

Consumer Myopia in Vehicle Purchases: Evidence from a Natural Experiment*

Kenneth Gillingham

Yale University

Sébastien Houde

Grenoble Ecole de Management

Arthur A. van Benthem

University of Pennsylvania

September 18, 2019

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 that provides exogenous variation in fuel-economy ratings. We examine the short-run equilibrium effects of a 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 16-39 cents in the purchase price when discounting at 4%. This undervaluation persists under a wide range of assumptions.

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

JEL classification codes: D12, H25, L11, L62, L71, Q4

*Gillingham: Yale University, 195 Prospect Street, New Haven, CT 06511, United States, phone: +1-203-436-5465, e-mail: kenneth.gillingham@yale.edu. Houde: Grenoble Ecole de Management, 12 Rue Pierre Semard, 38000, Grenoble, France, e-mail: sebastien.houde@grenoble-em.com. van Benthem: The Wharton School, University of Pennsylvania, 1354 Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104, United States, phone: +1-215-898-3013, e-mail: arthurv@wharton.upenn.edu. The authors are grateful to Hunt Allcott, Mark Jacobsen, Chris Knittel, Matt Kotchen, Josh Linn, Benjamin Leard, Arik Levinson, Erica Myers, Mathias Reynaert, John Rust, Jim Sallee and Joe Shapiro for their comments and suggestions. We also thank seminar participants at the 2019 AERE conference, the 2019 EMEE workshop, the 2019 Northeast Workshop on Energy Policy and Environmental Economics, the 2019 Workshop on Climate Economics, Innovation and Policy, Georgetown University, RWI Essen, Simon Fraser University, Stanford University, University of Connecticut, University of Copenhagen, University of Maryland, University of Pennsylvania, VU Amsterdam, 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 apparent 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 to available information 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 from a natural experiment providing exogenous variation in the fuel-economy ratings

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

²We follow a common terminology in the existing literature (e.g., Hausman 1979; Busse, Knittel, and Zettelmeyer 2013) and use the term “myopia” to describe a range of behavioral phenomena leading to undervaluation, which could include biased beliefs, lack of salience, rational inattention, and present bias.

³Our study also relates to papers that 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).

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 to 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 16-39 cents in the vehicle purchase price (a “valuation parameter” of 0.16-0.39) depending on the affected model year, and using a discount rate of 4%. We 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 and the presence of imperfect competition, as we illustrate in a bounding exercise.

Previous studies estimating the consumer valuation of fuel economy use several different identification strategies, but most leverage changes in gasoline prices to test

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).

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 in future fuel costs 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 Sallee, 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) employ cross-sectional variation in engine technologies and find even lower values; most of their estimates are below 0.30.

We contribute to this literature by bringing a new identification strategy to shed light on this unsettled question. An appealing feature of using our natural experiment to understand consumer valuation of future fuel costs is that we can rely on a sudden and exogenous shifter of the official fuel-economy rating, yet be assured that the vehicles

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

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

Our estimates are especially relevant for informing the intense debate on whether fuel-economy standards are justified from a private perspective.⁷ If consumers 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 by other policies. We use a novel approach to provide guidance to policymakers on this critical parameter for understanding the costs and benefits of fuel-economy standards. Our natural experiment—a revision of fuel-economy ratings—may be particularly relevant to studying more stringent fuel-economy standards, as consumers would be informed of the higher fuel economy through the ratings.

We also contribute by highlighting two new issues in this literature that help reconcile discrepancies across estimates. First, we demonstrate the quantitative importance of estimating a fuel-economy valuation parameter directly, rather than approximating it using average changes in equilibrium prices, quantities, and discounted changes in fuel expenditures—an approach commonly taken in the literature. In our sample, the approximation yields a valuation parameter that is more than double the correct value, which is large enough to substantially alter the conclusions of a valuation study. Second, we show that if there is market power in the automobile market, willingness-to-pay estimates that ignore this will overestimate the valuation of fuel economy.

Our undervaluation result suggests that a variety of behavioral channels may be at play, although we cannot quantify their relative importance. Inattention to fuel-economy

⁷In the U.S., the Trump Administration is in the process of weakening the standards based on a benefit-cost analysis that explicitly incorporates assumptions about the degree of consumer valuation of fuel economy (Bento et al. 2018). See <https://www.nytimes.com/2018/08/02/climate/trump-auto-emissions-california.html>

ratings, a lack of sophistication to correctly process fuel-economy information, and a variety of (incorrect) beliefs about fuel economy potentially paired with slow updating towards the true value are all possible explanations for why consumers on average are not willing to pay the full discounted benefits of higher fuel economy vehicles.

The remainder of this paper is organized as follows. We next describe the natural experiment. In Section 3, we discuss the data. Section 4 presents the empirical strategy and main results that show how the market responded to the information shock provided by the restatement. In Section 5, we estimate consumers' valuation of future fuel costs and discuss and interpret our estimates. The final section concludes.

2 The 2012 Fuel-Economy Rating Restatement

In many countries around the world automakers are required to report the fuel-economy performance of all new vehicles offered on the market. In the United States, this reported value is randomly audited by the EPA and considered a reasonable estimate of the true on-road fuel economy of the vehicle. This EPA rating plays a prominent role: it is used by automakers in advertising, is used in auto-shopping websites, and is required to be conspicuously displayed on every new vehicle at the dealer lot as part of an EPA fuel-economy label.⁸

On November 2, 2012, the EPA issued a press release stating that “in processing test data, Hyundai and Kia allegedly chose favorable results rather than average results from a large number of tests.”⁹ 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-

⁸See Appendix A for more details on the ratings and the label, including an example label. Note that the EPA ratings are different from the compliance ratings for the CAFE fuel-economy standards. These compliance ratings are based on a laboratory test established in 1978. The EPA revised the consumer ratings downward in 1986, and again in 2008, to more accurately reflect real-world driving conditions and fuel economy. However, to determine automakers' compliance with CAFE the government continues to use fuel-economy values based on the 1978 test procedure.

⁹The 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>.

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.”¹⁰

Immediately after the EPA press release, the fuel-economy ratings for all affected vehicles 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. 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. Tables A.1 and A.2 in Appendix A provide a list of the restated models and the change in miles-per-gallon for each. Combined ratings, which reflect an average of city and highway driving, were adjusted downward by up to four miles-per-gallon; highway ratings went down by up to six miles-per-gallon.

Prior to the restatement, Hyundai and Kia often mentioned the high fuel economy of their vehicles as a selling point.¹¹ 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.

¹⁰See <https://www.autoblog.com/2014/11/03/hyundai-kia-300-million-mpg-penalties/>.

¹¹Consider 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.”

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 nearly all dealer and manufacturer incentives.¹² The data do not allow us to observe movement on other dealer’s margins, such as preferential 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 bring in data on annual

¹²The data include all incentives that are reported for sales tax purposes, which include all dealer incentives and most manufacturer incentives. E.g., for California see <https://www.cdtfa.ca.gov/formspubs/pub34.pdf>.

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

	Affected Models		Not Affected Models		
	Hyundai	Kia	Hyundai	Kia	Others
	(1)	(2)	(3)	(4)	(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.

nationwide gasoline prices from the U.S. Energy Information Administration (EIA), on vehicle survival rates from Jacobsen and van Benthem (2015), and on average vehicle miles traveled from the 2017 National Household Travel Survey (NHTS). In sensitivity analysis, we also provide estimates for miles driven and survival rates using the 2006 NHTS, following Busse, Knittel, and Zettelmeyer (2013) as well as EIA’s gasoline prices at the monthly-national level and at the year-state level.

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:

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

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. Our parameter of interest, β , is the coefficient on the interaction of these two indicator variables. Our 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. We 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 trends or shocks that equally affect all models from each automaker. We include DMA indicators (η_r) and their interaction with the post-restatement indicator

$(\eta_r \times 1(\text{Post Restatement})_t)$ to control for potential compositional changes in the population of consumers buying a vehicle before and after the restatement. Finally, ω_j are VIN10 fixed effects.¹³ We weight the regressions by monthly sales¹⁴ and cluster standard errors at the VIN10 level.¹⁵ Finally, 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.

In our specification, 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. Therefore, β is capturing the effect of the restatement on the affected models—our desired effect—rather than any diminished brand perception from the restatement that affects all Hyundai and Kia models equally (such effects on the brands would be captured by $\mu_t \times \text{Make}_j$). One advantage of this specification is that it readily facilitates exploring different sources of variation to identify β . In our primary specification, we include all non-affected models in the relevant vehicle classes, but we also examine cases where we remove close substitute non-affected vehicles from the sample (to test for robustness to price spillovers within or across brands) or remove all other automakers besides Hyundai and Kia (to further confirm that effects on brand equity are not influencing our results).

