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

October 15, 2020

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, 327 Vance Hall, 3733 Spruce Street, Philadelphia, PA 19104, United States, phone: +1-215-898-3013, e-mail: arthurv@wharton.upenn.edu. The authors are grateful to many conference and 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, NHH Bergen, RWI Essen, Simon Fraser University, Stanford University, University of California Davis, University of Connecticut, University of Copenhagen, University of Maryland, University of Pennsylvania, Uppsala University, USC, VU Amsterdam, and Yale University.

The transportation sector is now the largest contributor of carbon dioxide emissions in the United States and emissions from petroleum constituted 46% of all energy-related carbon dioxide emissions in 2019 (U.S. Energy Information Administration 2020). 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).[1] 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).[2]

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 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-

[1]We follow a common terminology in the existing literature (e.g., Hausman 1979; Busse, Knittel and Zettelmeyer 2013*a*) 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.

[2]Our 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 (Jessoe and Rapson 2014), and educational investment (Jensen 2010).

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 (Kia Motors America, Inc. 2020). 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.[3] Using detailed microdata on all new vehicle transactions in the United States over the period August 2011 to June 2014 and exploiting variation across affected and unaffected vehicles produced by Hyundai and Kia, we find a 1.2% decline in the equilibrium prices of the affected models (just under $300). We do not find any evidence of diminished overall brand perception for Hyundai and Kia vehicles around the restatement. The change in equilibrium price demonstrates that the rated fuel economy of vehicles is valued by market participants. We then proceed by putting these results into context by 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. We also show that the undervaluation is unlikely to be explained by strategic price spillovers to non-affected models, consumer selection, slow updating of beliefs, or reduced trust and willingness to rely on EPA ratings following

[3] In 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).

the restatement—the undervaluation persists even when only including car buyers who started their search months after the restatement and were likely unaware that the ratings had ever changed.

Previous studies estimating the consumer valuation of fuel economy use several different identification strategies, but most leverage changes in gasoline prices to test whether vehicle prices fully adjust with the changes in the expected discounted present value of future fuel costs. This basic approach was used as early as the 1980s, with Kahn (1986) finding that used car prices adjust only one third to one half the amount that would be expected based on the changes 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. 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. Allcott and Wozny (2014) exploit variation in gasoline prices and estimate a central-case valuation parameter of 0.76 for used vehicle purchasers in the United States. These results suggest 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*a*) 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. Taken together, these studies suggest modest undervaluation at most.[4] 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 in two main ways. First, we demonstrate that the

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

fuel-economy rating itself is indeed valued in equilibrium using variation an exogenous and sudden shifter of the official fuel economy rating, in a context that is appealing because the vehicles themselves are identical before and after the change. Previous studies have used fuel economy ratings to construct a measure of fuel operating costs, but could not test if the market participants respond to the rating itself. 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. Second, we are the first to quantify the valuation of fuel economy using a natural experiment that provides policy-relevant variation in expected future fuel costs through changes in the rating itself, rather than changes in gasoline prices.

Our estimates are especially relevant for informing the intense debate on whether fuel-economy standards are justified from a private perspective.[5] 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

[5] In 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 Davenport (2018)

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

The remainder of this paper is organized as follows. We next describe the natural experiment. In Section 2, we discuss the data. Section 3 presents the empirical strategy and main results that show how the market responded to the information shock provided by the restatement. In Section 4, we estimate consumers' valuation of future fuel costs and discuss and interpret our estimates. The final section concludes.

1 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.[6]

On November 2, 2012, the EPA issued a press release stating that "in processing test

[6]See 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.

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

Immediately after the EPA press release, the fuel-economy ratings for all affected 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.[9] 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.

[7] The incident was covered by the press, e.g., see New York Times (2012).

[8] See Autoblog.com (2014).

[9] Consider this quote from a November 2, 2012 article (Autoblog.com 2012): "Hyundai aggressively advertised the fact that the brand offers four models that boast 40 mpg, but that claim is no longer true."

2 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 (R.L. Polk 2011-2014). 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.[10] 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 VIN10 uniquely identifies the vehicle trim, engine size, and further characteristics. The transaction price is the final price reported to the Department of Motor Vehicles of each state. This final price is net of all dealer-to-customer incentives in all cases we could verify, so if the vehicle costs $40,000 and the dealer offers a $2,000 markdown, the price reported is $38,000. Manufacturer-to-dealer incentives will be at least partly or entirely passed through to the final consumer price, so the DMV transaction price data should reflect these incentives.[11] A limitation of the DMV price data is that it may not include manufacturer-to-customer incentives, and if Hyundai and Kia responded to the restatement by offering enhanced rebates for the affected vehicles directly to customers, we could understate the price changes. Though we were unable to obtain comprehensive data on manufacturer-to-customer incentives, we conducted an extensive review of sources and found no evidence that incentives change significantly around the time of the restatement. See Appendix A for further details.

