- Dupas, Pascaline. 2011. "Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya." *American Economic Journal: Applied Economics* 3 (1):1–34.
- Gillingham, Kenneth, Richard Newell, and Karen Palmer. 2009. "Energy Efficiency Economics and Policy." *Annual Review of Resource Economics* 1:597–619.
- Goda, Gopi Shah, Colleen Flaherty Manchester, and Aaron J. Sojourner. 2014. "What Will My Account Really Be Worth? Experimental Evidence on How Retirement Income Projections Affect Saving." *Journal of Public Economics* 119:80–92.

- Goldberg, Pinelopi. 1998. "The Effects of the Corporate Average Fuel Efficiency Standards in the US." *Journal of Industrial Economics* 46 (1):1–33.
- Grigolon, Laura, Mathias Reynaert, and Frank Verboven. 2018. "Consumer Valuation of Fuel Costs and Tax Policy: Evidence from the European Car Market." *American Economic Journal: Economic Policy* 10 (3):193–225.
- Hausman, Jerry. 1979. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *Bell Journal of Economics* 10 (1):33–54.
- Hendel, Igal, Saul Lach, and Yossi Spiegel. 2017. "Consumers' Activism: The Cottage Cheese Boycott." *RAND Journal of Economics* 48 (4):972–1003.
- Houde, Sébastien and Erica Myers. 2019. "Heterogeneous (mis-) Perceptions of Energy Costs: Implications for Measurement and Policy Design." Working Paper 25722, National Bureau of Economic Research.
- Jacobsen, Mark R., Christopher R. Knittel, James M. Sallee, and Arthur A. van Benthem. 2019. "The Use of Regression Statistics to Analyze Imperfect Pricing Policies." *Journal of Political Economy* Forthcoming.
- Jacobsen, Mark R. and Arthur A. van Benthem. 2015. "Vehicle Scrappage and Gasoline Policy." *American Economic Review* 105 (3):1312–1338.
- Jensen, Robert. 2010. "The (Perceived) Returns to Education and the Demand for Schooling." *Quarterly Journal of Economics* 125 (2):515–548.
- Jessoe, Katrina and David Rapson. 2014. "Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review* 104 (4):1417–1438.
- Kahn, James. 1986. "Gasoline Prices and the Used Automobile Market." *Quarterly Journal of Economics* 101:323–340.
- Lacetera, Nicola, Devin Pope, and Justin Sydnor. 2012. "Heuristic Thinking and Limited Attention in the Car Market." *American Economic Review* 102 (5):2206–2236.

- Leard, Benjamin, Joshua Linn, and Katalin Springel. 2019. "Pass-through and Welfare Effects of Regulations That Affect Product Attributes." Working Paper.
- Leard, Benjamin, Joshua Linn, and Yichen Zhou. 2018. "How Much Do Consumers Value Fuel Economy and Performance? Evidence from Technology Adoption." Resources for the Future Working Paper.
- Newell, Richard and Juha Siikamaki. 2014. "Nudging Energy Efficiency Behavior: The Role of Information Labels." *Journal of the Association of Environmental and Resource Economists* 1 (4):555–598.
- Parry, Ian and Kenneth Small. 2005. "Does Britain or the United States Have the Right Gasoline Tax?" *American Economic Review* 95 (4):1276–1289.
- Parry, Ian, Margaret Walls, and Winston Harrington. 2007. "Automobile Externalities and Policies." *Journal of Economic Literature* 45:374–400.
- Sallee, James, Sarah West, and Wei Fan. 2016. "Do Consumers Recognize the Value of Fuel Economy? Evidence from Used Car Prices and Gasoline Price Fluctuations." *Journal of Public Economics* 135:61–73.
- Tadelis, Steven and Florian Zettelmeyer. 2015. "Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment." *American Economic Review* 105 (2):886–905.
- Verboven, Frank. 2002. "Quality-Based Price Discrimination and Tax Incidence: Evidence from Gasoline and Diesel Cars." *RAND Journal of Economics* 33 (2):275–297.

Tables & Figures

Table 1: Mean Sales, Prices, and Characteristics Across Automakers

	Affected Models		Not Affected Models		
	Hyundai (1)	Kia (2)	Hyundai (3)	Kia (4)	Others (5)
Panel A: Sales and Transaction Prices					
Total Sales (1000s)	1,041	516	944	1,001	26,300
Price (1000s \$)	21.6	20.0	24.1	23.5	28.6
# of Models by Model-Year	16	10	49	36	1,131
Panel B: Selected Vehicle Characteristics					
Fraction Sport	0.01	0.00	0.03	0.00	0.04
Fraction Small Car	0.71	0.18	0.16	0.22	0.33
Fraction Large Car	0.09	0.03	0.62	0.41	0.31
Fraction Crossover	0.19	0.80	0.19	0.36	0.33
Engine Cylinders	4.17	4.00	4.23	4.25	4.70
Displacement (liters)	2.02	1.98	2.39	2.34	1.72
Gross Vehicle Weight	2.89	2.96	3.28	3.23	3.47
MSRP (1000s \$)	20.8	18.9	24.1	22.8	28.7
Fuel Economy (miles/gallon)	29.5	25.8	27.0	27.0	26.4

Notes: Data cover August 2011 to June 2014 and include only classes of vehicles that have at least one affected model. A unit of observation is a year-month-DMA-VIN10, and these summary statistics are unweighted. The number of models by model-year refers to all model \times model-year combinations in each category (note some models have both affected and unaffected trims, and thus they may fall into both the affected and unaffected categories). DMA refers to a Nielsen Designated Market Area, which is an area covering several counties. MSRP refers to the manufacturer suggested retail price. All dollars are nominal dollars.

Table 2: Effect of Restatement on Transaction Prices

	(1)	(2)	(3)	(4)	(5)	(6)
	Logs			Levels		
Panel A: Primary Results						
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	-0.010 (0.004)	-0.010 (0.004)	-0.012 (0.003)	-150 (80)	-259 (94)	-294 (91)
Year-Month \times Class FE		Y	Y		Y	Y
Year-Month \times Make FE	Y	Y	Y	Y	Y	Y
VIN10 FE	Y	Y	Y	Y	Y	Y
DMA FE	Y		Y	Y		Y
$1(Post\ Restatement) \times DMA\ FE$	Y		Y	Y		Y
R-squared	0.95	0.92	0.95	0.96	0.95	0.96
N	1.52m	1.52m	1.52m	1.52m	1.52m	1.52m
Panel B: Robustness Checks for SUTVA Assumption						
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	-0.011 (0.004)	-0.014 (0.003)	-0.013 (0.003)	-261 (94)	-365 (83)	-342 (84)
Year-Month \times Class FE	Y	Y	Y	Y	Y	Y
Year-Month \times Make FE	Y	Y	Y	Y	Y	Y
VIN10 FE	Y	Y	Y	Y	Y	Y
DMA FE	Y	Y	Y	Y	Y	Y
$1(Post\ Restatement) \times DMA\ FE$	Y	Y	Y	Y	Y	Y
Exclude close substitutes of same make	Y			Y		
Exclude close substitutes of other makes		Y			Y	
Exclude all close substitutes			Y			Y
R-squared	0.95	0.95	0.95	0.96	0.96	0.96
N	1.50m	1.41m	1.39m	1.50m	1.41m	1.39m

Notes: Dependent variable is log or level of the transaction price (in dollars). An observation is a year-month-DMA-VIN10. VIN10 refers to the VIN prefix, which is a trim-engine combination. DMA refers to a Nielsen Designated Market Area, which is an area covering several counties. Class refers to the vehicle class. *Post Restatement* refers to the year-month being during or after November 2012. All estimations are weighted by monthly sales. Standard errors clustered by VIN10.

Table 3: Heterogeneous Effects of the Restatement on Transaction Prices

	Primary		Model Year		Δ GPM	
	(1)	(2)	(3)	(4)	(5)	(6)
	logs	levels	logs	levels	logs	levels
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	-0.012 (0.003)	-294 (91)				
$1(Post\ Restatement)_t \times 1(2011 - 2012\ Affected\ Model)_j$			-0.017 (0.006)	-544 (128)		
$1(Post\ Restatement)_t \times 1(2013\ Affected\ Model)_j$			-0.011 (0.004)	-259 (98)		
$1(PostRestatement)_t \times 1(Affected\ Model)_j \times \Delta GPM$					-2.92 (0.90)	-66544 (22470)
Year-Month \times Class FE	Y	Y	Y	Y	Y	Y
Year-Month \times Make FE	Y	Y	Y	Y	Y	Y
VIN10 FE	Y	Y	Y	Y	Y	Y
DMA FE	Y	Y	Y	Y	Y	Y
$1(Post\ Restatement) \times DMA\ FE$	Y	Y	Y	Y	Y	Y
R-squared	0.95	0.96	0.95	0.96	0.95	0.96
N	1.52m	1.52m	1.52m	1.52m	1.52m	1.52m

Notes: Dependent variable is log or level of the transaction price (in dollars). An observation is a year-month-DMA-VIN10. VIN10 refers to the VIN prefix, which is a trim-engine combination. DMA refers to a Nielsen Designated Market Area, which is an area covering several counties. Class refers to the vehicle class. *Post Restatement* refers to the year-month being during or after November 2012. Δ GPM refers to the change in the gallons-per-mile from the restatement. All estimations are weighted by monthly sales. Standard errors clustered by VIN10.