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. 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 price level. Columns 3 and 6 are the most flexible and therefore our preferred specifications. The coefficients become

¹³Our identification follows recent studies that use disaggregated panel data. For example, Allcott and Wozny (2014) and Busse, Knittel, and Zettelmeyer (2013) use monthly temporal variation in gasoline prices after conditioning on model year fixed effects. Sallee, West, and Fan (2016) exploit variation in odometer readings within a model year while controlling for VIN10-year-month.

¹⁴This is equivalent to running regressions at the microdata level (i.e., every car sale is a separate observation).

¹⁵Clustering at the VIN10 level allows for arbitrary forms of serial correlation patterns in the error terms, both over time and across DMAs. In addition, the treatment is (approximately) at the VIN10 level. Clustering at the model level generates very similar, and often slightly smaller, standard errors in Table 2.

Table 2: Effect of Restatement on Transaction Prices

	(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)	-0.012 (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

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.

slightly larger 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.

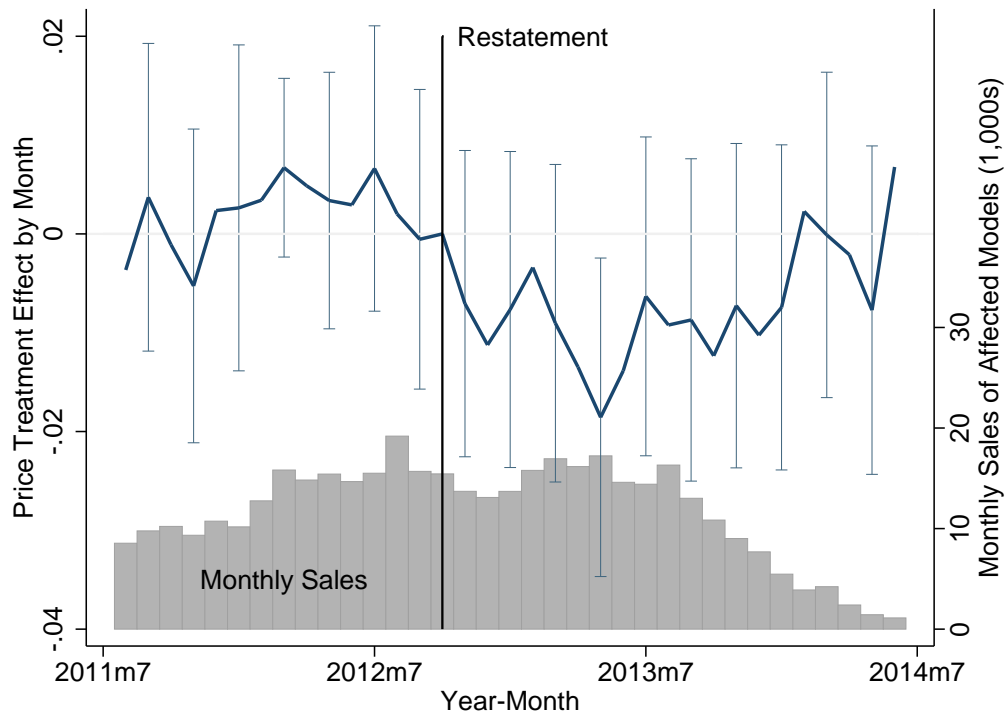


Figure 1: The Price Effect of the Restatement on Affected Models by Month Along with the Monthly Sales of Affected Models

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.

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).¹⁶ SUTVA can be violated in our context if there are spillovers between the treated and control (e.g., from strategic pricing in a market with differentiated products, either by Hyundai and Kia and/or by their competitors) or if there are general equilibrium effects due to the treatment, such as broader effects on the Hyundai and Kia brands. For example, suppose Hyundai and Kia recognize that demand for the affected vehicles would decrease, leading to an increase in demand for close substitutes. If the firms are profit-maximizing, they may find it beneficial to increase the price of their non-affected close substitutes. This would imply that our estimated coefficients would be *overestimates* of the effect of the restatement on the equilibrium prices (and later, as we will see, on the valuation of fuel economy, implying that such spillovers to close substitutes would lead to even greater undervaluation of fuel economy than we estimate). The same situation could also occur with close substitutes from other automakers.

We perform several robustness checks to exploit different sources of variation to confirm that SUTVA holds in our case. Table 3 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.

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

¹⁶The 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.

Table 3: Robustness Checks for SUTVA Assumption

	(1)	(2)	(3)	(4)	(5)	(6)
	Logs			Levels		
$1(\text{Post Restatement})_t \times 1(\text{Affected Model})_j$	-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.

data from Edmunds.com and MotorTrend.com.¹⁷ 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 Table 2. The coefficients excluding substitutes are all close to our primary specification, indicating that the slight change in the competitive landscape from the restatement had little influence on the pricing of close substitute models.

In Appendix B, we explore alternative sets of fixed effects and find that the results are robust. These alternative fixed effects slightly change the variation being used to identify our coefficients. Specifically, Appendix Table B.1 includes sets of vehicle class fixed effects where we use finer or coarser definitions of vehicle class, which essentially changes how we control for the relative time trends in the prices of affected and non-affected vehicles. We find that our results are highly robust to all of these alternative specifications. Ap-

¹⁷Edmunds.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.

pendix Table B.2 also adds quarter-of-age \times make fixed effects to capture the cyclical nature in the vehicle market that depends on the time since a vintage of a vehicle was introduced to the market; this hardly changes the estimates.

The robustness checks so far confirm that spillover effects to close substitutes appear to be limited, with relatively small changes in our estimated equilibrium price response across the checks. In addition to effects on close substitutes, one might also be concerned that the widely-publicized restatement had an effect on the overall Hyundai and Kia brand equity. If the overall brand equity for the two automakers is affected, then the equilibrium prices may be changing due to a diminished brand perception that affects all Hyundai and Kia models in addition to the response to the lower fuel-economy ratings on the affected models. As explained above, our year-month \times automaker indicator variables assure that we are exploiting variation across affected and unaffected vehicles after conditioning on a common price trend for each automaker, so this concern should not affect our estimates of interest. To provide further support that this is not a concern, we also estimate the model removing all other automakers besides Hyundai and Kia, so that we are exploiting *only* variation within the two automakers across affected and non-affected vehicles. We again find very similar results. This estimation, along with further robustness checks on sample selection, can be found in Appendix Tables B.3 and B.4.

Finally, Appendix Table B.5 shows how the point estimates change when we omit up to 12 months following the restatement. The reason we might want to do this robustness check is two-fold. First, one might be worried about an unusual selection of car buyers for the affected models just after the restatement. Presumably this would dissipate for new car purchases several months later. Second, it is possible that new car buyers just after the restatement base their decision (at least in part) on the earlier ratings they had seen prior to the restatement when they compared vehicles in preparation for the purchase (see Section 5.3 for a detailed discussion about consumer beliefs about fuel-economy ratings and realized fuel economy). As more months pass, it becomes increasingly unlikely that new car buyers are basing their decision on the older ratings. The estimates in Appendix

Table B.5 are quite similar to the full-sample estimate. Not surprisingly, standard errors increase as we shrink the sample. Removing transactions close to the restatement date ensures that the effect is coming from new car buyers who were unlikely to have seen both the pre- and post-restatement fuel-economy ratings—in other words, they are unlikely to respond to the *change* in ratings but rather process the level of the new, lower, rating only.

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 change in the fuel-economy rating. In Table 4, we explore heterogeneous treatment effects with respect to these variables.¹⁸ 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.¹⁹ The negative coefficient indicates that the price reductions are larger for models that faced a greater reduction in fuel economy (i.e., an 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 but do respond in the expected direction on average.

¹⁸Appendix Tables B.6 and B.7 explore heterogeneity by make and vehicle class.

¹⁹The results have nearly identical implications if we use miles-per-gallon.

Table 4: Heterogeneous Effects of the Restatement on Transaction Prices

	Primary		Model Year		Δ GPM	
	(1) Logs	(2) Levels	(3) Logs	(4) Levels	(5) Logs	(6) Levels
$1(\text{Post Restatement})_t \times 1(\text{Affected Model})_j$	-0.012 (0.003)	-294 (91)				
$1(\text{Post Restatement})_t \times 1(\text{2011} - \text{2012 Affected Model})_j$			-0.017 (0.006)	-544 (128)		
$1(\text{Post Restatement})_t \times 1(\text{2013 Affected Model})_j$			-0.011 (0.004)	-259 (98)		
$1(\text{Post Restatement})_t \times 1(\text{Affected Model})_j \times \Delta\text{GPM}$					-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
$1(\text{Post Restatement}) \times \text{DMA FE}$	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.