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

[10]The data include all vehicle transactions, including leases. For leased vehicles, the leasing company buys the vehicle and the transaction price is recorded.

[11]One caveat is that if manufacturer-to-dealer incentives were altered after the restatement, this could potentially shift dealer effort towards different vehicles. In our search of the automobile trade press, we did not find any evidence this occurred in our setting.

slightly below the non-affected models. Both automakers specialize in smaller cars that are priced below the average for other automakers.

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. MSRP, fuel economy, displacement and weight from DataOne (2011-2014). All dollars are nominal dollars.

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 nationwide gasoline prices from the U.S. Energy Information Administration (EIA) (U.S. Energy Information Administration 2011-2014), on vehicle survival rates from R.L. Polk

(1993-2009), and on average vehicle miles traveled from the National Highway Traffic Safety Administration (NHTSA) (National Highway Traffic Safety Administration 2018). In sensitivity analysis, we also provide estimates for miles driven and survival rates from Busse, Knittel and Zettelmeyer (2013*b*) as well as EIA's gasoline prices at the monthly-national level and at the year-state level (U.S. Energy Information Administration 2011-2014).

3 The Equilibrium Effects of the Restatement

In our empirical investigation, we first proceed by using a reduced-form estimator to show how our natural experiment provides internally valid and robust estimates of the impact of the restatement on various outcomes. This empirical strategy does not require making assumptions on how consumers perceive future fuel operating costs.

3.1 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:

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

where $Price$ is either the log or level of the transaction price for a VIN10 j sold in region r (DMA) in year-month t. $1(Post\ Restatement)_t$ is an indicator variable for after the restatement in November 2012 and $1(Affected\ Model)_j$ is an indicator variable for an affected model. 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. These allow us to focus on variation across affected and unaffected vehicles produced by Hyundai and Kia (after controlling for nonparametric automaker-specific time trends to capture any time-varying changes, such as to reputation). We include DMA indicators (η_r) and their interaction with the post-restatement indicator ($\eta_r \times 1(Post\ Restatement)_t$) 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.[12] We weight the regressions by monthly sales[13] and cluster standard errors at the VIN10 level.[14] 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.

Our identifying variation thus comes from within-model and within-region price changes across affected and unaffected vehicles produced by Hyundai and Kia, conditional on flexible time price trends for each vehicle make and class. The source of the variation in the covariate we care about is the restatement itself, which leads some vehicles to be affected and others unaffected in a plausibly random way. 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 $\mu_t \times Make_j$). 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).

[12]Our identification follows recent studies that use disaggregated panel data. For example, Allcott and Wozny (2014) and Busse, Knittel and Zettelmeyer (2013*a*) 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.

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

[14]Clustering 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	-0.010	-0.012	-150	-259	-294
	(0.004)	(0.004)	(0.003)	(80)	(94)	(91)
Year-Month × Class FE		Y	Y		Y	Y
Year-Month × 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) × 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.

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 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 only few treated vehicles are left and the treatment effect reverts 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.

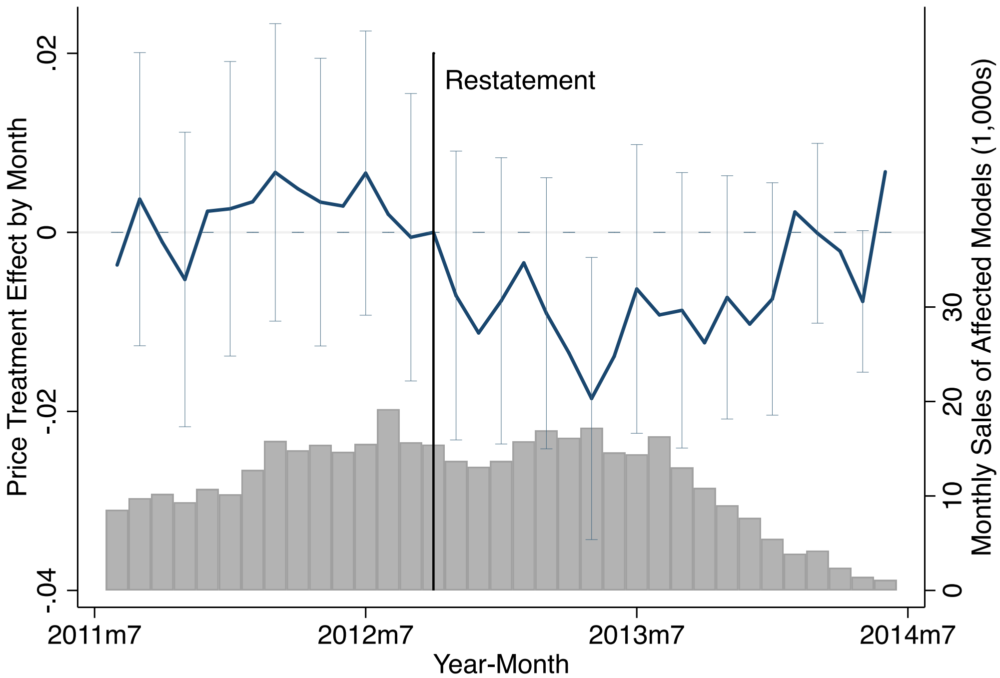

Figure 1: The Price Effect of the Restatement on Affected Models by Month Along with the Monthly Sales of Affected Models

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.