Table 4: Estimates of the Valuation of Fuel Economy

Panel A: Exact Valuation Parameter Estimation Results from the Restatement								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$r = 1\%$		$r = 4\%$		$r = 7\%$		$r = 12\%$	
$1(\Delta Lifetime\ Fuel\ Costs)_{jt} \times$ $1(Affected\ Model)_j$	-0.14		-0.17		-0.20		-0.25	
	(0.05)		(0.06)		(0.06)		(0.08)	
$1(\Delta Lifetime\ Fuel\ Costs)_{jt} \times$ $1(2011 - 2012\ Affected\ Model)_j$		-0.32		-0.38		-0.44		-0.56
		(0.16)		(0.19)		(0.23)		(0.29)
$1(\Delta Lifetime\ Fuel\ Costs)_{jt} \times$ $1(2013\ Affected\ Model)_j$		-0.13		-0.15		-0.18		-0.22
		(0.05)		(0.05)		(0.06)		(0.08)
Year-Month \times Class FE	Y	Y	Y	Y	Y	Y	Y	Y
Year-Month \times Make FE	Y	Y	Y	Y	Y	Y	Y	Y
VIN10 FE	Y	Y	Y	Y	Y	Y	Y	Y
DMA FE	Y	Y	Y	Y	Y	Y	Y	Y
1(Post Restatement) \times DMA FE	Y	Y	Y	Y	Y	Y	Y	Y
R-squared	0.96	0.96	0.96	0.96	0.96	0.96	0.96	0.96
N	1.52m	1.52m	1.52m	1.52m	1.52m	1.52m	1.52m	1.52m
Panel B: Comparison with Recent Studies								
Studies using exact valuation parameter				r	valuation parameter			
Sallee, West, and Fan (2016)				5%	1.01			
Allcott and Wozny (2014)				6%	0.76			
Own Estimate from Restatement				5%	[0.16-0.40]			
Own Estimate from Restatement				6%	[0.17-0.42]			
Studies using approximate valuation parameter								
Busse, Knittel, and Zettelmeyer (2013)				6%	1.33			
Grigolon, Reynaert, and Verboven (2018)				6%	0.91			
Leard, Linn, and Zhou (2018)				1.3%	0.54			
Own Estimate from Restatement				6%	[0.39-0.95]			
Own Estimate from Restatement				1.3%	[0.31-0.78]			

Notes: Dependent variable is the transaction price (in nominal dollars). Lifetime fuel costs are computed using annual U.S. gasoline prices, survival probabilities from Jacobsen and van Benthem (2015), and VMT from NHTSA (2018). In Panel A, the results are reported for different discount rates (r). A coefficient of -1 implies that a one-dollar increase in lifetime fuel costs reduces the transaction price by one dollar. Values between -1 and 0 imply that consumers undervalue future fuel costs. An observation is a year-month-DMA-VIN10. VIN10 refers to the VIN prefix, which is a trim-engine combination. DMA refers to a Nielsen Designated Market Area, which is an area covering several counties. Class refers to the vehicle class. *Post Restatement* refers to the year-month being during or after November 2012. All estimations are weighted by monthly sales. Standard errors clustered by VIN10. In Panel B, we report a range of our own estimates that accounts for heterogeneity between model-years 2011-2012 vs. 2013.

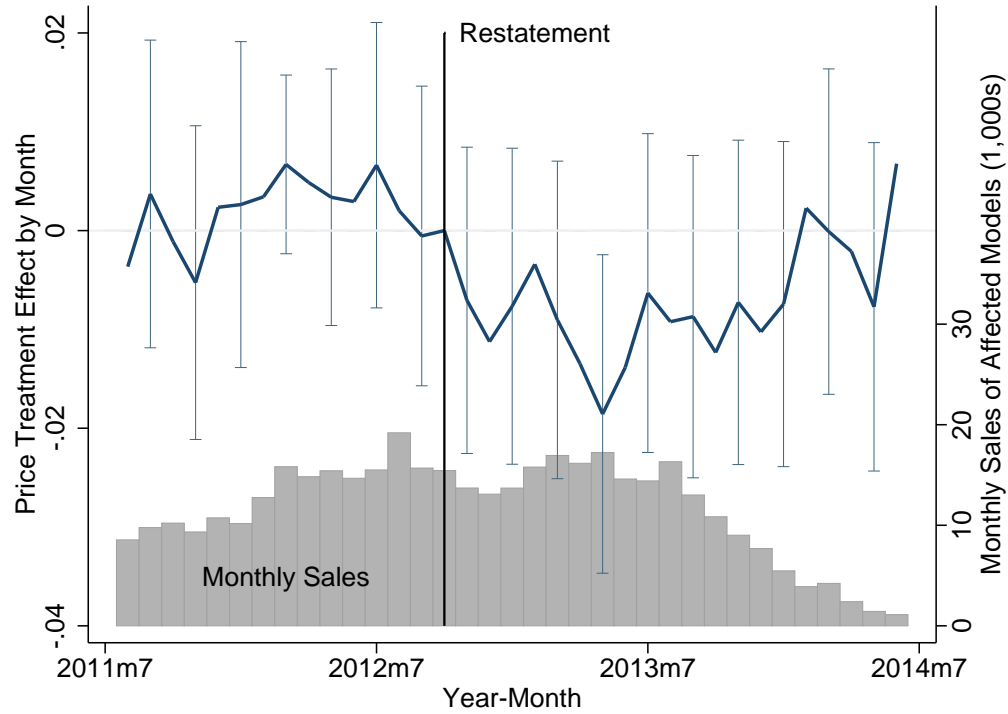


Figure 1: The Price Effect of the Restatement on Affected Models by Month Along with the Monthly Sales of Affected Models

Notes: The black vertical line indicates the fuel-economy restatement date. Treatment effects on price are on the left vertical axis; monthly sales of affected models are on the right vertical axis. The standard error for every other month is shown by the bars and whiskers. Note that the overall pre-post treatment effect is statistically significant (Table 2), although the monthly treatment effects are noisily estimated.

APPENDIX

A Fuel-Economy Label and Affected Vehicles

This appendix provides further details on the compensation offered to previous buyers, provides a complete list of affected vehicles, and gives an example of a fuel-economy label.

While there have been other fuel-economy restatements for a small number of vehicle models before (e.g., Ford restated the fuel economy for six models in 2014, and similar issues arose in 2019), the restatement by Hyundai and Kia was by far the largest in history and the first example of a restatement that affected many models. To make amends after this restatement, Hyundai and Kia provided owners of the affected vehicles purchased prior to the restatement with a lifetime offer of reimbursement based on the difference between the original and restated EPA fuel-economy rating (plus a 15% premium as an apology).²⁰ This compensation was announced only after the news about the restatement became public. Buyers were compensated via prepaid debit cards given at dealerships based on odometer readings and the fuel costs for the region in which they live. For example, a 1 mile-per-gallon adjustment amounted to a refund of approximately \$88 for an owner who drove 15,000 miles.

Through a class-action lawsuit, with a settlement finally approved by the courts on July 6, 2015, a second reimbursement option was added allowing affected customers to receive a single cash lump-sum payment (so customers could avoid having to return to the dealership frequently to have mileage verified).²¹ An appellate court put this settlement on hold in January 2018, ruling that a lower court had made errors in approving the settlement. As a result, there is still a class-action lawsuit working its way through the

²⁰From <https://www.autoblog.com/2012/11/02/hyundai-kia-admit-exaggerated-mileage-claims-will-compensate-o/>

²¹From <https://www.consumerwatchdog.org/courtroom/us-court-appeals-rejects-hyundaikia-settlement-fuel-economy-scandal>

courts as of January 2019.²²

Note that both the initial compensation and any later payments resulting from class-action lawsuits only affected vehicles that had already been sold before the restatement date, and did not affect new vehicle buyers afterwards. As such, the new car transaction prices that we analyze do not involve or include compensation or settlement payments.

Next we move to the list of all of the Hyundai and Kia vehicles affected by the restatement. Table A.1 contains a complete list of all of the Hyundai affected vehicles, along with selected vehicle characteristics. Table A.2 provides the same information for the Kia affected vehicles. 80,000 of the vehicles sold had their combined (city and highway) rating drop by 3-4 miles-per-gallon, while 240,000 dropped by 2 miles-per-gallon, and 580,000 dropped by 1 mile-per-gallon.²³ Note that for some models, the change in the combined miles-per-gallon rating is zero, even if the city or highway ratings changed. In Table B.4 below, we show a robustness check in which we run our primary specifications while excluding such minimally affected models to confirm that they are not affecting our results.