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. Such quantity adjustments are important, since they affect how our estimates translate into the willingness-to-pay for fuel economy, and thus our conclusions about undervaluation. We therefore carefully consider how quantity effects affect our calculations of consumer valuation in Section 5 below.

First, it is important to point out that our setting is quite different from Busse, Knittel, and Zettelmeyer (2013). By November 2012, automakers had already completed production of model year 2011 and 2012 vehicles and had moved on to producing model year 2013 vehicles. All remaining vehicles from model years 2011 and 2012 were already on dealer lots. Thus, it would be physically impossible for production of this vintage to adjust to the restatement. The only quantity adjustment possible would be in dealers

shifting sales to a later time. But this is likely to be an unappealing option for dealers because of non-negligible inventory costs from holding older model year vehicles on the dealer lot.

Model year 2013 vehicles were still midway through their production cycle at the time of the restatement. It is certainly possible that Hyundai and Kia could adjust production of these 2013 vehicles due to the restatement. However, such adjustments in production are typically costly. They require physical adjustments to assembly lines and renegotiation of contracts with suppliers. Therefore, supply was very likely inelastic for model year 2011 and 2012 vehicles but possibly somewhat more elastic for model year 2013 vehicles.

Estimating the equilibrium effects of the restatement on quantities turns out to be challenging in our context. In Appendix C.1, we examine quantity responses 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. Appendix Table C.1 shows that, in our preferred specification, the estimated effect of the restatement on sales is 0.05 (standard error 0.04). While we can only take this noisy evidence as suggestive, we certainly do not find clear evidence for a negative equilibrium quantity effect. Still, our noisy estimates do not entirely rule out substantial negative quantity effects. Fortunately, we do not have to take a strong stance on the magnitude of the quantity response for our key conclusion about substantial undervaluation to hold, as we will show in detail in Section 5.2.

Besides effects on sales, another possible adjustment in response to the restatement 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 fuel 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.

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:

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

where $Price_{jrt}$ is the vehicle transaction price and ΔG_{jt} is the change in the discounted lifetime fuel cost due to the restatement.²⁰ In Appendix D.1, we motivate Equation (2) from a random utility model and show that γ can be interpreted as the valuation parameter, which quantifies how consumers trade off discounted future energy operating costs with the purchase price.²¹ If sales do not adjust, we can interpret a value of -1 as

²⁰Note $\Delta G_{jt} = 0$ 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 (which varies by vehicle model). The only other source of variation in ΔG_{jt} could be from changes in expected future 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.1).

²¹Much 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 (e.g., Allcott and Wozny 2014; Sallee, West, and Fan 2016; Grigolon,

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 Section 5.2.

There are three empirical challenges to interpreting an estimate of γ in Equation (2) as a causal estimate of undervaluation. First, the change in the expected discounted future fuel costs ΔG_{jt} must be constructed based on assumptions about future driving, vehicle survival probabilities, expected future gasoline prices, and the car buyer’s discount rate. We follow the existing literature in using an extensive set of assumptions to better understand the plausible range of γ . Second and relatedly, ΔG_{jt} is potentially subject to measurement error (see Appendix D.1 for details). Our natural experiment helps to overcome some of the measurement issues in ΔG_{jt} because the restatement is perfectly observed. Third, if there is a quantity effect, such that sales (and thus market shares) also respond to the restatement, then γ would not be estimating the willingness-to-pay. We discuss this in more detail in the next subsection and perform a bounding analysis to show the influence of quantity effects on our findings.

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 Jacobsen and van Benthem (2015), and expected gasoline prices being held constant in real terms at the level at time t (a martingale assumption, following evidence from Anderson, Kellogg, and Sallee (2015)). Table 5 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.

The results show that the equilibrium price changes induced by the restatement correspond to substantial undervaluation of fuel economy: the increase in the expected net

Reynaert, and Verboven 2018; Leard, Linn, and Zhou 2018) in presenting a valuation parameter conditional on an assumed discount rate (and the same set of assumptions about other factors). This is an expositional choice.

Table 5: The Valuation of Fuel Economy Based on the Equilibrium Price Change

	(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)		(0.07)		(0.08)	
$1(\Delta Lifetime Fuel Costs)_{jt} \times$ $1(2011 - 2012 Affected Model)_j$		-0.33		-0.39		-0.46		-0.58
		(0.17)		(0.20)		(0.24)		(0.30)
$1(\Delta Lifetime Fuel Costs)_{jt} \times$ $1(2013 Affected Model)_j$		-0.13		-0.16		-0.18		-0.23
		(0.05)		(0.06)		(0.07)		(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

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). 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.

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.²² 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 0.33, 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 0.58 and the 2013 model year has a parameter of 0.23.

Our preferred estimates use a middle ground 4% discount rate. This gives a valuation parameter of 0.16 for model year 2013 and 0.39 for model years 2011-2012. A value of 4% falls in the middle of the range of discount rates assumed in the preferred specifications from other recent studies, which vary from 1.3% to 6% (see Table 7). In our context, using a relatively low discount rate appears reasonable because we study new-vehicle buyers who are likely not capital constrained, have access to cheap car loans, and can likely borrow at low rates in general. The real borrowing rate represents the opportunity costs of the lease or loan payments for those who lease or finance their new-vehicle purchases; for those who pay cash, this rate is the opportunity costs of not being able to invest in other investments with a similar risk-return tradeoff. This rate was quite low during our sample period.²³

We cannot emphasize enough that with different sets of assumptions, the undervaluation parameter would change. For a wide enough range of assumptions, the valuation

²²For the pooled sample, an implicit discount rate of approximately 80% would be required to bring the valuation parameter to one. Put in terms of payback period (the metric used most often by industry), our pooled-sample result implies a payback period of about three years.

²³Leard, Linn, and Zhou (2018) report a real borrowing rate of 1.3% for the period October 2009 to September 2014. Nominal auto loan rates for new vehicles were in the 4.5-5% range during our sample period (https://www.federalreserve.gov/releases/g19/HIST/cc_hist_tc_levels.html); after accounting for CPI increases of 1.5-2.1% for the period 2012-2014, the real auto loan rate was approximately 3%. The federal funds rate in November 2012 was 0.16% (<https://www.macrotrends.net/2015/fed-funds-rate-historical-chart>). Using the Allcott and Wozny (2014) approach, we also calculate the average discount rate for auto loans from the Consumer Expenditure Survey, but updated for our sample period. This yields a discount rate in the range of 4%.

parameter can be as low as zero or as high as one. However, we conduct a fairly exhaustive sensitivity analysis to investigate the robustness of our results in Appendix Table D.1 and conclude that, using reasonable sets of assumptions for constructing ΔG_{jt} that closely follow the existing literature, these assumptions do not change our main result of substantial undervaluation.

5.2 Bounding Analysis

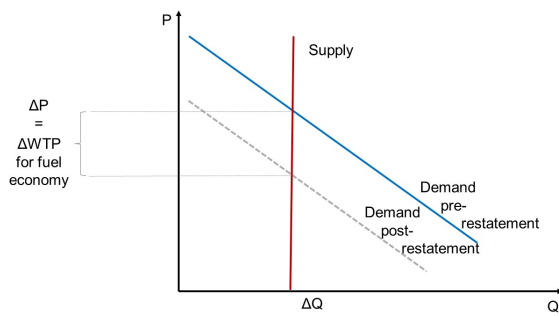
5.2.1 Conceptual Framework

Our valuation analysis so far is based entirely on changes to the equilibrium prices. However, if sales also respond to the restatement, the parameter γ in Equation (2) no longer represents consumers' willingness-to-pay for fuel economy. In this section, we present a simple framework to provide intuition for why the change in willingness-to-pay and the change in equilibrium prices diverge and illustrate how to calculate the willingness-to-pay in such cases.

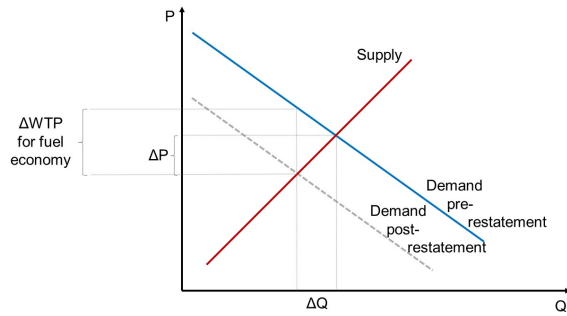
When the supply of vehicles is at least somewhat elastic, such that there are non-negligible quantity effects, the difference between the change in willingness-to-pay and equilibrium prices depends on the slopes of the supply curve, the (residual) demand curve, and the underlying market structure. The panels in Figure 2 illustrate four possible scenarios for how the supply of vehicles could influence the difference between the change in willingness-to-pay and prices. In all four, the restatement shifts demand downward towards the origin and this vertical shift represents the change in willingness-to-pay.²⁴ The first three panels provide the intuition under perfect competition, which is useful to fix ideas and is the common assumption in the literature (e.g., Busse, Knittel, and Zettelmeyer 2013). The fourth panel allows for imperfect competition.