Finally, we explore whether the restatement had an appreciable effect on the unaffected Hyundai and Kia vehicles. One might hypothesize that negative press relating to the restatement affected all vehicles by the two automakers, including those unaffected by the restatement. To explore this, we estimate (1) but replace the year-month by make indicators $\mu_{t \times Make_j}$ with $1(Hyundai\ or\ Kia)_j$ interacted with year-month indicators to be able to plot a single effect for the two affected brands. Figure 2 plots this effect on the unaffected Hyundai and Kia vehicles over time. We observe a modest upward trend prior to the restatement that continues after the restatement through the end of 2013. (The trend reverses in 2014 but recall that very few affected vehicles were remaining at that time.) Just after the restatement, it appears that this trend is slightly amplified, suggesting that the prices of unaffected Hyundai and Kia vehicles increased somewhat more quickly after the restatement than would have been expected by the trend—i.e., a temporary, small but positive price effect. One possible explanation is Hyundai and Kia are multi-product firms with competing models, so profit maximization could have led to higher prices for the unaffected vehicles.[15] Importantly, there is no evidence of a negative brand effect.

3.1.1 Robustness Checks

Any effects that occur at the overall brand level are controlled for in our primary specification and in any event, such effects do not appear to be important in our setting (Figure 2). However, identification could still be compromised in other ways. 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

[15]As we will discuss later, this could bias our valuation parameter upwards, but our robustness checks suggest such a bias is unlikely.

[16]The 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.

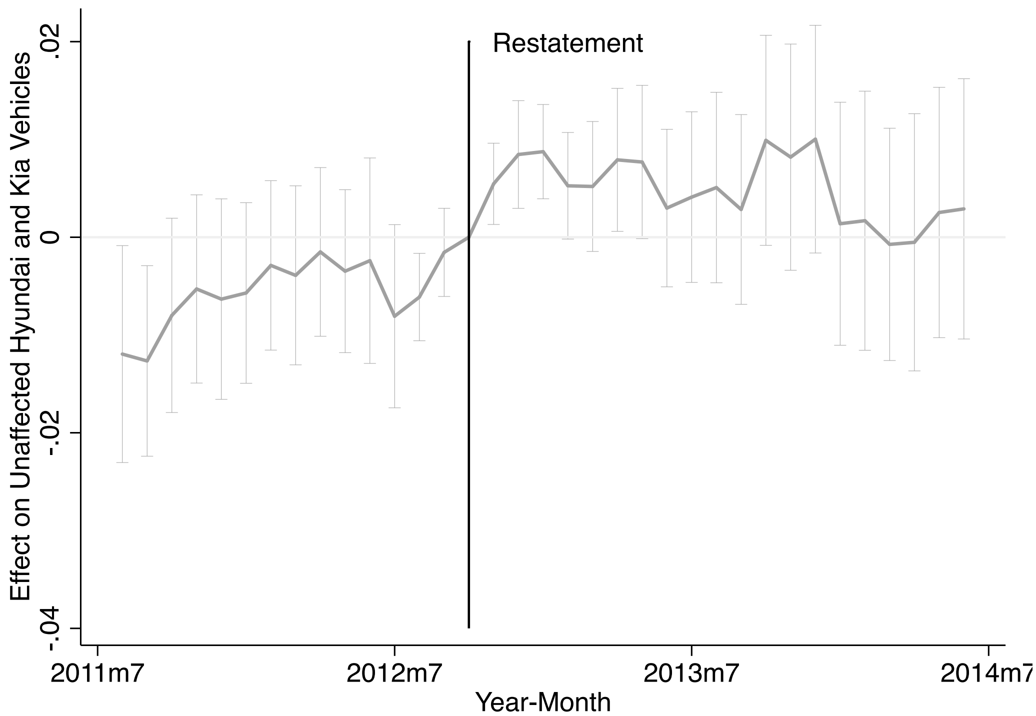

Figure 2: The Price Effect of the Restatement on Unaffected Hyundai and Kia Vehicles

Notes: The black vertical line indicates the fuel-economy restatement date. The standard error for every other month is shown by the bars and whiskers.

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.[17]

We thus 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

[17]In theory, there could also be a secondary response by Hyundai and Kia to the increased prices of close substitutes, which could perhaps counter the overestimate from the initial response.

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 data from Edmunds.com and MotorTrend.com.[18] 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 close 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.