We now move to a discussion of the fuel-economy label. Fuel-economy labels on all new vehicles indicate the combined city/highway fuel economy of the vehicle in large block letters, include an estimate of the projected annual fuel cost from running that vehicle in large letters, include a dollar value savings (or spending) in fuel costs over the next five years relative to the average new vehicle, and also provide the vehicle's tailpipe greenhouse gas rating and a smog rating.²⁴ The EPA-rated fuel economy on the labels is also presented on websites widely used by car buyers, such as www.fueleconomy.gov and www.edmunds.com. In any comparison between vehicles, the EPA-rated fuel economy values will play prominently.

²²Hyundai and Kia also settled with the U.S. EPA and agreed to pay \$100 million in civil penalties, the largest such fines in EPA history up to that date, in addition to relinquishing emissions credits worth around \$200 million and offering previous buyers compensation. See <https://www.epa.gov/enforcement/hyundai-and-kia-clean-air-act-settlement>

²³Source: <https://www.autoblog.com/2012/11/02/hyundai-kia-admit-exaggerated-mileage-claims-will-compensate-o/>.

²⁴The combined city/highway fuel-economy estimate is based on U.S. EPA test ratings. The annual fuel cost estimates and fuel savings estimates are based on on-road fuel economy and an assumed 15,000 miles driven annually.

Table A.1: Hyundai Affected Models

(1)	(2)	(3)	(4)	(5)	(6)	(7) Original Rating			(10) Restated Rating		
Model	Model Year	Trim	Engine	Drive	Tran.	City MPG	Hwy MPG	Comb. MPG	City MPG	Hwy MPG	Comb. MPG
Elantra	2011		1.8L		Automatic	29	40	33	28	38	32
Elantra	2011		1.8L		Manual	29	40	33	28	38	32
Sonata HEV	2011		2.4L		Automatic	35	40	37	34	39	36
Accent	2012		1.6L		Automatic	30	40	33	28	37	31
Accent	2012		1.6L		Manual	30	40	34	28	37	32
Azera	2012		3.3L		Automatic	20	29	23	20	28	23
Elantra	2012		1.8L		Automatic	29	40	33	28	38	32
Elantra	2012		1.8L		Manual	29	40	33	28	38	32
Genesis	2012		3.8L		Automatic	19	29	22	18	28	22
Genesis	2012		4.6L		Automatic	17	26	20	16	25	19
Genesis	2012		5.0L		Automatic	17	26	20	17	25	20
Genesis	2012		5.0L R-Spec		Automatic	16	25	19	16	25	18
Sonata HEV	2012		2.4L		Automatic	35	40	37	34	39	36
Tucson	2012		2.0L	2WD	Automatic	23	31	26	22	29	25
Tucson	2012		2.0L	2WD	Manual	20	27	23	20	26	22
Tucson	2012		2.4L	2WD	Automatic	22	32	25	21	30	25
Tucson	2012		2.4L	4WD	Automatic	21	28	23	20	27	23
Veloster	2012		1.6L		Automatic	29	38	32	27	35	30
Veloster	2012		1.6L		Manual	28	40	32	27	37	31
Accent	2013		1.6L		Automatic	30	40	33	28	37	31
Accent	2013		1.6L		Manual	30	40	34	28	37	32
Azera	2013		3.3L		Automatic	20	30	24	20	29	23
Elantra	2013		1.8L		Automatic	29	40	33	28	38	32
Elantra	2013		1.8L		Manual	29	40	33	28	38	32
Elantra	2013	Coupe	1.8L		Automatic	28	39	32	27	37	31
Elantra	2013	Coupe	1.8L		Manual	29	40	33	28	38	32
Elantra	2013	GT	1.8L		Automatic	28	39	32	27	37	30
Elantra	2013	GT	1.8L		Manual	27	39	31	26	37	30
Genesis	2013		3.8L		Automatic	19	29	22	18	28	22
Genesis	2013		5.0L R-Spec		Automatic	16	25	19	16	25	18
Santa Fe	2013		2.0L Turbo	2WD	Automatic	21	31	25	20	27	23
Santa Fe	2013		2.4L	2WD	Automatic	22	33	26	21	29	24
Santa Fe	2013		2.0L Turbo	4WD	Automatic	20	27	22	19	24	21
Santa Fe	2013		2.4L	4WD	Automatic	21	28	23	20	26	22
Tucson	2013		2.0L	2WD	Automatic	23	31	26	22	29	25
Tucson	2013		2.0L	2WD	Manual	20	27	23	20	26	22
Tucson	2013		2.4L	2WD	Automatic	22	32	25	21	30	25
Tucson	2013		2.4L	4WD	Automatic	21	28	23	20	27	23
Veloster	2013		1.6L		Automatic	29	40	33	28	37	31
Veloster	2013		1.6L Turbo		Automatic	25	34	29	24	31	28
Veloster	2013		1.6L		Manual	28	40	32	27	37	31
Veloster	2013		1.6L Turbo		Manual	26	38	30	24	35	28

Source: <https://hyundaiampginfo.com/customerinfo/affected-models> and <https://kiampginfo.com/overview/affected-models>. MPG denotes miles-per-gallon.

Table A.2: Kia Affected Models

(1)	(2)	(3)	(4)	(5)	(6)	(7) (8) (9) Original Rating			(10) (11) (12) Restated Rating		
Model	Model Year	Trim	Engine	Drive	Tran.	City MPG	Hwy MPG	Comb. MPG	City MPG	Hwy MPG	Comb. MPG
Optima HEV	2011		2.4L	2WD	Automatic	35	40	37	34	39	36
Rio	2012		1.6L	2WD	Automatic	30	40	33	28	36	31
Rio	2012		1.6L	2WD	Manual	30	40	34	29	37	32
Sorento	2012	GDI	2.4L	2WD	Automatic	22	32	25	21	30	24
Sorento	2012	GDI	2.4L	4WD	Automatic	21	28	23	20	26	22
Soul	2012		1.6L	2WD	Automatic	27	35	30	25	30	27
Soul	2012		1.6L	2WD	Manual	27	35	30	25	30	27
Soul	2012		2.0L	2WD	Automatic	26	34	29	23	28	25
Soul	2012		2.0L	2WD	Manual	26	34	29	24	29	26
Soul	2012	ECO	1.6L	2WD	Automatic	29	36	32	26	31	28
Soul	2012	ECO	2.0L	2WD	Automatic	27	35	30	24	29	26
Sportage	2012		2.0L	2WD	Automatic	22	29	24	21	28	24
Sportage	2012		2.4L	2WD	Automatic	22	32	25	21	30	25
Sportage	2012		2.4L	2WD	Manual	21	29	24	20	27	23
Sportage	2012		2.0L	4WD	Automatic	21	26	23	20	25	22
Sportage	2012		2.4L	4WD	Automatic	21	28	24	20	27	23
Optima HEV	2012		2.4L	2WD	Automatic	35	40	37	34	39	36
Rio	2013		1.6L	2WD	Automatic	30	40	33	28	36	31
Rio	2013		1.6L	2WD	Manual	30	40	34	29	37	32
Rio	2013	ECO	1.6L	2WD	Automatic	31	40	34	30	36	32
Sorento	2013	GDI	2.4L	2WD	Automatic	22	32	25	21	30	24
Sorento	2013	GDI	2.4L	4WD	Automatic	21	28	23	20	26	22
Soul	2013		1.6L	2WD	Automatic	27	35	30	25	30	27
Soul	2013		1.6L	2WD	Manual	27	35	30	25	30	27
Soul	2013		2.0L	2WD	Automatic	26	34	29	23	28	25
Soul	2013		2.0L	2WD	Manual	26	34	29	24	29	26
Soul	2013	ECO	1.6L	2WD	Automatic	29	36	32	26	31	28
Soul	2013	ECO	2.0L	2WD	Automatic	27	35	30	24	29	26
Sportage	2012		2.0L	2WD	Automatic	22	29	24	21	28	24
Sportage	2012		2.4L	2WD	Automatic	22	32	25	21	30	25
Sportage	2012		2.4L	2WD	Manual	21	29	24	20	27	23
Sportage	2012		2.0L	4WD	Automatic	21	26	23	20	25	22
Sportage	2012		2.4L	4WD	Automatic	21	28	24	20	27	23

Source: <https://hyundaimpginfo.com/customerinfo/affected-models> and <https://kiampginfo.com/overview/affected-models>. MPG denotes miles-per-gallon.

In May 2011, the Environmental Protection Agency and National Highway Traffic Safety Administration updated the label and it became widely used by nearly all automakers starting with model-year 2012. It was mandatory starting with model-year 2013. Figure A.1 provides an example of the post-2011 fuel-economy label required to be posted on all new vehicles at the dealership. The fuel economy listed on the label for each affected Hyundai or Kia vehicle was updated immediately at the beginning of November in 2012.