We begin with the case of perfectly inelastic supply (i.e., a zero quantity effect). Panel A shows that under perfectly inelastic supply the change in equilibrium price (our γ) is

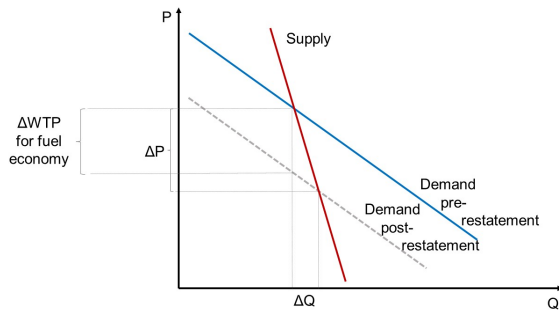
²⁴We assume locally parallel shifts in the demand curve, which is supported by the limited role for consumer selection as discussed in Section 4.1.



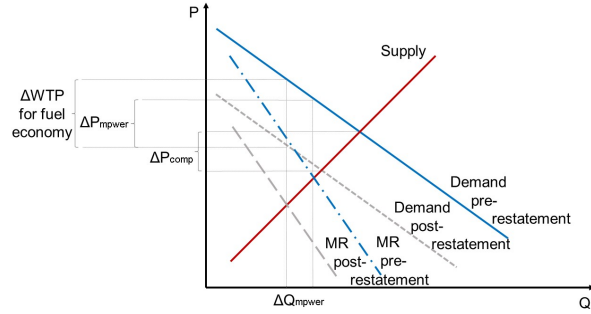
(a) Inelastic Supply



(b) Upward Sloping Supply



(c) Downward Sloping Supply



(d) Market Power and Upward Sloping Supply

Figure 2: Interpretation of the Equilibrium Effect

Notes: The first three panels present a particular scenario with respect to the slope of the supply curve and how it impacts the interpretation of the equilibrium price effect under competitive pricing. The fourth panel assumes upward-sloping supply and compares the change in equilibrium price for the competitive case versus the market power case.

exactly equal to the change in willingness-to-pay for fuel economy. This intuition also holds under imperfect competition, so if we have perfectly inelastic supply, then our results in Table 5 can be interpreted as the willingness-to-pay regardless of the nature of competition in the market.

Next, we assume upward-sloping supply, which would imply a negative quantity effect from the restatement. This is a standard assumption, even if we find no evidence to support it in our data (although we cannot rule it out either). Panel B shows that under upward-sloping supply, the change in equilibrium price underestimates the willingness-to-pay for fuel economy. In the next subsection, we will perform a set of bounding calculations to provide guidance on how one might adjust the estimates in Table 5 based on different assumptions of the slope of supply.

Panel C assumes downward-sloping supply, which is consistent with a positive quantity effect from the restatement. This is less likely, but it is possible that there are local economies of scale in the production of vehicles. The point estimates of our quantity effects also suggest this could be the case. In this scenario, the change in equilibrium price overestimates the willingness-to-pay. This would suggest that our estimates in Table 5 are biased upwards and that the true willingness-to-pay is even closer to zero.

Panel D allows for imperfect competition and for the sake of exposition examines the case of upward-sloping supply, which is the standard assumption. When there is imperfect competition, the marginal revenue lies below the residual demand, allowing firms to earn a markup. Therefore, with upward-sloping supply, the change in price when there is market power will always be greater than the change in price in the competitive market. This means that when we have imperfect competition and upward-sloping supply, the change in equilibrium price will still be an underestimate of the willingness-to-pay, but not as much of an underestimate as it would have been under perfect competition. We will discuss this further in our bounding analysis below.²⁵ In Appendix D.3 we derive the

²⁵The less likely case of downward-sloping supply is more complicated under imperfect competition and the bias from ignoring imperfect competition could go either way depending on the relative slopes of the supply and demand curves.

results discussed in this section more formally with a simple analytical model.

5.2.2 Bounds on the Valuation of Fuel Economy

We can use the theoretical observations about the influence of a quantity effect on our valuation parameter to inform a simple bounding analysis. We begin with the implications of upward- or downward-sloping supply while assuming perfect competition. Recall that in Appendix Table C.1, we found noisy estimates for the effect of the restatement on sales, with a slightly positive point estimate of 0.05 (standard error 0.04). As this cannot rule out either a positive or negative quantity effect, we use a wide range of values for what the quantity effect might be. If we assume standard upward-sloping demand, as in Panel B of Figure 2, then we should see a negative quantity effect. We examine quantity effects down to -5%. For context, since we found a precisely estimated price effect of -1%, a -5% quantity effect would be quite large relative to the price effect. If we assume economies of scale are such a dominant force that they induce a downward-sloping supply curve, as in Panel C, then we should see a positive quantity effect. We examine quantity effects up to +5%.

To estimate willingness-to-pay, we further have to assume a price elasticity of demand.²⁶ Berry, Levinsohn, and Pakes (1995) find vehicle model-level 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 but at a higher level of vehicle-model aggregation. 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. Moreover, our data are highly disaggregated; an observation is even more detailed than make-model-trim-vintage (VIN10), thus affording ample opportunities for consumers to substitute to a similar vehicle, leading to more elastic demand. Accordingly, we first calculate our estimates using a demand elasticity of -6, but we also perform the analysis using a smaller estimate of -4. We also need to assume

²⁶In Appendix D.3, we show how to translate a given change in the equilibrium price into a change in willingness-to-pay using demand and supply elasticities alone under common assumptions.

an average vehicle price pre-restatement, and for this we use \$24,500 (this is calculated as $\$294/0.012$ for consistency with our main results in Table 2; it is also reasonably closely aligned with the summary statistics on vehicle prices for Hyundai and Kia in Table 1).²⁷

Using these assumptions, the adjustment formulas in Appendix D.3, and the \$294 reduction in equilibrium price due to the restatement, Table 6 shows that, for a 5% reduction in quantity, the willingness-to-pay is \$498 when using a demand elasticity of -6 and \$600 when using an elasticity of -4 (under perfect competition). The latter is roughly a doubling of the estimated equilibrium price change. Conceptually, we are just moving along the demand curve by the percentage change in quantity. For smaller quantity effects—e.g., in the -1% range—a \$294 reduction in equilibrium price translates in a willingness-to-pay of \$335 (under the -6 elasticity), which is a much tighter bound. If we assume a +5% quantity effect, then the \$294 reduction in equilibrium price corresponds to a willingness-to-pay of only \$90 when using a demand elasticity of -6 and is even below zero when using a demand elasticity of -4. Overall, these illustrative calculations suggest that the estimated valuation parameters could be either twice as large or close to zero for these particular quantity effects.

Table 6: Interpretation of Equilibrium Change in Prices w.r.t. Different Supply Curves

Quantity Effect (%)	Willingness-to-Pay (\$)	
	$\eta_D = -6$	$\eta_D = -4$
-5	498	600
-1	335	355
0	294	294
1	253	233
5	90	-12

Notes: The table shows how a given equilibrium change in price translates into willingness-to-pay for fuel economy (under perfect competition). 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.

If imperfect competition is at play, but we calculate the willingness-to-pay for fuel

²⁷Note that when we use a lower pre-restatement price, such as \$20,000, the range of results narrows substantially.

economy assuming perfect competition, the results with upward-sloping supply would be biased upwards, since the change in price is not as much of an underestimate of the willingness-to-pay. Thus, the results in Table 6 showing the willingness-to-pay for quantity effects of -5% and -1% should be seen as an upper bounds. These upper bounds indicate that with even a large quantity effect of -5% (which is not justified by our data), the willingness-to-pay should be no more than double the equilibrium price change.

Combined with Table 5, the results in Table 6 demonstrate that our main conclusions about substantial undervaluation hold up to a wide range of quantity effects. For instance, consider the pooled sample and a 12% discount rate in Table 5. Further, suppose that the supply curve is highly elastic such that it translates to a doubling of the valuation parameter from 0.25 to 0.50. For our preferred 4% discount rate, a doubling of the valuation parameter corresponds to an adjusted value of 0.34. For the valuation parameter for the 2011 and 2012 model years, a doubling of the estimate would yield a value of 0.78. Of course, for those model years a highly elastic supply is very unlikely.²⁸ Assuming a supply elasticity closer to zero, the effect on the valuation parameter should be much more modest. In Table 6, a quantity effect of -1% leads to an underestimate of the willingness-to-pay of only 12% using a demand elasticity of -6 (calculated as $(294-334)/334$) or 17% using a demand elasticity of -4. When applied to the model years 2011 and 2012, the valuation of fuel economy falls below 0.5, suggesting substantial undervaluation.