Table 3: Robustness Checks for SUTVA Assumption

	(1)	(2)	(3)	(4)	(5)	(6)
		Logs			**Levels**	
$1(Post\ Restatement)_t \times 1(Affected\ Model)_j$	-0.011	-0.014	-0.013	-261	-365	-342
	(0.004)	(0.003)	(0.003)	(94)	(83)	(84)
Year-Month × Class FE	Y	Y	Y	Y	Y	Y
Year-Month × 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) × 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.

In Appendix B, we explore alternative sets of fixed effects and find that the results are

[18] Edmunds.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.

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. Appendix Table B.2 also adds quarter-of-age $\times$ make fixed effects to capture the cyclicality 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, we deal with several potential concerns and shed light on the interpretation of the estimates. A key issue is if car buyers were aware of the restatement. This likely changes over time. Those who bought a car soon after the restatement may have been aware of the actual restatement and may even have seen the old ratings. As time passes, it becomes increasingly unlikely that car shoppers know about the restatement; most people started their search after the restatement had happened, just saw the new fuel-economy ratings and never knew they had been changed.

It is important for several reasons to establish if our results are driven by buyers who were likely aware of the restatement. 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, complicating the interpretation of the price effect. Yet, as more months pass, it becomes increasingly unlikely that new car buyers are aware of and basing their decision on the older ratings. Third, as detailed in Section 4.3, the interpretation of the estimates also depends on whether car buyers were aware of the restatement as this might impact their beliefs about and trust in fuel-economy ratings and realized fuel economy.

Table 4 presents evidence that our results are not driven by the period shortly after the restatement. The point estimates do not change much when we omit up to 12 months following the restatement. The estimates in columns 2-12 of Table 4 are quite similar to the full-sample estimate in column 1. 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. This suggests that our results apply more broadly to settings in which fuel-economy ratings change without any known issues of misreporting.

3.1.2 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 5, we explore heterogeneous treatment effects with respect to these variables.[19] Columns 1 and 2 replicate our preferred specification from Table 2. Columns 3 and 4 allow the treatment effect to vary by model year.

[19] Appendix Tables B.5 and B.6 explore heterogeneity by make and vehicle class.

Table 4: Effect on Transaction Prices Excluding the Months Closest to the Restatement

	(1)	(2)	(3)	(4)	(5)	(6)
	Number of Post-Months Excluded					
	0 (base)	1	2	3	4	5
$1(Post\ Restatement)_t$	-0.012	-0.012	-0.013	-0.013	-0.015	-0.016
$\times 1(Affected\ Model)_j$ (Logs)	(0.003)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
$1(Post\ Restatement)_t$	-294	-310	-324	-341	-389	-408
$\times 1(Affected\ Model)_j$ (Levels)	(91)	(98)	(106)	(115)	(125)	(136)
N	1.52m	1.47m	1.43m	1.38m	1.34m	1.29m
	(7)	(8)	(9)	(10)	(11)	(12)
	Number of Post-Months Excluded					
	6	7	8	9	10	11
$1(Post\ Restatement)_t$	-0.016	-0.013	-0.012	-0.012	-0.011	-0.010
$\times 1(Affected\ Model)_j$ (Logs)	(0.006)	(0.006)	(0.006)	(0.007)	(0.008)	(0.009)
$1(Post\ Restatement)_t$	-415	-368	-347	-364	-330	-320
$\times 1(Affected\ Model)_j$ (Levels)	(147)	(158)	(171)	(185)	(204)	(229)
N	1.25m	1.20m	1.16m	1.11m	1.07m	1.02m
Year-Month × Class FE	Y	Y	Y	Y	Y	Y
Year-Month × 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) × DMA FE	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 evel regressions equals 0.95-0.96. Standard errors clustered by VIN10.

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 3.2 for details) or automakers facing customers with different demand elasticities for the newest model year vehicles.

Table 5: Heterogeneous Effects of the Restatement on Transaction Prices

	Primary		**Model Year**		**Δ GPM**	
	(1)	(2)	(3)	(4)	(5)	(6)
	Logs	Levels	Logs	Levels	Logs	Levels
$1(\textit{Post Restatement})_t \times 1(\textit{Affected Model})_j$	-0.012	-294				
	(0.003)	(91)				
$1(\textit{Post Restatement})_t \times 1(2011-2012\ \textit{Affected Model})_j$			-0.017	-544		
			(0.006)	(128)		
$1(\textit{Post Restatement})_t \times 1(2013\ \textit{Affected Model})_j$			-0.011	-259		
			(0.004)	(98)		
$1(\textit{Post Restatement})_t \times 1(\textit{Affected Model})_j \times \Delta GPM$					-2.92	-66544
					(0.90)	(22470)
Year-Month × Class FE	Y	Y	Y	Y	Y	Y
Year-Month × 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) × 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. ΔGPM refers to the change in the gallons-per-mile from the restatement. All estimations are weighted by monthly sales. Standard errors clustered by VIN10.