The fuel-economy rating—featured prominently on all major automotive websites and on the labels on dealer lots—is the primary source of information for potential car buyers to compare fuel economy across different vehicles. Fuel economy is likely a highly salient vehicle specification during the car-buying process, as consumers explicitly consider tradeoffs between vehicle models and are repeatedly presented with the same information about fuel economy. Our experiment leverages a change in these miles-per-gallon ratings, while other studies exploit changes in gasoline prices. Both fuel-economy ratings and gasoline prices should inform the consumer’s estimate of future fuel savings. The salience of the miles-per-gallon rating during the car-buying process adds to the appeal of our setting; whether used car drivers know or remember the fuel economy of their vehicles is an unsettled question.

There is a growing literature on the extent to which consumers pay attention to labels about the energy efficiency of products. For example, Newell and Siikamaki (2014) find that the EnergyGuide label for appliances that provides simple information on the monetary value of energy savings appears to come close to guiding cost-efficient decisions. Davis and Metcalf (2015) show that more precise information from EnergyGuide labels can lead to significantly better choices. Houde and Myers (2019) also show heterogeneity in the response to energy information in appliance purchases. In one of the few papers on fuel-economy labels, Alberini, Bareit, and Filippini (2016) find that discrete fuel-economy grades (‘A’-‘G’) on mandatory labels for new vehicles in Switzerland influence equilibrium prices. This literature allows us to hypothesize that a large change in the listed fuel

economy on the labels will influence equilibrium outcomes in the new vehicle market.²⁵ Moreover, in our context, it is not just the label that changed, but actually the EPA fuel-economy rating, which affects everywhere that fuel economy is mentioned.



Figure A.1: An Example of a Fuel-Economy Label

²⁵The fact that Allcott and Knittel (2019) show that interventions to provide information about fuel economy (in addition to the fuel-economy labels) have little effect on behavior casts some doubt on the effectiveness of informational interventions, but is still consistent with consumers basing their beliefs on the rated fuel economy posted on the vehicle and found on websites and in manufacturer brochures.

B Robustness Checks

This section provides a series of robustness checks on our primary results. We begin by focusing on several different sets of fixed effects, which slightly change the variation being used to identify our coefficients. Table B.1 provides the first set of robustness results by including quarter-of-age by make fixed effects to capture the cyclicalities in the vehicle market that depends on the time since a vintage of a vehicle was introduced to the market.

Table B.1: Robustness Checks with Quarter-of-Age Fixed Effects

	(1)	(2)	(3)	(4)	(5)	(6)
	Logs			Levels		
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	-0.012 (0.003)	-0.012 (0.003)	-0.011 (0.003)	-294 (91)	-294 (92)	-276 (89)
Year-Month \times Class FE	Y	Y	Y	Y	Y	Y
Year-Month \times Make FE	Y	Y	Y	Y	Y	Y
VIN10 FE	Y	Y	Y	Y	Y	Y
DMA FE	Y	Y	Y	Y	Y	Y
$1(Post\ Restatement) \times DMA\ FE$	Y	Y	Y		Y	Y
Quarter-of-Age FE		Y			Y	
Quarter-of-Age \times Make FE			Y			Y
R-squared	0.95	0.95	0.95	0.96	0.96	0.96
N	1.52m	1.52m	1.52m	1.52m	1.52m	1.52m

Notes: Dependent variable is log or level of the transaction price (in dollars). Columns 1 and 4 are our primary specification from Table 2. An observation is a year-month-DMA-VIN10. VIN10 refers to the VIN prefix, which is a trim-engine combination. DMA refers to a Nielsen Designated Market Area, which is an area covering several counties. Class refers to the vehicle class. *Post Restatement* refers to the year-month being during or after November 2012. Quarter-of-age refers to the number of quarters since the introduction of a new VIN10. All estimations are weighted by monthly sales. Standard errors are clustered at the VIN10 level.

In Appendix Table B.2, we perform further robustness checks that include different sets of fixed effects for month-of-sample interacted with vehicle class. Specifically, we change the definition of a vehicle class to be a finer vehicle class definition than the one used in our main specification, where we do not distinguish luxury and non-luxury brands. In this robustness test, we use the exact segment definition proposed by R.L. Polk, which distinguishes luxury and non-luxury brands (which we label “finer class fixed effects”). We also use a coarser set of class fixed effects, which combine compact, mid size and full size crossover utility vehicles (into “crossover”); compact, mid size and full size

sport utility vehicles (into “SUV”); subcompacts and compacts (into “small cars”); and mid size and full size (into “large cars”). These checks slightly change the variation being used, which amounts to effectively changing the composition of the control group the affected models are compared to. We find that our results are highly robust to these alternative specifications.

Table B.2: Robustness Checks with Alternate Class Fixed Effects

	(1)	(2)	(3)	(4)	(5)	(6)
	Logs			Levels		
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	-0.012 (0.003)	-0.011 (0.004)	-0.011 (0.004)	-294 (91)	-283 (93)	-240 (90)
Year-Month \times Class FE	Y			Y		
Year-Month \times Finer Class FE		Y			Y	
Year-Month \times Coarser Class FE			Y			Y
Year-Month \times Make FE	Y	Y	Y	Y	Y	Y
VIN10 FE	Y	Y	Y	Y	Y	Y
DMA FE	Y	Y	Y	Y	Y	Y
$1(Post\ Restatement) \times DMA\ FE$	Y	Y	Y	Y	Y	Y
R-squared	0.95	0.95	0.95	0.96	0.96	0.96
N	1.52m	1.52m	1.52m	1.52m	1.52m	1.52m

Notes: Dependent variable is log or level of the transaction price (in dollars). Columns 1 and 4 are our primary specification. An observation is a year-month-DMA-VIN10. VIN10 refers to the VIN prefix, which is a trim-engine combination. DMA refers to a Nielsen designated market area, which is an area covering several counties. Class refers to the vehicle class. *Post Restatement* refers to the year-month being during or after November 2012. All estimations are weighted by monthly sales. Standard errors are clustered at the VIN10 level.

We also perform a further set of robustness checks. First, we perform a series of checks relating to decisions we made in creating our dataset. We see what happens if we do not drop vehicles with transaction prices below \$5,000 (3,203 additional vehicles are retained). We view transaction prices less than \$5,000 with suspicion, as they are likely miscoded. We also examine the effect of excluding price outliers by only including vehicle transactions within a price ratio around the mean price for that model-trim over the whole sample period between 0.67 and 1.5. Finally, we restrict the sample to include Hyundais and Kias only, allowing us to focus only on variation between affected and non-affected models for these two automakers. In Table B.3 we see some minor differences, but by-and-large, we find that our results are robust across these specifications.

In our final set of robustness checks, we run all of the primary specifications but we exclude affected models where the change in the rated fuel economy is minimal (defined as only changes in city and/or highway ratings, but no change in the combined rating). There were a fair number of these models, and one might be concerned that they skew our results. Table B.4 excludes these minimally treated models from the sample. Again, the results are remarkably similar.

One may also be interested in heterogeneity based on the automaker being affected. Was Hyundai or Kia affected more? These are not robustness checks per se, but they provide further insight to the heterogeneity of our results. Table B.5 examines the heterogeneous treatment effect on transaction prices by automaker. The point estimates suggest a slightly larger effect for Hyundai than Kia, but the difference in the effect between the two is not statistically significant.

In Table B.6, we examine heterogeneous effects on transaction prices by vehicle class. We observe a larger effect for large cars than small cars. For vehicles in the crossover and sport classes, the effect is not statistically significant. Our take-away from this is that large cars and small cars are the dominant force behind the equilibrium price change, which could correspond to consumers interested in these car classes being sensitive to fuel-economy information.

Table B.3: Further Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)
	Logs			Levels		
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	-0.016 (0.005)	-0.010 (0.003)	-0.011 (0.004)	-295 (92)	-279 (89)	-336 (81)
Year-Month \times Class FE	Y	Y	Y	Y	Y	Y
Year-Month \times Make FE	Y	Y	Y	Y	Y	Y
VIN10 FE	Y	Y	Y	Y	Y	Y
DMA FE	Y	Y	Y	Y	Y	Y
$1(Post\ Restatement) \times DMA\ FE$	Y	Y	Y	Y	Y	Y
Include prices \leq \$5,000	Y			Y		
Exclude price outliers		Y			Y	
Hyundais and Kias only			Y			Y
R-squared	0.86	0.98	0.92	0.96	0.98	0.93
N	1.52m	1.48m	0.14m	1.52m	1.48m	0.14m

Notes: Dependent variable is log or level of the transaction price (in dollars). The “exclude price outliers” specification excludes outliers less than 67% of the mean price and greater than 150% of the mean price. An observation is a year-month-DMA-VIN10. VIN10 refers to the VIN prefix, which is a trim-engine combination. DMA refers to a Nielsen Designated Market Area, which is an area covering several counties. Class refers to the vehicle class. *Post Restatement* refers to the year-month being during or after November 2012. All estimations are weighted by monthly sales. Standard errors are clustered at the VIN10 level.