5.3 Comparison to Previous Literature

Table 7 summarizes the range of our results along with several notable papers that perform a similar valuation exercise. The table divides studies into those estimating an exact valuation parameter or an approximate valuation parameter, a distinction we discuss further below. 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,

²⁸As 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 costs of holding vehicles in inventory on the dealer lot.

which implies near-full valuation. 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. Interestingly, our estimates also align with automakers' beliefs about how consumers value fuel economy. For instance, our valuation estimate of 0.39 corresponds to a payback time of a little less than three years, where payback time is defined as the number of years that consumers fully value fuel economy after which they do not value it at all. Automakers report that their planning decisions are based on an assumed consumer payback time of one to four years. This finding is based on years of focus groups with potential car buyers and other market research (National Research Council 2015; McAlinden et al. 2016).²⁹

Table 7: Comparison of Estimates with Other Studies

<i>Studies using exact valuation parameter</i>	<i>r</i>	<i>valuation parameter</i>
Sallee, West, and Fan (2016)	5%	1.01
Allcott and Wozny (2014)	6%	0.76
Own Estimate from Restatement	5%	[0.17-0.42]
Own Estimate from Restatement	6%	[0.18-0.44]
<i>Studies using approximate valuation parameter</i>		
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
Leard, Linn, and Springel (2019)	2.9-5.3%	0.06-0.76
Own Estimate from Restatement	6%	[0.40-1.01]
Own Estimate from Restatement	1.3%	[0.31-0.77]

Notes: For our own estimates, we report a range that highlights the heterogeneity between model years 2011-2012 versus 2013. The lower value of the range represents the valuation parameter for model years 2011-2012. The upper value corresponds to model year 2013.

²⁹This estimate is also consistent with Allcott and Knittel (2019), who find a required payback period of two years or less using stated-preference survey data.

5.4 Possible Explanations for Our Lower Valuation Estimates

There are several possible explanations for why our estimates are lower than most others. Broadly speaking, the explanations fall into three categories: differences in empirical setting, differences in the variation being used, and differences in methodology.

5.4.1 Differences in Empirical Setting

The focus of our analysis is on new cars from Hyundai and Kia during the period 2011 to 2014. Several of the other studies provide estimates from different markets and time frames.

Some of the recent studies estimate the valuation parameter for used car buyers. For example, 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). However, Grigolon, Reynaert, and Verboven (2018) uses data from the European automobile market, which differs from the market in the United States.

Also, our analysis is based primarily on Hyundai and Kia new car buyers, and it is possible that these buyers are different from other new car buyers. On the one hand, it seems likely that Hyundai and Kia, which are known for smaller, more fuel-efficient cars, draw a segment of buyers that are more attentive to fuel economy and 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. Some of the earlier papers use data covering a time period that ends before ours begins. Our data start in 2011 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. It is possible that 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.

5.4.2 Differences in Identifying Variation

One major difference is that our study is the first to use variation from a natural experiment that exogenously changed fuel-economy ratings; most previous studies leverage changes in gasoline prices. This feature of our analysis is very useful, as it assures that other vehicle attributes are held constant, and it leverages exogenous variation in fuel economy, which is exactly the attribute that would change under revised fuel-economy standards.

The variation from the fuel-economy rating restatement that we exploit could affect the interpretation of our results in several ways. First, it is possible that consumers are slow to update prior beliefs about the true fuel economy of the vehicles after the ratings changed. Perhaps consumers base beliefs on information from many sources, such as conversations with other car owners or advertisements they had seen previously, and update over time to eventually reach the true fuel-economy value. In this case, estimates of the valuation parameter would increase over time. However, the vast majority of consumers in our sample bought their vehicles either before or at least several months after the Hyundai and Kia restatement. Should consumers update, it is likely that new car buyers several months after the restatement would have already updated their beliefs about fuel economy.³⁰ Yet in Appendix Table B.5, we find that if we exclude car buyers in the months just after the restatement (for up to one year in length), our estimates are only modestly affected. One would have to believe that updating is extremely slow for this to change the interpretation of our estimates.

Even so, the results in Appendix Table B.5 do not completely rule out extremely slow updating of beliefs about fuel economy. If that were the case, our results would directly apply to the first few years after fuel-economy standards are tightened, as eventually

³⁰Also note that many of these car buyers likely started their search after the restatement and only saw the newer, lower fuel-economy rating on car comparison websites, which did not indicate that the rating had changed.

consumers would correctly update. However, if fuel-economy standards continue to be tightened year-on-year, our results would continue to apply for the further increases in the standards. It is also useful to recognize that if new car buyers are slow to update their beliefs about fuel economy, we might expect the same new car buyers to be slow to update their beliefs about future gasoline prices when current gasoline prices change. If so, then the results from most of the previous studies would also only be useful for understanding medium-run consumer responses.

Second, it is possible that Hyundai or Kia new car buyers were already aware that the affected 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 were not affected by the restatement, including other unaffected Hyundai and Kia models. The reason for these common complaints is that individual driving behavior also influences fuel economy, so there is heterogeneity in the actual on-road fuel economy achieved.³¹ But in general, the EPA fuel-economy rating for each vehicle is carefully designed and monitored to be correct on average and thus is used in all car comparison articles, websites, and apps that we are aware of. The rating is also widely used in the academic literature, including in all of the valuation studies in Table 7, to provide an unbiased estimate of the true fuel economy (e.g., Allcott and Knittel 2019; Jacobsen et al. 2019).³² All things considered, it appears highly unlikely that consumers already knew about the restatement in advance and already adjusted their priors for the fuel economy of the exact models and trims that were affected.

Third, it could be possible that the restatement itself had an impact on consumers'

³¹One particular way in which consumer beliefs about fuel economy may deviate from the EPA ratings is if car buyers believe that the ratings are overstated for all vehicles by a fixed percentage. In that case, consumers should anticipate a larger change in discounted fuel costs than predicted by the change in EPA ratings, which would imply that our—already low—estimate of undervaluation is too high. If instead consumers believe all ratings were off by a fixed amount in gallons-per-mile, our estimate should be unaffected.

³²Indeed, Allcott and Knittel (2019) provide evidence that consumer beliefs about fuel economy are noisy, but unbiased on average.

overall trust and willingness to rely on EPA ratings. This latter explanation would affect our estimates if the restatement induces trust to vary asymmetrically between affected models versus non-affected models within Hyundai and Kia (recall Appendix Table B.3 shows that our results are robust to using only variation from Hyundai and Kia). We view this as unlikely because new car buyers would have to do substantial research to determine which models and trims were affected. When new car buyers go to a car comparison website or the dealership, there would be no indication of the restatement, only the new numbers for the fuel economy of the affected vehicles.

Fourth, as suggested in survey evidence in Allcott (2013), it is possible that, for at least some consumers, when they compare pairs of similar vehicles, they mis-categorize them as having exactly the same fuel economy, but when evaluating across vehicle pairs with very different fuel-economy ratings (e.g., in different vehicle classes), they perceive a difference in fuel economy. Should this be true, it might imply that consumers undervalue fuel economy for small changes in ratings, but come closer to correctly valuing fuel economy for larger changes. Our empirical analysis is based on relatively small changes in fuel economy (from 1 to 6 miles-per-gallon, but with most restatements around 1-3 miles-per-gallon). These relatively small changes in fuel economy are especially useful for understanding fuel-economy standards because they are in the order of magnitude of recent year-over-year increases in standards in the United States. In other words, for policy relevance, we are most interested in the consumer response to such changes.

While our study provides evidence of substantial undervaluation, we cannot differentiate between different possible explanations. For example, consumers may be inattentive to fuel-economy ratings, exhibit a lack of sophistication in correctly processing fuel-economy information, have large cognitive costs when considering future fuel expenditures, display a variety of incorrect beliefs about fuel economy, or update their beliefs slowly when new information arrives (see above). All of these explanations—or some combination of them—could help explain why consumers on average are not willing to pay the full discounted benefits of higher fuel-economy vehicles.