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.[20] 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). These results suggest that consumers do not respond to the magnitude of the restatement perfectly proportionately (otherwise the

[20]The results have nearly identical implications if we use miles-per-gallon.

mean change in gallons-per-mile should have led to the mean change in price) but do respond in the expected direction on average.

3.2 Effects on Other Outcomes

In equilibrium, it is possible for there to be other adjustments as well. Busse, Knittel and Zettelmeyer (2013*a*) 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 4 below.

First, it is important to point out that our setting is quite different from Busse, Knittel and Zettelmeyer (2013*a*). 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, especially in the short run. They require physical adjustments to assembly lines and renegotiation of contracts with suppliers. These factors would tend to dampen the response in model year 2013 sales, but even so, we would expect some reduction in sales (i.e., a negative elasticity). In contrast, supply was very likely inelastic for model year 2011 and 2012 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 specific models phase in and phase out. 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 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 4.2.

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.

4 Implications for the Valuation of Fuel Economy

4.1 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:

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

where $Price_{jrt}$ is the vehicle transaction price and ΔG_{jt} is the change in the discounted lifetime fuel cost due to the restatement.[21] 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.[22] If sales do not adjust, we can interpret a value of -1 as full valuation—where an increase in expected future fuel costs is entirely reflected by a decrease in the purchase price—but discuss the implications of elastic supply in Section 4.2.

There are four 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. This

[21] We use ΔG_{jt} to denote that we are focusing on the variation in G_{jt} that is coming from the change in fuel economy due to the restatement (which varies by vehicle model). ΔG_{jt} is thus equal to zero for all non-affected models and it is also equal to zero in the pre-restatement period for affected vehicles. The only other source of variation in ΔG_{jt} could be from changes in expected future gasoline prices at the time of purchase of an affected vehicle. This variation is modest given that gasoline prices were similar around the time of 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).

[22] Much 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, 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.

is because the micro-foundation of Equation (2) is an inverted market share equation. 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. Finally, one may be concerned about SUTVA, but our robustness checks for our estimation of Equation (1) show that SUTVA violations should not be an issue in our setting.

We first estimate Equation (2) using a baseline set of assumptions in constructing ΔG_{jt}: expected driving from the NHTSA, vehicle survival probabilities from Jacobsen and van Benthem (2015) (source: R.L. Polk (1993-2009)), 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 6 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 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.[23] 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. A natural interpretation of this difference is that there is considerably less elastic supply for the (already produced) 2012 model year than for the 2013 vintage. 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

[23]For 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.

Table 6: The Valuation of Fuel Economy Based on the Equilibrium Price Change

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$r = 1\%$		$r = 4\%$		$r = 7\%$		$r = 12\%$	
$1(\Delta Lifetime\ Fuel\ Costs)_{jt} \times$	-0.14		-0.17		-0.20		-0.25	
$1(Affected\ Model)_j$	(0.05)		(0.06)		(0.07)		(0.08)	
$1(\Delta Lifetime\ Fuel\ Costs)_{jt} \times$		-0.33		-0.39		-0.46		-0.58
$1(2011 - 2012\ Affected\ Model)_j$		(0.17)		(0.20)		(0.24)		(0.30)
$1(\Delta Lifetime\ Fuel\ Costs)_{jt} \times$		-0.13		-0.16		-0.18		-0.23
$1(2013\ Affected\ Model)_j$		(0.05)		(0.06)		(0.07)		(0.08)
Year-Month × Class FE	Y	Y	Y	Y	Y	Y	Y	Y
Year-Month × 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) × 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 R.L. Polk (1993-2009), 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.

parameter of 0.39 for model years 2011-2012 and 0.16 for model year 2013. 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 8). 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.[24]

[24] Leard, Linn and Zhou (2018) report a real borrowing rate of 1.3% for the period October 2009 to September 2014. Using the same approach as Allcott and Wozny (2014), we find that nominal auto loan rates for new vehicles were in the 4% range during our sample period (Board of Governors of the Federal Reserve System 2010-2013); after accounting for CPI increases of 1.5-2.1% for the period 2012-2014 (Integrated Pub-

We cannot emphasize enough that with different sets of assumptions, the undervaluation parameter would change. For a wide enough range of assumptions, the valuation parameter can be as low as zero or as high as one. However, 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.

4.2 Bounding Analysis

4.2.1 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 3 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.[25] 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*a*). The fourth panel allows for imperfect competition.

We begin with the case of perfectly inelastic supply (i.e., a zero quantity effect). Panel

lic Use Microdata Series 1962-2017), the real auto loan rate was approximately 2.5%. The federal funds rate in November 2012 was 0.16% (Macrotrends.com 2020).

[25]We assume locally parallel shifts in the demand curve, which is supported by the limited role for consumer selection as discussed in Section 3.1.