Table B.4: Robustness Check Excluding Minimally Treated Observations

	(1)	(2)	(3)	(4)	(5)	(6)
	Logs			Levels		
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	-0.010 (0.004)	-0.010 (0.004)	-0.011 (0.004)	-147 (84)	-253 (97)	-286 (94)
Year-Month \times Class FE		Y	Y		Y	Y
Year-Month \times Make FE	Y	Y	Y	Y	Y	Y
VIN10 FE	Y	Y	Y	Y	Y	Y
DMA FE	Y		Y	Y		Y
$1(Post\ Restatement) \times DMA\ FE$	Y		Y	Y		Y
R-squared	0.95	0.91	0.95	0.96	0.95	0.96
N	1.51m	1.51m	1.51m	1.51m	1.51m	1.51m

Notes: Dependent variable is log or level of the transaction price (in dollars). An observation is a year-month-DMA-VIN10. VIN10 refers to the VIN prefix, which is a trim-engine combination. DMA refers to a Nielsen Designated Market Area, which is an area covering several counties. Class refers to the vehicle class. *Post Restatement* refers to the year-month being during or after November 2012. All estimations are weighted by monthly sales. Standard errors are clustered at the VIN10 level.

Table B.5: Heterogeneous Effects on Transaction Prices by Automaker

	Primary		Automaker	
	(1)	(2)	(3)	(4)
	logs	levels	logs	levels
$1(PostRestatement)_t \times 1(Affected\ Model)_j$	-0.012 (0.004)	-294 (91)		
$1(PostRestatement)_t \times 1(Hyundai\ Affected\ Model)_j$			-0.014 (0.005)	-365 (123)
$1(PostRestatement)_t \times 1(Kia\ Affected\ Model)_j$			-0.010 (0.004)	-212 (114)
Year-Month \times Class FE	Y	Y	Y	Y
Year-Month \times Make FE	Y	Y	Y	Y
VIN10 FE	Y	Y	Y	Y
DMA FE	Y	Y	Y	Y
$1(Post\ Restatement) \times DMA\ FE$	Y	Y	Y	Y
R-squared	0.95	0.96	0.95	0.96
N	1.52m	1.52m	1.52m	1.52m

Notes: Dependent variable is log or level of the transaction price (in dollars). An observation is a year-month-DMA-VIN10. VIN10 refers to the VIN prefix, which is a trim-engine combination. DMA refers to a Nielsen Designated Market Area, which is an area covering several counties. Class refers to the vehicle class. *Post Restatement* refers to the year-month being during or after November 2012. All estimations are weighted by monthly sales. Standard errors are clustered at the VIN10 level.

Table B.6: Heterogeneous Effects on Transaction Prices by Vehicle Class

	(1) logs	(2) levels
$1(PostRestatement)_t \times 1(SmallCar\ Affected\ Model)_j$	-0.013 (0.005)	-320 (123)
$1(PostRestatement)_t \times 1(LargeCar\ Affected\ Model)_j$	-0.025 (0.004)	-702 (134)
$1(PostRestatement)_t \times 1(Crossover\ Affected\ Model)_j$	-0.007 (0.004)	-190 (98)
$1(PostRestatement)_t \times 1(Sport\ Affected\ Model)_j$	-0.002 (0.005)	239 (220)
Year-Month \times Class FE	Y	Y
Year-Month \times Make FE	Y	Y
VIN10 FE	Y	Y
DMA FE	Y	Y
$1(Post\ Restatement) \times DMA\ FE$	Y	Y
R-squared	0.95	0.96
N	1.52m	1.52m

Notes: Dependent variable is log or level of the transaction price (in dollars). An observation is a year-month-DMA-VIN10. VIN10 refers to the VIN prefix, which is a trim-engine combination. DMA refers to a Nielsen Designated Market Area, which is an area covering several counties. Class refers to the vehicle class. *Post Restatement* refers to the year-month being during or after November 2012. All estimations are weighted by monthly sales. Standard errors are clustered at the VIN10 level.

C Effect of Restatement on Other Outcomes

C.1 Effect on Quantities

As mentioned in Section 4.2, we might not expect to see much of an effect of the restatement on the sales of affected vehicles for several reasons. First, by November 2012, Hyundai and Kia had completed production of all 2011 and 2012 model-year vehicles and had moved on to producing model-year 2013 vehicles. Thus, it would be physically impossible for production of these model-year vehicles to adjust to the restatement. The only quantity adjustment possible would be in dealers shifting sales to a later time, but there are non-negligible inventory costs to holding older model-year vehicles on the dealer lot, which makes this likely an unappealing option for dealers.

In November 2012, Hyundai and Kia were producing their 2013 model-year vehicles, so it is certainly possible that they could adjust production due to the restatement. However, it is important to note that the restatement was a one-time negative shock for specific model-year vehicles, rather than a longer-term change in demand, such as due to changes in gasoline prices. Changing production requires costly physical adjustments to assembly lines and renegotiations of contracts with suppliers. Firms can make these changes in response to major swings in gasoline prices (e.g., see Busse, Knittel, and Zettelmeyer 2013), but it is unclear if they would make these costly changes in response to a one-time negative shock. Furthermore, from a long-run profit maximization standpoint, maintaining market share (and thus minimal changes to production decisions) might well be the optimal managerial strategy in response to such a one-time negative shock. For example, it could allow firms to maintain market share that would be difficult to recapture later, allowing for higher profits in the longer-run, even if short-run profits are reduced. Classic models in industrial organization with switching costs can rationalize such behavior. For example, see discussion of the many possible pricing behaviors firms may display in response to financial distress in Borenstein and Rose (1995).

Because it is possible that Hyundai and Kia adjusted production of affected models

at the same time that prices were adjusted, we also estimate several models exploring the effect of the restatement on sales. Such estimations are likely to provide little useful evidence, since automobile sales are very noisy. For example, model-trims have highly variable temporal phase-in and phase-out patterns and there are niche model-trims that are rarely sold.

Table C.1 confirms our intuition that automobile sales are very noisy. In this table, we estimate a model aggregated at the VIN10-DMA-year-month level and regress the sales of each model on $1(Post\ Restatement)_t \times 1(Affected\ Model)_j$ and the same set of fixed effects that we include in the price regressions in Tables 2-4, which are year-month by class fixed effects, year-month by make fixed effects, VIN10 fixed effects, DMA fixed effects, and post restatement by DMA fixed effects. To address the possibility of niche models that are rarely sold and model-years that are being phased in or phased out (and thus showing large percentage changes in sales) unduly affecting our results, in columns 2 through 5 we present the results where we exclude observations if the monthly sales are less than some percentage of average monthly sales for a particular model-trim. In column 2, that percentage is 50% of monthly sales, in column 3 it is 40%, in column 4 it is 30%, and in column 5 it is 25%.

The coefficients in Table C.1 are all *positive*, suggesting that the restatement *increased* sales, which is a counter-intuitive result. However, they are all imprecisely estimated. We recognize that the lack of a statistically significant effect (either positive or negative) may be due to a lack of power, although all estimations include over three million observations. Note that the highly variable phase-in and phase-out patterns make estimating an effect on quantities especially challenging when there is not a strong signal in the data. In our case, it appears that there is not a strong negative quantity effect of the restatement, as one might have anticipated. In Appendix D.2, we discuss the implications of negative or positive quantity effects on estimates of the valuation of fuel economy. We find our conclusions about undervaluation to be robust to a wide range of quantity effects.

Table C.1: Effect of Restatement on Sales

	(1) Main	(2) <50%	(3) <40%	(4) <30%	(5) <25%
$1(PostRestatement)_t \times 1(Affected\ Model)_j$	0.15 (0.08)	0.05 (0.04)	0.04 (0.05)	0.05 (0.05)	0.06 (0.05)
Year-Month \times Class FE	Y	Y	Y	Y	Y
Year-Month \times Make FE	Y	Y	Y	Y	Y
VIN10 FE	Y	Y	Y	Y	Y
DMA FE	Y	Y	Y	Y	Y
$1(Post\ Restatement) \times DMA\ FE$	Y	Y	Y	Y	Y
R-squared	0.46	0.53	0.53	0.52	0.51
N	4.00m	3.52m	3.62m	3.70m	3.75m

Notes: Dependent variable is log of sales. Columns 2-5 present the results excluding observations if the monthly sales are less than some percentage of average sales, as given in the heading. An observation is a year-month-DMA-VIN10. VIN10 refers to the VIN prefix, which is a trim-engine combination. DMA refers to a Nielsen designated market area, which is an area covering several counties. Class refers to the vehicle class. *Post Restatement* refers to the year-month being during or after November 2012. Standard errors are clustered at the VIN10 level.

C.2 Effect on Advertising

In this subsection, we examine adjustments in advertising by the two affected automakers. For example, the automakers could have increased advertising expenditures to make up for the bad publicity. To examine this, we use data from Kantar Media on advertising expenditures by automaker. In the two figures below, we find no evidence of changes in either advertising expenditures or the number of advertisements by Hyundai and Kia after the restatement. We have also run simple regressions and find no statistically significant effects, with the point estimate quite close to zero. We thus conclude that the quantity of advertising did not change after the restatement.

Of course, Hyundai and Kia are required by law to update any advertisement that specifies the fuel economy of the vehicle, so the content of advertisements must change at least somewhat. This is analogous to the change in advertising around fuel economy that occurs during gasoline price shocks, underscoring that our estimated effect is an equilibrium effect in the same way that the rest of the literature is estimating an equilibrium effect.