5.4.3 Differences in Methodology

A final 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 discounted 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 (as in the approximation) is usually not the same as the mean of their ratio (the estimate of γ in Equation (2)) 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 Table 7, we divide up the recent studies based on the approach taken. We also provide our own estimates using the same discount rates used in the previous studies and show how the approximation impacts the valuation parameters for model years 2011-2012 versus 2013. To compute the approximated valuation parameter, 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 are below 0.5 when we estimate the exact valuation parameter, 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 near one, as in several previous papers. Thus, the choice between exact versus approximated valuation parameter is consequential in our empirical application—it more than doubles the estimate of the valuation parameter. The approximation bias is large enough to significantly alter the main conclusions from the analysis, potentially leading the researcher to incorrectly conclude that consumers do not undervalue future fuel savings. In addition, our 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, although the magnitude of the bias could differ across studies.

6 Conclusions

This paper exploits an unexpected restatement in the EPA-rated fuel economy for over a million 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 source of exogenous variation in future fuel costs expected by consumers. The 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.

In our preferred set of estimates, we find that consumers are indifferent between one dollar in discounted future gasoline costs and 16-39 cents in the vehicle purchase price, where the higher estimate is when we restrict the sample to 2011-2012 model year vehicles. This result suggests that consumers undervalue future fuel savings when they purchase new vehicles. We further perform an extensive sensitivity analysis to show that even under a wide range of assumptions about factors such as consumer expectations, discounting, and expected future driving, we continue to find undervaluation. We also perform a bounding analysis using different assumptions about supply elasticities, demand elasticities, and market structure to illustrate that for very broad ranges of assumptions, we continue to find substantial undervaluation.

Such undervaluation of fuel economy could come about from a mix of behavioral factors, such as (rational) inattention, lack of salience of fuel economy, or present bias in the vehicle purchasing decision. We cannot disentangle these factors, but from a policy perspective, it is crucial to know if and to what extent consumers are undervaluing fuel economy. Our finding is consistent with long-standing beliefs in the automobile industry, but differs from some—but not all—of the recent literature. Our analysis highlights that our results differ less after accounting for whether the study estimates the exact valuation

parameter or an approximation. But 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. We suspect that similar policy considerations carry over to other settings. For example, the presence of behavioral biases in valuing important attributes might apply more generally to contexts with products that have back-loaded costs or benefits, such as solar panels, energy efficiency upgrades, health care plans, and retirement savings, among many others.

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. 2013. "The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market." *American Economic Journal: Economic Policy* 5 (3):30–66.
- 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.
- Bento, Antonio M., Kenneth Gillingham, Mark R. Jacobsen, Christopher R. Knittel, Benjamin Leard, Joshua Linn, Virginia McConnell, David Rapson, James M. Sallee,

- Arthur A. van Benthem, and Kate S. Whiitefoot. 2018. "Flawed Analyses of U.S. Auto Fuel Economy Standards." *Science* 362 (6419):1119–1121.
- 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.
- 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 Metcalf. 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.
- McAlinden, Sean P., Yen Chen, Michael Schultz, and David J. Andrea. 2016. *The Potential Effects of the 2017-2025 EPA/NHTSA GHG/Fuel Economy Mandates on the U.S. Economy*. Center for Automotive Research: Ann Arbor, MI.
- National Research Council. 2015. *Cost, Effectiveness, and Deployment of Fuel Economy Technologies for Light-Duty Vehicles*. The National Academies Press: Washington, DC.
- 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.

APPENDIX

A Fuel-Economy Label and Affected Vehicles

This appendix provides further details on the restatement and 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 (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).³³ 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 in the region where they live.

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).³⁴ 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 courts as of January 2019.³⁵

³³From <https://www.autoblog.com/2012/11/02/hyundai-kia-admit-exaggerated-mileage-claims-will-compensate-o/>.

³⁴From <https://www.consumerwatchdog.org/courtroom/us-court-appeals-rejects-hyundaikia-settlement-fuel-economy-scandal>.

³⁵Hyundai 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 compensation to previous buyers. See <https://www.epa.gov/enforcement/hyundai-and-kia-clean-air-act-settlement>.

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.³⁶ 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.³⁷ 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.

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

³⁶Source: <https://www.autoblog.com/2012/11/02/hyundai-kia-admit-exaggerated-mileage-claims-will-compensate-o/>.

³⁷The 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) Model	(2) Model Year	(3) Trim	(4) Engine	(5) Drive	(6) Tran.	(7) Original Rating			(8) Restated Rating		
						(7) City MPG	(8) Hwy MPG	(9) Comb. MPG	(10) City MPG	(11) Hwy MPG	(12) 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) Model	(2) Model Year	(3) Trim	(4) Engine	(5) Drive	(6) Tran.	(7) Original Rating			(8) Restated Rating		
						(9) City MPG	(10) Hwy MPG	(11) Comb. MPG	(12) City MPG	(13) Hwy MPG	(14) 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.

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.

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.³⁸ 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.

³⁸The 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.



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

B Robustness Checks

This section provides a series of results to explore the robustness and heterogeneity in our primary findings. 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 different sets of fixed effects for month-of-sample interacted with vehicle class. Specifically, we change the definition of a vehicle class to be finer 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 how we control for relative time trends in the price of affected and non-affected vehicles across segments. We find that our results are highly robust to these alternative specifications.

Table B.2 provides further 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.

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, or 0.02% of observations). 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 dif-

Table B.1: 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 Coarser Class FE		Y			Y	
Year-Month \times Finer 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. Columns 2 and 5 use a coarse definition of vehicle classes where we only distinguish: small car, large car, minivan, crossover, SUV, and pickup. Columns 3 and 6 use a finer definition of vehicle classes, relative to the main specifications, where luxury and non-luxury vehicles are distinguished. The definition of vehicle classes in those specifications closely follows R.L. Polk nomenclature. 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 B.2: 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	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 clustered by VIN10.

ferences, but by-and-large, we find that our results are robust across these specifications.

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 clustered by VIN10.

We also run all of the primary specifications after excluding 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). One might be concerned that these skew our results. Table B.4 excludes these minimally treated models from the sample. Again, the results are remarkably similar.

Table B.5 shows how the point estimates change when we omit up to 12 months following the restatement. The reason we might want to do this robustness check is two-fold. First, one might be worried about an unusual selection of car buyers for the affected models just after the restatement. Presumably this would dissipate for new car purchases many months later. Second, it is possible that new car buyers just after the restatement do not look at the updated ratings and base their decision on the earlier ratings they had seen prior to the restatement when they were new car comparison shopping in preparation for the purchase. But it is unlikely that new car buyers will still be basing decisions

Table B.4: Robustness Check Excluding Minimally Treated Observations

	(1)	(2)	(3)	(4)	(5)	(6)
	Logs			Levels		
$1(\text{Post Restatement})_t \times 1(\text{Affected Model})_j$	-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
$1(\text{Post Restatement}) \times \text{DMA FE}$	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 clustered by VIN10.

on information acquired many months before. Fortunately, the estimates are quite stable: the specifications with omitted months do not differ much from the baseline specification in column 1.

Table B.5: Effect on Transaction Prices Excluding the Months Closest to the Restatement

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Number of Post-Months Excluded											
	0 (base)	1	2	3	4	5	6	7	8	9	10	11
$1(\text{Post Restatement})_t \times 1(\text{Affected Model})_j$ (Logs)	-0.012 (0.003)	-0.012 (0.004)	-0.013 (0.004)	-0.013 (0.004)	-0.015 (0.005)	-0.016 (0.005)	-0.016 (0.006)	-0.013 (0.006)	-0.012 (0.006)	-0.012 (0.007)	-0.011 (0.008)	-0.010 (0.009)
$1(\text{Post Restatement})_t \times 1(\text{Affected Model})_j$ (Levels)	-294 (91)	-310 (98)	-324 (106)	-341 (115)	-389 (125)	-408 (136)	-415 (147)	-368 (158)	-347 (171)	-364 (185)	-330 (204)	-320 (229)
N	1.52m	1.47m	1.43m	1.38m	1.34m	1.29m	1.25m	1.20m	1.16m	1.11m	1.07m	1.02m
Year-Month \times Class FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Year-Month \times Make FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VIN10 FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
DMA FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
$1(\text{Post Restatement}) \times \text{DMA FE}$	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Notes: Each row and column represents the results from a different regression, for twenty-four total. For all regressions the dependent variable is either the 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. The R-squared for all log and level regressions equals 0.95-0.96. Standard errors clustered by VIN10.

One may also be interested in automaker heterogeneity. Was Hyundai or Kia affected more? Table B.6 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.

Table B.6: Heterogeneous Effects on Transaction Prices by Automaker

	Primary		Automaker	
	(1) Logs	(2) Levels	(3) Logs	(4) Levels
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	-0.012 (0.004)	-294 (91)		
$1(Post\ Restatement)_t \times 1(Hyundai\ Affected\ Model)_j$			-0.014 (0.005)	-365 (123)
$1(Post\ Restatement)_t \times 1(Kia\ Affected\ Model)_j$			-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 clustered by VIN10.