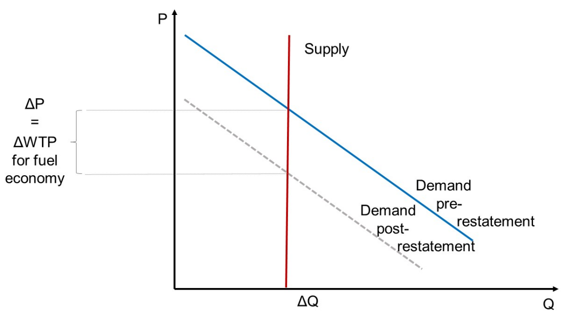

(a) Inelastic Supply

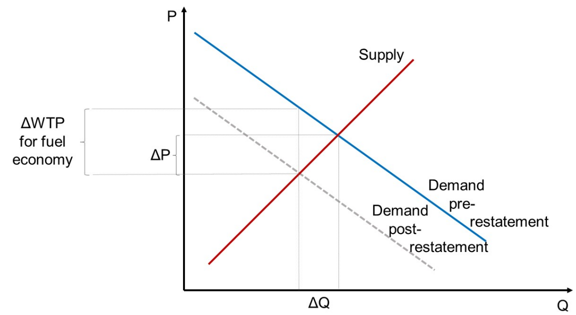

(b) Upward-Sloping Supply

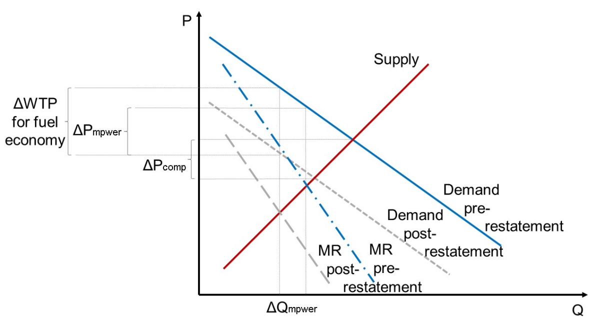

(c) Market Power and Upward-Sloping Supply

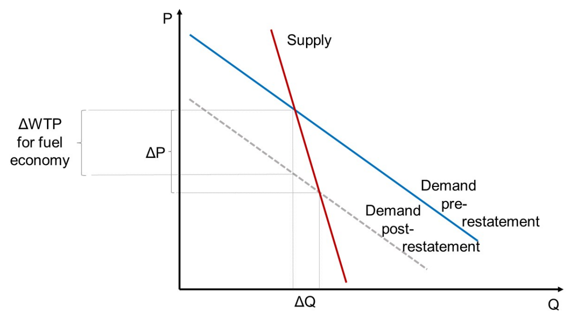

(d) Downward-Sloping Supply

Figure 3: Interpretation of the Equilibrium Effect

Notes: Panels (a), (b), and (d) 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. Panel (c) compares the change in equilibrium price for the competitive case versus the market power case.

A shows that under perfectly inelastic supply the change in equilibrium price (our γ) is 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 6 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 6 based on different assumptions of the slope of supply.

Panel C allows for imperfect competition with an upward-sloping supply curve. When there is imperfect competition, the marginal revenue lies below the residual demand, allowing firms to earn a markup. Therefore 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, 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. In Appendix D.3 we derive the results discussed in this section more formally with a simple analytical model.

Finally, Panel D assumes the less likely case of downward-sloping supply under perfect competition.[26] We cover this case for completeness, as it is consistent with our positive (though not statistically significant) point estimate of the quantity effect from the restatement. Localized economies of scale are one possible economic justification for downward-sloping supply, but we recognize this would be atypical. In this scenario, the change in equilibrium price overestimates the willingness-to-pay. This would suggest that our estimates in Table 6 are biased upwards and that the true willingness-to-pay

[26]The 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.

is even closer to zero.

4.2.2 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 supply, as in Panel B of Figure 3, 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.[27] Berry, Levinsohn and Pakes (1995) find vehicle model-level own-price demand elasticities ranging to -6.5, while Busse, Knittel and Zettelmeyer (2013*a*) 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 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

[27] In 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.

aligned with the summary statistics on vehicle prices for Hyundai and Kia in Table 1).[28]

Using these assumptions, the adjustment formulas in Appendix D.3, and the $294 reduction in equilibrium price due to the restatement, Table 7 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 (less likely) +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 7: Interpretation of Equilibrium Change in Prices w.r.t. Different Supply Curves

Quantity Effect (%)	Willingness-to-Pay ($) η_D = -6	Willingness-to-Pay ($) η_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 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 7 showing the willingness-to-pay for quan-

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

tity 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 6, the results in Table 7 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 6. 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.[29] Assuming a supply elasticity closer to zero, the effect on the valuation parameter should be much more modest. In Table 7, 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.