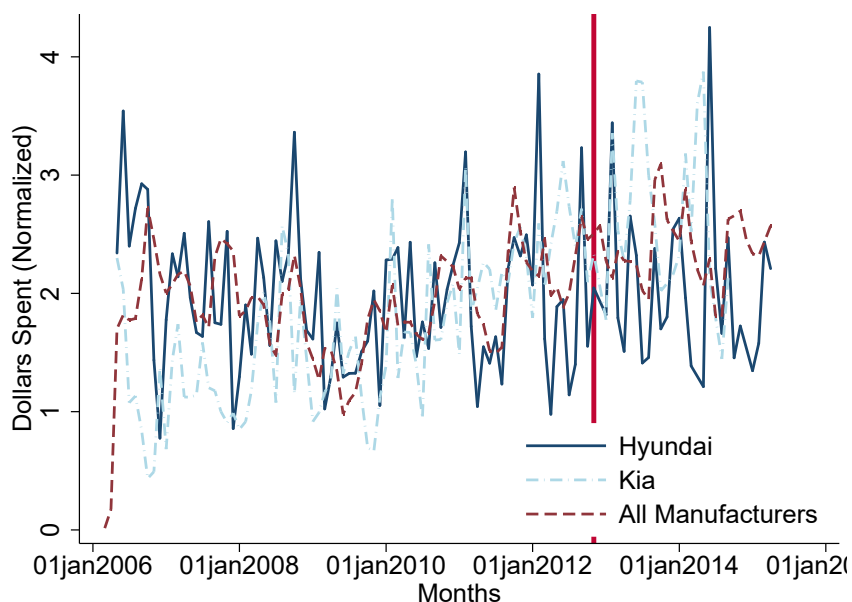


Figure C.1: Spending on Advertising by Different Automakers

Notes: The red line is the date of the restatement.

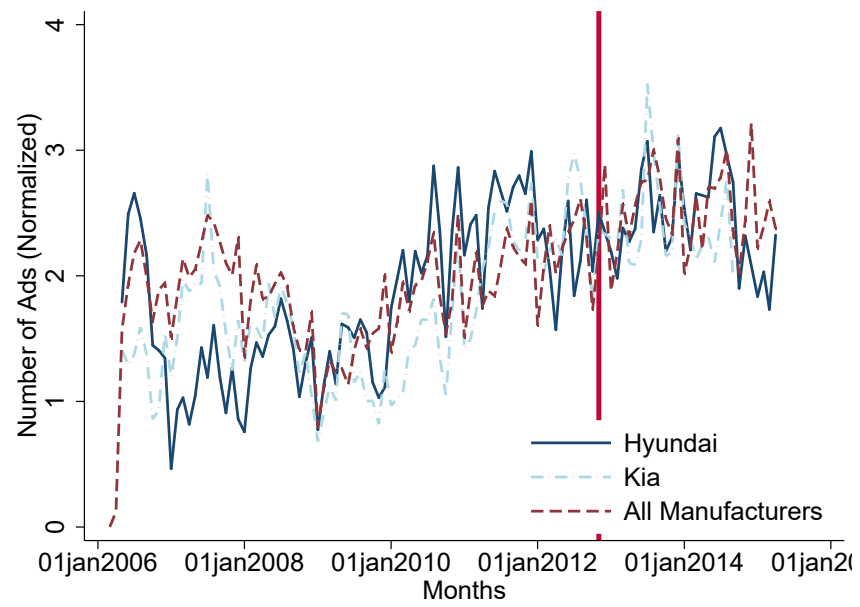


Figure C.2: The Number of Advertisements by Different Automakers

Notes: The red line is the date of the restatement.

D Further Details on the Valuation Calculations

D.1 Motivation from a Discrete Choice Model

This subsection motivates equation (2) from a discrete choice model. For this, we closely follow Allcott and Wozny (2014). The starting point is a random utility model, where the alternative-specific indirect utility of product j at time t , U_{jt} , is a linear function of income (Y), the purchase price (P_{jt}), discounted fuel operating costs (G_{jt}), other controls (X_{jt}), and unobservables ($\tilde{\xi}_{jt}$):

$$U_{jt} = \delta(Y - P_{jt} - \eta G_{jt}) + X_{jt}\beta + \tilde{\xi}_{jt}.$$

With the assumption of an i.i.d Type I extreme value error $\xi_{jt} \equiv \tilde{\xi}_{jt} + \delta Y$, we have a multinomial logit specification, implying that

$$s_{jt} = \frac{e^{U_{jt}}}{\sum_k e^{U_{kt}}},$$

where s_{jt} is average probability of purchase of the representative consumer, or the market share. Further, under this assumption of the errors, we have the standard identity:

$$\log(s_{jt}) - \log(s_{0t}) = -\delta P_{jt} - \theta G_{jt} + X_{jt}\beta + \xi_{jt},$$

where we define $\theta \equiv \delta\eta$. Then in this framework, the definition of the valuation parameter is the ratio θ/δ . This parameter quantifies the trade-off between how consumers value an extra dollar spent on the upfront purchase price (through δ) and a dollar spent on expected future fuel costs (through θ).

To directly estimate this valuation parameter, Allcott and Wozny (2014) invert the market share equation as follows:

$$P_{jt} = \gamma G_{jt} + X_{jt}\tilde{\beta} + \epsilon_{jt}, \tag{D.1}$$

where $\gamma \equiv -\theta/\delta$ is the quantity of interest and the structural error term is $\epsilon_{jt} = \frac{1}{\delta}(\log(s_{0t}) - \log(s_{jt}) + \xi_{jt})$. Similarly, define $\tilde{\beta} \equiv \frac{1}{\delta}\beta$. Note that X_{jt} here contains various controls required for identification, including a variety of fixed effects. In our context, we include year-month by class fixed effects ($\rho_{t \times Class_j}$), year-month by make fixed effects ($\mu_{t \times Make_j}$), region fixed effects (η_r) and their interaction with an indicator for the post restatement period ($\eta_r \times 1(Post\ Restatement)_t$), and VIN10 fixed effects (ω_j). With these fixed effects included, equation (D.1) is effectively the same as equation (2).

Interpreting the estimate of γ as an estimate of the valuation of fuel economy requires that the structural error term is not correlated with the regressors. As defined above, ϵ_{jt} includes the market share at time t for product j . In our setting the identification of γ thus requires that the contemporaneous market shares for each product j should not be correlated with the change in discounted fuel costs induced by the restatement. Our identification argument relies on the timing and one-time nature of the restatement: it was an unexpected shock to the automakers that made it unlikely that they would change product-line production decisions (indeed, for the 2011-2012 model-year vehicles, production line decision changes would be impossible). Furthermore, we examined evidence for quantity adjustments by automakers, and find no support for such adjustments. In comparison, Allcott and Wozny (2014)'s identification argument also relies on a timing assumption, but exploits the fact that the used vehicle market, and in particular scrap rates, should not vary in response to gasoline prices. This assumption is most likely to hold for vehicles that are not too old.

Another key difference between our estimation and the empirical strategies that have recently been used in this literature is that our estimating equation, unlike equation (D.1), does not use the level of discounted fuel costs as a regressor (captured by the variable G), but the difference in G induced by the restatement. Therefore, because the restatements were known and salient, measurement error in the level of G is not a major concern in our setting. This is in contrast to Saltee, West, and Fan (2016) and Allcott and Wozny (2014), whose empirical strategy essentially requires constructing the average G that each

consumer faces, which will be a noisy estimate of its true value. Allcott and Wozny (2014) address the issue using an instrumental variables strategy. Despite not explicitly including an estimate of G in their estimation, Busse, Knittel, and Zettelmeyer (2013)'s empirical strategy is also prone to measurement error due to the fact that they must impute the average gasoline price that each consumer faces. They show that this issue is not important in their setting by using different levels of aggregation in average gasoline prices.

Our natural experiment and approach allow us to circumvent the measurement error issue to a certain extent by focusing on estimating the behavioral response to a change in G induced by the restatement and publicized by the EPA, which is perfectly observed. Note that the size of the change in G that each consumer faced is, of course, function of the gasoline prices consumers paid, driving behavior, and other assumptions required to construct G . We show, however, that our estimate of the valuation parameters are robust to these assumptions (Table D.2).

D.2 Bounding Analysis

In Section 4.2 and Appendix C.1, we discuss why the supply of affected vehicle models may be inelastic due to physical limitations and costly adjustments in the context of our one-time restatement shock. In Table C.1, we also explore the effect of the restatement on sales, but our estimates are too noisy to rule out either a positive or negative quantity effect. In this appendix, we discuss how either a positive or a negative quantity effect affects our estimates of the valuation parameter. The key insight from this discussion is that our primary conclusions about substantial undervaluation hold for a wide range of supply elasticities.