In Table B.7, 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.7: Heterogeneous Effects on Transaction Prices by Vehicle Class

	(1)	(2)
	Logs	Levels
$1(\text{Post Restatement})_t \times 1(\text{Small Car Affected Model})_j$	-0.013 (0.005)	-320 (123)
$1(\text{Post Restatement})_t \times 1(\text{Large Car Affected Model})_j$	-0.025 (0.004)	-702 (134)
$1(\text{Post Restatement})_t \times 1(\text{Crossover Affected Model})_j$	-0.007 (0.004)	-190 (98)
$1(\text{Post Restatement})_t \times 1(\text{Sport Affected Model})_j$	-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
$1(\text{Post Restatement}) \times \text{DMA FE}$	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 clustered by VIN10.

C Effect of Restatement on Other Outcomes

C.1 Effect on Quantities

In this appendix, we 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, leading to large month-on-month relative changes in sales.

Table C.1 confirms our intuition that automobile sales are very noisy. In Panel A, 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 2-5, 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. The full-sample result in column 1 is especially noisily estimated. We therefore need 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. Accordingly, we focus on the specifications in columns 2 through 5, which 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 25% of monthly sales, in column 3 it is 30%, in column 4 it is 40%, and in column 5 it is 50%. Panels B and C show the estimates from regressions with fewer fixed effects, analogous to columns 1 and 2 in Table 2. The point estimates are very similar.

The coefficients in Table C.1 are all *positive*, suggesting that the restatement *increased* sales, which may appear to be a counter-intuitive result. However, they are all imprecisely estimated. The precision of the estimates improves somewhat as we apply a stricter exclusion criterion; we favor column 5 for that reason. 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. As the estimates in Table C.1 do not allow us to rule out either sizable positive or negative quantity effects, we discuss the implications of negative or positive quantity effects on estimates of the valuation of fuel economy in Section 5.2. 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)	(2)	(3)	(4)	(5)
	Incl. Outliers	<25%	<30%	<40%	<50%
Panel A: Main Specifications					
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	0.15 (0.08)	0.06 (0.05)	0.05 (0.05)	0.04 (0.05)	0.05 (0.04)
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
Panel B: Without Year-Month \times Class FE					
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	0.19 (0.08)	0.07 (0.05)	0.06 (0.05)	0.05 (0.04)	0.06 (0.04)
Panel C: Without DMA FE and $1(Post\ Restatement) \times DMA$ FE					
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	0.11 (0.06)	0.06 (0.04)	0.05 (0.04)	0.04 (0.04)	0.05 (0.04)
R-squared	0.22-0.46	0.23-0.51	0.24-0.53	0.24-0.53	0.24-0.53
N	4.01m	3.75m	3.70m	3.62m	3.53m

Notes: Dependent variable is log of sales. Panels A-C correspond to the identical fixed-effect structure as in Table 2. Columns 2-5 present the results eliminating outliers by 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. The R-squared row shows the range across the three panels. Standard errors clustered by VIN10.

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.

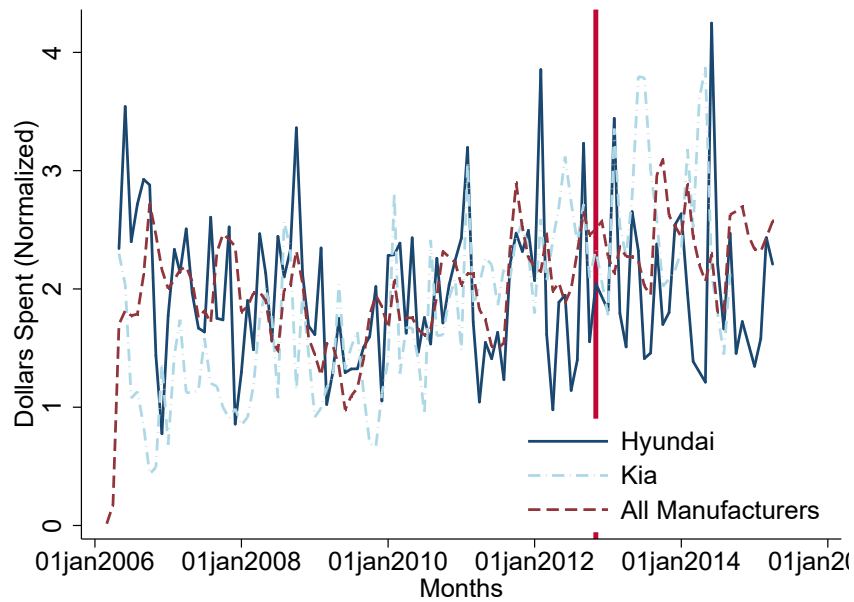


Figure C.1: Spending on Advertising by Different Automakers

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

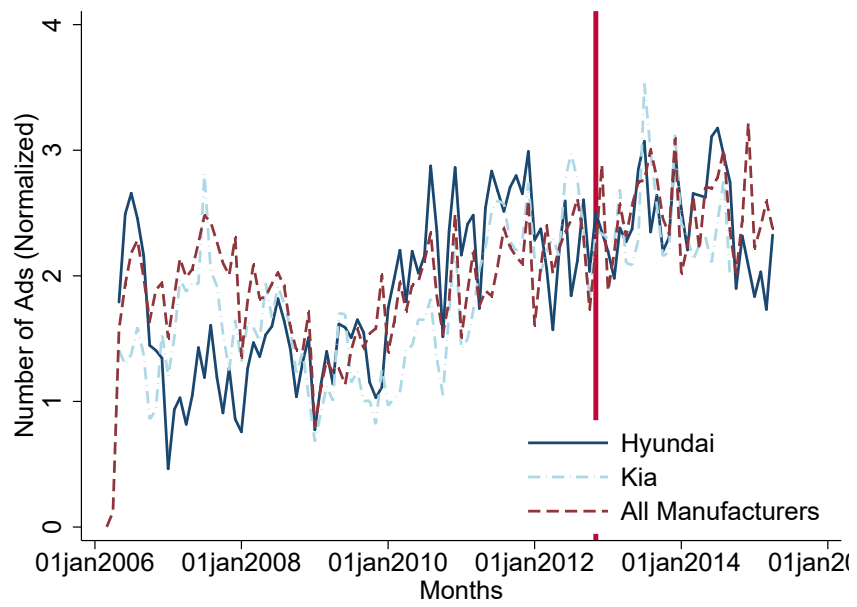


Figure C.2: The Number of Advertisements by Different Automakers

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 the 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 tradeoff 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. This is true when supply is completely inelastic and fixed. If supply cannot change, then the market shares are exogenous. (Note that, if supply responses do occur, we need to adjust the estimate of our valuation parameter. We do this in Section 5.2 and conclude that our finding of undervaluation of fuel economy is robust a wide range of quantity adjustments.)

Another 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 ΔG is not a major concern in our setting. This is in contrast to Allcott and Wozny (2014) and Sallee, West, and Fan (2016), 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, a function of the gasoline prices consumers paid, driving behavior, and other assumptions required to construct G . We show, however, that our estimates of the valuation parameter are robust to these assumptions (Table D.1).

D.2 Sensitivity Analysis of the Valuation Parameter

To estimate the valuation parameter, we need to construct the discounted change in future fuel costs of each vehicle 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. Table D.1 outlines 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 purchase. 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 the 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 years 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 base 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). We also estimate the result using vehicle survival probabilities specific to Hyundai and Kia (which are somewhat higher than for most other brands), using data provided by Jacobsen and van Benthem. Data for vehicle miles traveled come from the NHTS. We compare results using the 2006 and the 2017 wave of the NHTS.