4.3 Comparison to Previous Literature

Table 8 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*a*), Sallee, West and Fan (2016), and Grigolon, Reynaert and Verboven (2018) are all close to one, 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

[29] As 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.

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).[30] That the payback time implied by our results is similar to the payback time reported by automakers is striking, but of course is not direct evidence that either our results or automakers' assumptions are correct.

Table 8: Comparison of Estimates with Other Studies

Studies using exact valuation parameter	*r*	*valuation parameter*
Sallee, West and Fan (2016)	5%	1.01
Allcott and Wozny (2014)	6%	0.76
Own Estimate from Restatement	5%	[0.17-0.42]
Own Estimate from Restatement	6%	[0.18-0.44]
Studies using approximate valuation parameter		
Busse, Knittel and Zettelmeyer (2013*a*)	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.

4.4 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.

[30]This 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.

4.4.1 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*a*) 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.

4.4.2 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 useful, as it assures that other vehicle attributes are held constant, and it leverages exogenous variation in the fuel economy ratings, which is highly relevant for understanding the consumer response to fuel-economy standards. Using variation from the fuel-economy rating restatement could affect the interpretation of our results in several ways, which we discuss in this section.

First, it is important to consider how variation in the fuel-economy rating influences new car buyers' beliefs about future fuel costs. New car buyers likely form their beliefs about the future fuel costs of each potential vehicle choice using a variety of sources of information. The fuel-economy ratings are likely to be the most important source for most consumers, as they are widely recognized, on all car-comparison websites, and displayed prominently on the label on each new car in the dealer lot. But it is also likely that at least some consumers draw upon other sources of information, including conversations with other car owners, advertisements, or their own past experience with vehicles of similar attributes. Our study exogenously changes only the fuel-economy ratings, leaving other sources of information unchanged.

How might this affect the interpretation of our results? There are a few possibilities. As a starting point, it is useful to note that survey evidence from multiple studies suggests that consumer beliefs about future fuel costs—based on all existing information—on average are close to unbiased when using the official fuel economy ratings (Allcott 2013; Allcott and Knittel 2019; Andor et al. 2020). This evidence suggests that it is unlikely that there is some general set of biased beliefs about future fuel costs.

However, there is the possibility is that at least some of the new car buyers prior to the restatement were already aware that the affected vehicle ratings were incorrect, and thus already incorporated this information into their car-buying decision process. Should this be the case, our estimation would understate the price effect and thus understate the valuation of fuel economy. However, for it to be the case, the new car buyers would have to have known exactly which vehicles were affected and unaffected by the restatement. This seems unlikely to us. The restatement came as a complete surprise to the automotive community—we could find no articles in the mainstream media or trade press about it prior to November 2012. That said, it is possible that new car buyers had another source of information that somehow told them the true fuel economy of the affected vehicles. We explored this possibility extensively by looking for information that was available prior to the restatement that would support it. We did find blog posts by automobile aficionados prior to the restatement that indicated that they were having a hard time achieving the rated fuel economy, but 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. Thus, while we cannot rule out some other source of information that informed new car buyers about the true fuel economy of the affected Hyundai and Kia vehicles, we found no evidence to support this hypothesis.

It is also possible that new car buyers after the restatement bring in other information from prior to the restatement. For example, consumers may still remember advertisements describing the high fuel economy of some of the Hyundai and Kia affected vehicles from prior to the restatement, and may not fully update their prior beliefs after the restatement. If this was the case, we would expect new car buyers to update their beliefs over time, which would suggest that our valuation parameter would increase over time as consumers more fully adjust. However, in Table 4, 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. The fact that we do not see a change in the valuation parameter over time since the restatement suggests that belief updating would have to be very slow for prior information to be a major explanation for our results.[31]

Even so, the results in Table 4 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 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

[31] In reality, we believe it is most likely that most new car buyers after the restatement started their search well after the restatement occurred and only saw the newer, lower fuel-economy rating on car comparison websites—so they very well could be unaware that that the rating had changed at all.

results from most of the previous studies would also only be useful for understanding medium-run consumer responses.

Second, a closely related issue is lack of trust in fuel-economy ratings. It is possible that the restatement itself had an impact on consumers' overall trust and willingness to rely on the official ratings. Because our empirical strategy compares affected to unaffected vehicles, for a change in trust to change the interpretation of our valuation parameter, there would need to be an asymmetric change in trust between the affected and unaffected vehicles.[32] Note that we do not need to assume full trust in the ratings for our identification strategy to work. For example, mistrust does not affect our estimates if consumers on average uniformly discount all ratings, but trust *changes* in ratings. Since our specification is based on changes in gallons-per-mile, this particular type of mistrust does not affect our estimates. If instead consumers on average do not trust a change in the fuel-economy ratings (e.g., the change is 2 miles-per-gallon but consumers believe that the change was 1 mile-per-gallon), this would constitute another type of biased beliefs. We view this as unlikely because new car buyers would have to do substantial research to determine which models and trims received an updated rating. 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. Should it be the case, it could be another mechanism for why we observe undervaluation of fuel economy. Note that when fuel-economy standards are increased, the government adjusts the official ratings. Thus, biased beliefs about what changes in ratings imply for future fuel costs that lead to undervaluation (perhaps due to other sources of information that influence the decision process) would have direct relevance to policy.