How the equilibrium price effect translates to an estimate of consumers' willingness-to-pay for fuel economy depends on the relationship between demand and supply. Figure D.1 presents three possible scenarios. In all three cases, the effect of the restatement is represented by a downward shift in the demand curve for the affected models—if consumers value fuel economy, a downward adjustment of the miles-per-gallon rating reduces the

private consumer surplus that consumers would expect from such models. When the supply is perfectly inelastic (Panel A), the downward shift in demand results in a change in equilibrium prices that is exactly equal to the willingness-to-pay for fuel economy. In our setting, we argue that this is a likely scenario. Given the unexpected nature of the restatement and the limited scope for production adjustments (especially for the 2011-2012 model-years), the supply curve for affected models should have been almost completely inelastic, and thus the change in equilibrium price and willingness-to-pay should nearly coincide.

When supply is elastic in the expected direction (Panel B), the change in equilibrium price underestimates the change in willingness-to-pay. In the unlikely case that the supply curve is downward sloping (Panel C)—for instance, because of strong economies of scale in the production process—the opposite holds: the change in equilibrium price overestimates the change in willingness-to-pay. Note that the point estimates in Table C.1 point to this case. In these two scenarios, the magnitude of the discrepancy between the change in equilibrium price and the true willingness-to-pay for fuel economy depends on the supply and demand elasticities.

To bound our estimates, consider the following scenario where we first consider highly elastic supply that would result in a large quantity effect. Based on our quantity regressions, we begin by assuming a reduction in quantity of 5%, which is approximately the lower bound of the 95% confidence interval we report in Table C.1 (column 2). We also have to assume a price elasticity of demand. Berry, Levinsohn, and Pakes (1995) find own-price demand elasticities ranging to -6.5, while Busse, Knittel, and Zettelmeyer (2013) consider demand elasticities that range from -2 to -5, in part based on Berry, Levinsohn, and Pakes (1995)'s estimates. Hyundai and Kia are in the smaller car segment of the market, so one might expect more elastic demand, which would suggest a number closer to -6. Accordingly, we first calculate our estimates using a demand elasticity of -6, but we also complete the analysis using a lower estimate of -4. We also need to assume an average vehicle price pre-restatement, and for this we use \$24,500 for our illustrative

calculation (this is calculated as $\$294/0.012$ for consistency with our main results; it is also reasonably closely aligned with the summary statistics on vehicle prices for Hyundai and Kia). Note that when we use a smaller number for the pre-restatement price, such as $\$20,000$, the range of results narrows substantially, so our bounding analysis is conservative in this sense.

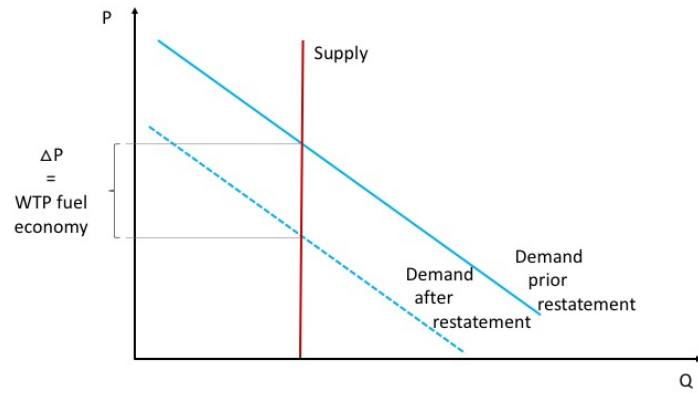
We can then make a set of simple calculations directly applying the demand elasticity to give us the difference between the post-restatement price ($\$24,500 - \294) and what the price would have been had the quantity not changed. Conceptually, we are just moving along the demand curve by the percentage change in quantity.²⁶ Using these assumptions and the $\$294$ reduction in equilibrium price due to the restatement, we find the willingness-to-pay for the 5% reduction in quantity is $\$496$ when using a demand elasticity of -6 and $\$597$ when using an elasticity of -4 . This is roughly a doubling of the estimated equilibrium price change.

In the unlikely case that economies of scale are such a dominant force that they induce a downward sloping supply curve, the equilibrium change in price is an overestimate of the willingness-to-pay. For example, suppose that we observe a positive quantity effect of $+5\%$ (as our quantity estimation point estimates suggest). Then a $\$294$ reduction in equilibrium price corresponds to a willingness-to-pay of only $\$92$ when using a demand elasticity of -6 and is even below zero when using a demand elasticity of -4 . Overall, these illustrative bounds suggest that the valuation parameters that we estimate could be either twice as large or close to zero for these particular scenarios.

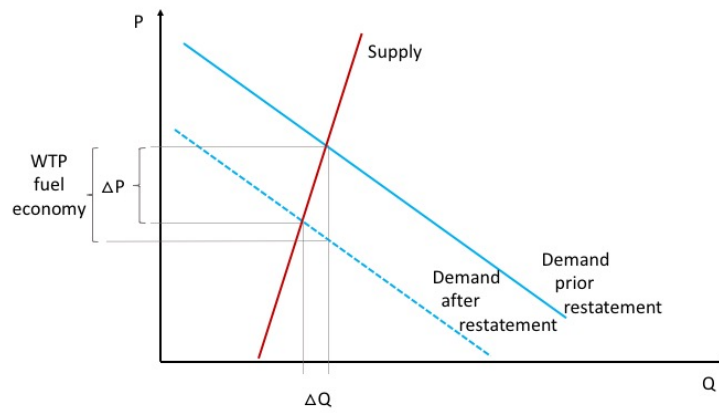
Note that a quantity effect of $+5\%$ in this context is quite large relative to the $\$294$ price reduction, which is slightly over 1% of the transaction price. Under more reasonable scenarios for our setting where the quantity effects are in the $\pm 1\%$ range, a $\$294$ reduction in equilibrium price translates in a willingness-to-pay of either $\$334$ or $\$254$ (under the -6 elasticity), respectively, which are much tighter bounds.

Table D.1 summarizes these results. Combined with Table 4, these results demonstrate

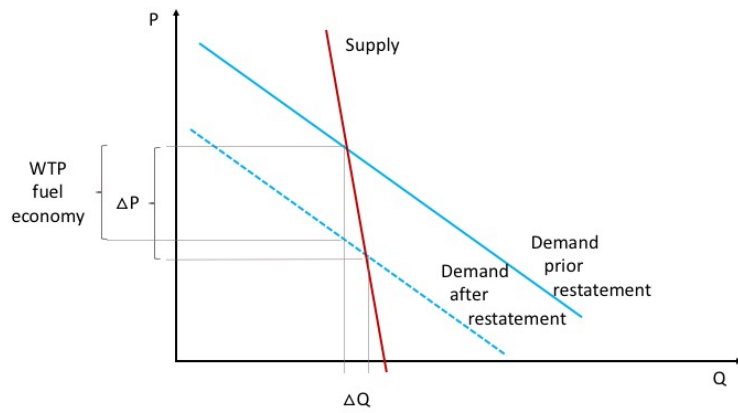
²⁶One can show that the total willingness-to-pay is given by $\Delta P + \Delta P \times \frac{\Delta Q}{\eta_D}$, where η_D is the price elasticity of demand.



(a) Inelastic Supply



(b) Upward Sloping Supply



(c) Downward Sloping Supply

Figure D.1: Interpretation of the Equilibrium Effect

Notes: Each panel presents a particular scenario with respect to the slope of the supply curve and how it impacts the interpretation of the equilibrium price effect.

Table D.1: Interpretation of Equilibrium Change in Prices w.r.t. Different Supply Curves

Quantity Effect (%)	Willingness-to-Pay (\$) $\eta_D = -6$	Willingness-to-Pay (\$) $\eta_D = -4$
-5	496	597
-1	334	355
0	294	294
1	254	233
5	92	-9

Notes: The table shows how a given equilibrium change in price translates into willingness-to-pay (WTP) for fuel economy. η_D refers to the price elasticity of demand we use in our calculations. For all rows, we use an equilibrium change in transaction prices of \$294, following our primary results. These illustrative calculations are also based on an average pre-restatement price of \$24,500.

that our main conclusions about substantial undervaluation hold up to a wide range of supply elasticities, and hence, quantity effects. For instance, consider the pooled sample and a 12% discount rate. Further, suppose that the supply is highly elastic such that it translates in a doubling of the valuation parameter, as it does in Table D.1 for the willingness-to-pay estimate assuming a demand elasticity of -4 (for a demand elasticity of -6 the valuation parameter would need to increase by 69%). Then the valuation parameter would be exactly 0.5, because our estimate is 0.25. For our preferred case of a 4% discount rate, a doubling of the valuation parameter would correspond to 0.34. For the valuation parameter for the 2011 and 2012 model-years, a doubling of the estimate would yield a value of 0.76. Of course, for those models a highly elastic supply is very unlikely.²⁷ Assuming a supply elasticity closer to zero, the effect on the valuation parameter should be much more modest. In Table D.1, a quantity effect of -1% leads to an underestimate of the willingness-to-pay by only 12% using $\eta_D = -6$ (this comes from $(294-334)/334$) or 17% using $\eta_D = -4$. Applying this magnitude of bias to the valuation parameter for the model-years 2011 and 2012 would again imply valuation of fuel economy below 0.5.