Table D.1: Sensitivity Analysis: Valuation Parameters

Discount Rate	Gasoline Prices	VMT	Survival Probability	Ratio of Means	Valuation Parameter	Valuation Parameter: 2012 Model Year Only
4%	Year-US	NHTS 17	JvB	No	0.173	0.395
1%	Year-US	NHTS 17	JvB	No	0.144	0.329
12%	Year-US	NHTS 17	JvB	No	0.255	0.582
4%	2012-US	NHTS 17	JvB	No	0.169	0.389
4%	2012-2014-US	NHTS 17	JvB	No	0.174	0.402
4%	Month-US	NHTS 17	JvB	No	0.171	0.407
4%	Year-State	NHTS 17	JvB	No	0.201	0.384
4%	Year-US	NHTS 06	JvB	No	0.148	0.337
4%	Year-US	NHTS 17	BKZ	No	0.181	0.412
4%	Year-US	NHTS 17	Hyundai/Kia	No	0.198	0.453
4%	All	NHTS 06/17	BKZ/JvB	No	[0.169-0.201]	[0.384-0.417]
4%	Year-US	NHTS 17	BKZ	Yes	0.438	0.908

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. "Year" refers to annual data. "US" refers to national-level data. "State" refers to state-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 2017. For the survival probabilities, we use estimates provided by Jacobsen and van Benthem (2015) (JvB). We also consider the NHTS data as reported by Busse, Knittel, and Zettelmeyer (2013) (BKZ). For the scenario labelled "Hyundai/Kia," we use the survival probabilities specific to Hyundai and Kia calculated from JvB's data. 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.3 Imperfect Competition

Our calculations of the valuation parameter are based on the implicit assumption that the equilibrium prices were set in a competitive market. This assumption was also used in most of the other recent studies in Table 7, including Busse, Knittel, and Zettelmeyer (2013). However, the automobile market is traditionally modeled by economists as a market with differentiated products, where automakers can exercise some market power (Berry, Levinsohn, and Pakes 1995). In Panel D of Figure 2, we present the case with market power and upward-sloping supply to provide intuition for how market power may affect our valuation parameter estimate.

In this section, we present a stylized analytical model to provide further intuition for how market power influences the calculation of the willingness-to-pay. For illustrative purposes, we focus on comparing the two extreme cases of perfect competition and monopoly. This provides easily accessible intuition for the broader case of market power in a market with multiple firms; as such a market with imperfect competition would fall in between the two extremes.

Our main finding is that the equilibrium price effect under monopoly is (weakly) greater than the equilibrium price effect under perfect competition. Further, if we have elastic and upward-sloping supply (as in Panel D of Table 7), then the gap between the willingness-to-pay and the equilibrium price change is smaller when the market is a monopoly than under perfect competition. In other words, market power implies that an upward-sloping supply curve would affect our valuation calculations less.

D.3.1 Preliminaries

Consider the case of (locally) linear demand. This is a reasonable assumption given that we find relatively small price changes. To keep the exposition simple, we also focus on the single-product case where demand is given by:

$$P(Q) = \alpha^0 - \frac{Q}{\delta},$$

where α^0 and $\delta > 0$. We model the effect of the restatement as a reduction in the overall willingness-to-pay for the product by all consumers in the market.³⁹ Formally, this implies a downward parallel shift in demand:

$$P'(Q) = \alpha^1 - \frac{Q}{\delta},$$

where $0 < \alpha^1 < \alpha^0$, and the change in willingness-to-pay for fuel economy equals $\Delta WTP = \alpha^1 - \alpha^0$.

Finally, we assume that supply is elastic and upward-sloping:

$$MC(Q) = \beta + \frac{Q}{\sigma},$$

where $\beta^0 > 0$ and $\sigma > 0$.

D.3.2 Competitive Pricing

In the competitive case, the equilibrium price, before or after the restatement, is determined by the intersection of demand and supply: $P(Q) = MC(Q)$. Solving for quantities before and after the restatement, the change in equilibrium price ($\Delta P = P' - P$) is:

$$\Delta P = (\alpha^1 - \alpha^0) \frac{\delta}{\sigma + \delta} = \Delta WTP \frac{\delta}{\sigma + \delta}$$

It can also be useful to re-express this expression in terms of demand and supply elasticities. Using the linear case as an approximation of the demand and supply relationships, the demand elasticity is given by $\epsilon^D = -\delta \cdot P/Q$ and the supply elasticity is given by: $\epsilon^S = \sigma \cdot P/Q$. Replacing these two expressions in the expression above, we have:

$$\Delta P = \frac{-\epsilon^D}{\epsilon^S - \epsilon^D} \Delta WTP. \tag{D.2}$$

³⁹For larger shifts, it is possible that there is a rotation of the demand curve, but the parallel shift assumption is reasonable as a local approximation. It is also supported by our robustness checks suggesting that there is little evidence to support a compositional effect.

This expression formalizes the intuition in Figure 2. When supply is upward-sloping and elastic, i.e., $\epsilon^S > 0$, and demand is downward-sloping $\epsilon^D < 0$, the change in equilibrium price will always underestimate the change in willingness-to-pay given that $\frac{\epsilon^S - \epsilon^D}{-\epsilon^D} > 1$. When supply is perfectly inelastic, i.e., $\epsilon^S = 0$, the change in equilibrium price is exactly the change in willingness-to-pay: $\Delta P = \Delta WTP$.

D.3.3 Monopoly

In the case of a monopolist, the equilibrium price is determined by the intersection of the marginal revenue and the marginal cost: $MR(Q) = MC(Q)$. When we again assume (locally) linear demand, the marginal revenue curve is given by:

$$MR(Q) = \alpha - \frac{2Q}{\delta}.$$

Solving for quantities before and after the restatement, the change in equilibrium price is now:

$$\Delta P = \Delta WTP \frac{\delta + \sigma}{2\sigma + \delta},$$

which can be expressed as:

$$\Delta P = \frac{\epsilon^S - \epsilon^D}{2\epsilon^S - \epsilon^D} \Delta WTP. \quad (D.3)$$

Comparing the change in equilibrium price under both market structures, the following inequalities can easily be verified: $\frac{-\epsilon^D}{\epsilon^S - \epsilon^D} < \frac{\epsilon^S - \epsilon^D}{2\epsilon^S - \epsilon^D} < 1$ if $\epsilon^S > 0$ and $\epsilon^D < 0$. This implies that under imperfect competition, the change in equilibrium price implied by the restatement will always be larger than under the competitive case, but still it will be less than the full change in willingness-to-pay if supply is elastic ($\epsilon^S > 0$). Put simply, under imperfect competition, the firm has a greater ability to adjust prices to capture the consumer surplus associated with the valuation of fuel economy relative to a perfect competitive setting. Nonetheless, the firm cannot fully capture the surplus, unless supply is completely inelastic.

D.3.4 Bias from Ignoring Imperfect Competition

If we are in a setting with imperfect competition, but we calculate the change in willingness-to-pay assuming perfect competition (i.e., using the demand and supply parameters), the estimate for ΔWTP is given by re-arranging Equation (D.2):

$$\Delta WTP^{biased} = \frac{\epsilon^S - \epsilon^D}{-\epsilon^D} \Delta P, \quad (D.4)$$

where, with a slight abuse of notation, the superscript *biased* is a mnemonic that indicates that the change in willingness-to-pay is calculated using the wrong assumption about the underlying market structure.

If imperfect competition is at play and we are actually in a monopoly setting, then the true (i.e., unbiased) change in willingness-to-pay, which we denote ΔWTP^* , should be calculated using Equation (D.3):

$$\Delta WTP^* = \frac{2\epsilon^S - \epsilon^D}{\epsilon^S - \epsilon^D} \Delta P. \quad (D.5)$$

The bias from ignoring imperfect competition is simply the difference between the two expressions:

$$Bias = \Delta WTP^* - \Delta WTP^{biased} = \frac{(\epsilon^S)^2}{\epsilon^D(\epsilon^S - \epsilon^D)} \Delta P. \quad (D.6)$$

The bias is thus proportional to the size of the change in price and a scaling term that is always less than zero for $\epsilon^S > 0$ and $\epsilon^D < 0$. In our context, given that $\Delta P < 0$, the bias would be positive. This means that by ignoring imperfect competition, we are *overestimating the reduction in willingness-to-pay induced by the restatement*. Note that in this single-product case, the valuation parameter is simply the ratio of the change in willingness-to-pay over the change in expected fuel costs. An upward bias (in absolute value) in calculating ΔWTP thus biases the valuation parameter toward one (and more generally upward). Therefore, this simple illustration shows how ignoring imperfect competition can lead to an overestimate of the willingness-to-pay for fuel economy.

The basic logic here generalizes to other forms of imperfect competition besides monopoly. The key point for our setting is that under imperfect competition, upward-sloping supply is less influential in biasing our valuation parameter based on inelastic demand than in the perfect competition setting.

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

With the setup based on a discrete choice model presented in Section D.1 above, it is 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 simple 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: $\$2,550/\$3,000 = 0.85$. What we see 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 Equation (D.1) is to consistently estimate the true γ . Consider a case where there is heterogeneity over vehicles in γ , so that we can write the parameter as γ_j .⁴⁰ Our simple example above was one case where there was heterogeneity in γ across vehicles, and we showed in Table B.7 that there is heterogeneity across car classes, so we know that empirically there is indeed heterogeneity in γ across vehicles in our context. 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:

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

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

⁴⁰Houde and Myers (2019) consider heterogeneity over consumers, so it is γ_i being considered, but the same logic follows as here.

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.⁴¹

⁴¹It is an upward bias 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).