²⁷As discussed earlier, the supply for model-year 2011 and 2012 should be inelastic given the impossibility of adjusting the production of a model-year that has finished its production cycle and high cost of holding vehicles in inventory on the dealer lot.

D.3 Sensitivity Analysis of the Valuation Parameter

To estimate the valuation parameter, we need to construct the discounted future fuel costs of each model in our sample. This requires making assumptions about how consumers discount the future, drive their vehicles, forecast gasoline prices, and how long they expect their vehicles to last. We conduct fairly exhaustive sensitivity analysis for all these parameters to investigate the robustness of our results.

Table D.2 outlines the various sensitivity tests we have conducted, data sources, and comparisons with other studies. We find that the discount rate is the variable having the most important effect on valuation. We consider different data sources for gasoline prices. We further consider different scenarios where expected gasoline prices are being held constant in real terms at the levels at the time of registration. This martingale assumption implies that consumers use today's price as a forecast of future prices for the entire lifetime of their vehicle. We consider average price at the annual-national level, annual-state level, month-national level, and at the month-national level without seasonal trends. We also consider a scenario where we remove all variation in gasoline prices and use the gasoline price for the year 2012, 2013, 2014, the average of 2012 and 2013, or the average of 2012, 2013 and 2014 as the constant gasoline price that consumers use in their forecasting. Finally, we consider a scenario where consumers are able to make a perfect forecast of future gasoline prices, where we use realized prices up to 2017 and then the Energy Information Administration's forecasted gasoline prices for the other future years. Compared to previous studies, our different scenarios about expectations of gasoline prices broadly cover the range of assumptions that has been used. For instance, Busse, Knittel, and Zettelmeyer (2013) and Sallee, West, and Fan (2016) both use the martingale assumption. Allcott and Wozny (2014) use the martingale assumption, but also consider a scenario where consumers based their expectations on oil futures.

For vehicles' survival probabilities, we estimate the results separately using the data from Jacobsen and van Benthem (2015) and Busse, Knittel, and Zettelmeyer (2013), the latter of which were derived from the National Household Travel Survey (NHTS). Data

for vehicle miles travelled come also from the NHTS. We compare results using the 2006 or the 2017 wave of the NHTS.

Table D.2: Sensitivity Analysis: Valuation Parameters

Discount Rate	Gasoline Prices	VMT	Survival Probability	Ratio of Means	Valuation Parameter	Valuation Parameter: 2012 Model-Year Only
4%	Ann.-Nat.	NHTS 17	JvB	No	0.167	0.379
1%	Ann.-Nat.	NHTS 17	JvB	No	0.139	0.316
12%	Ann.-Nat.	NHTS 17	JvB	No	0.246	0.559
4%	2012-Nat.	NHTS 17	JvB	No	0.160	0.370
4%	2012-2014-Nat.	NHTS 17	JvB	No	0.169	0.389
4%	Month-Nat.	NHTS 17	JvB	No	0.165	0.391
4%	Ann.-State	NHTS 17	JvB	No	0.195	0.369
4%	Ann.-Nat.	NHTS 06	JvB	No	0.142	0.323
4%	Ann.-Nat.	NHTS 17	NHTS	No	0.172	0.390
4%	All	NHTS 06/17	NHTS/JvB	No	[0.137-0.195]	[0.315-0.422]
4%	Ann.-Nat.	NHTS 17	NHTS	Yes	0.419	0.855

Notes: Valuation parameters presented for different assumptions pertaining to the construction of the discounted fuel costs. Different levels of aggregation are considered for gasoline prices. “Ann.” refers to annual data. “Nat.” refers to national-level data. The row with “2012-National” uses the average U.S. nationwide gasoline price for the year 2012.: 3.68 USD/gallon. Similarly, the row with “2012-2014-National” uses the average U.S. nationwide gasoline price, where the average is taken over the years: 2012, 2013 and 2014: 3.56 USD/gallon. In those two scenarios, there is no variation in discounted fuel costs induced by gasoline prices. The VMT estimates are based on the NHTS survey. We use the data for the survey years 2006 or 2018. For the survival probabilities, we use the NHTS data as reported by Busse, Knittel, and Zettelmeyer (2013). We also consider the estimates provided by Jacobsen and van Benthem (2015). In the last row, we report the valuation parameters using the approximation that relies on the ratio of the mean change in prices over the mean change in discounted fuel costs. This approximation has a large impact on the valuation parameter and leads to an upward bias. The discount rate and whether we solely rely on the 2012 model-years are the two dimensions that induce the most variation in the results. The data source for the VMT, survival probabilities, and the level of aggregation in the gasoline prices have little effects on the results.

D.4 Bias from the “Ratio of the Means” Approximation

With the setup in Section D.1 above, it is even easier to understand the ratio of the means issue referred to in the main text. Before moving to the equations, it is illustrative to begin with a simplified example to fix ideas. Suppose that two different vehicle models were subject to a restatement in fuel economy: Model A, which has a price of \$50,000, and Model B, which has a price of \$10,000. Also, suppose that both models are equally popular, so we can ignore their relative market shares in this example. When the unexpected restatement occurs, this changes consumer expectations about the future fuel costs of each of the two vehicles. Suppose the restated EPA fuel-economy ratings correspond to

a change in discounted lifetime fuel costs of \$5,000 for Model A and \$1,000 for Model B. We are then interested in how the equilibrium prices and quantities change. Suppose that sales are held constant. And further suppose that the restatement leads to heterogeneous changes in equilibrium prices: \$5,000 for Model A, but only \$100 for Model B.

The valuation parameter implied by this illustrative event is $\$5,000/\$5,000 = 1$ for Model A and $\$100/\$1,000 = 0.1$ for Model B. The mean of the valuation ratio is thus the average of 1 and 0.1, which equals 0.55. This is the exact valuation parameter when both models are equally popular. Now consider the approximation, which is the ratio of the mean of the changes in prices over the mean of the changes in future fuel costs: $\$2,550/\$3,000 = 0.85$. The intuition is that the naive approximation puts too much weight on changes in the numerator or denominator that are large in absolute value.

The intuition for the issue should be clear: the ratio of the means is not necessarily the same as the mean of the ratios. Houde and Myers (2019) analyze the appliance energy efficiency context and show the conditions under which we would expect a bias more generally, and what the sign of the bias might look like. The insights from the appliance energy efficiency context carry over to our setting as well. To see the issue mathematically, note that the goal in estimating (D.1) is to consistently estimate the true value of γ . Consider a case where there is heterogeneity over vehicles in γ , so that we can write it as γ_j .²⁸ Our simple example above was one case where there was heterogeneity in γ across vehicles, and we showed in Table B.6 that there is heterogeneity across car classes, so we know that empirically there is indeed heterogeneity in γ across vehicles. We are interested in the mean effect, or $E[\gamma_j]$, where the mean is taken over the population of vehicles. However, by definition, this is the mean of a ratio: $E[\gamma_j] = E[\theta_j/\delta_j]$.

To see how this true value (a mean of a ratio) relates to the approximation (a ratio of means), consider the second-order Taylor expansion:

$$E[\theta_j/\delta_j] \approx E[\theta_j]/E[\delta_j] - \text{cov}(\delta_j, \theta_j)/E[\delta_j]^2 + \text{Var}(\delta_j)E[\theta_j]/E[\delta_j]^3.$$

²⁸Houde and Myers (2019) consider heterogeneity over consumers, so it is γ_i being considered, but the same logic follows as here.

Thus, the value of interest $E[\gamma_j]$ (the mean of the ratio) is only equal to $E[\theta_j]/E[\delta_j]$ (the ratio of the means) when the covariance and variance terms in the equation are equal to zero (this is a slightly weaker condition than assuming no heterogeneity in γ_j). Our results indicate that there is heterogeneity in γ_j and our calculations showing a difference in the results between the two approaches suggest that the higher order terms in the approximation are important.

Note that several papers in the literature that aim to estimate $E[\gamma_j]$ report a ratio that corresponds to $E[\theta_j]/E[\delta_j]$, as they separately estimate $E[\theta_j]$ and $E[\delta_j]$. This is true for studies that rely on reduced form methods (Busse, Knittel, and Zettelmeyer 2013; Leard, Linn, and Zhou 2018) and a similar issue could arise using structural methods (Grigolon, Reynaert, and Verboven 2018). A key point is that when there is heterogeneity across the population and a correlation between the response in upfront purchase price and the response in future fuel costs, this correlation will lead the exact measure of undervaluation to deviate from the approximation. If there is a positive correlation between θ and δ (e.g., vehicles for which consumers really do not like a change in upfront purchase price are more likely to be vehicles for which consumers really do not like a change in future fuel costs), then this equation would predict that the approximation would be biased upwards in terms of the valuation.²⁹

²⁹It is a bias upwards in the valuation because we subtract off the covariance term, so the coefficient becomes more negative, which means less undervaluation (recall -1 means full valuation, while zero means not valuing future fuel costs at all).