Third, by using variation in fuel-economy ratings, we may be capturing aspects of human behavior not captured in previous studies that use gasoline price variation. As

[32]For example, if consumers believe that ratings for all vehicles were off by a fixed amount in miles-per-gallon (strictly speaking, gallons-per-mile, but we ignore this subtlety), our estimate should be unaffected. To see this, consider the case where consumers believe that every rating is 3 miles-per-gallon too high. They would still correctly interpret a change from 25 to 23 miles-per-gallon as a 2 miles-per-gallon change, even if they believe the true fuel economy went from 22 to 20 miles-per-gallon. Alternatively, car buyers might believe that all ratings are overstated 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.

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 to 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.

In summary, by using variation in fuel-economy ratings, the interpretation of our coefficients may differ from those in studies primarily using gasoline price variation. By focusing on variation in fuel-economy ratings, we capture aspects of behavior that previous estimates using gasoline price variation do not. We cannot differentiate between several different possible behavioral explanations for our undervaluation result, including inattention to fuel-economy ratings, a lack of sophistication in processing fuel-economy information, large cognitive costs in considering future fuel expenditures, and incorrect beliefs, but all of these would have policy relevance.

4.4.3 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 8, 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.

5 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 exogeneous 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.

Abaluck, Jason, and Jonathan Gruber. 2016. "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance." *American Economic Review*, 106: 2145–2184.

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 Christopher Knittel. 2019. "Are Consumers Poorly Informed about Fuel Economy? Evidence from Two Experiments." *American Economic Journal: Economic Policy*, 11(1).

Allcott, Hunt, and Michael Greenstone. 2012. "Is There an Energy Efficiency Gap?" *Journal of Economic Perspectives*, 26(1): 3–28.

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.

Andor, Mark, Andreas Gerster, Kenneth Gillingham, and Marco Horvath. 2020. "Running a Car Costs Much More Than People Think." *Nature*, 580: 453–455.

Autoblog.com. 2012. "Hyundai, Kia Admit Exaggerated Mileage Claims, Will Compensate Owners." `https://www.autoblog.com/2012/11/02/hyundai-kia-admit-exaggerated-mileage-claims-will-compensate-o/` (Accessed on August 29, 2020).

Autoblog.com. 2014. "Hyundai-Kia Fuel Economy Errors Trigger $300M in Federal Penalties." https://www.autoblog.com/2014/11/03/hyundai-kia-300-million-mpg-penalties/ (Accessed on August 29, 2020).

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.

Board of Governors of the Federal Reserve System. 2010-2013. "2013 Survey of Consumer Finances." https://www.federalreserve.gov/econres/scf_2013.htm (Accessed on August 29, 2020).

Busse, Meghan, Christopher Knittel, and Florian Zettelmeyer. 2013*a*. "Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review*, 103(1): 220–256.

Busse, Meghan, Christopher Knittel, and Florian Zettelmeyer. 2013*b*. "Data for: Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review*, 103(1): 220–256. https://www.openicpsr.org/openicpsr/project/112589/version/V1/view.

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.

DataOne. 2011-2014. "VIN Decoding Data." unpublished data (Accessed on June 30, 2020).

Davenport, Coral. 2018. "Trump Administration Unveils Its Plan to Relax Car Pollution Rules." *New York Times*. August 2nd.

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." National Bureau of Economic Research Working Paper 25722.

Integrated Public Use Microdata Series. 1962-2017. "Note on Adjusting Dollar Amount Variables for Inflation. United States Department of Labor." `https://cps.ipums.org/cps/cpi99.shtml` (Accessed on September 2, 2019).

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.

Kia Motors America, Inc. 2020. "Kia MPG Information: Compensation." `https://kiampginfo.com/overview/compensation/` (Accessed on June 23, 2015).

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.

Macrotrends.com. 2020. "Federal Funds Rate - 62 Year Historical Chart." `https://www.macrotrends.net/2015/fed-funds-rate-historical-chart` (Accessed on August 29, 2020).

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 Highway Traffic Safety Administration. 2018. "Average Vehicle Miles Traveled." https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/ld_cafe_co2_nhtsa_2127-al76_epa_pria_181016.pdf (Accessed on August 27, 2020).

National Research Council. 2015. *Cost, Effectiveness, and Deployment of Fuel Economy Technologies for Light-Duty Vehicles.* The National Academies Press:Washington, DC.

New York Times. 2012. "Hyundai and Kia Acknowledge Overstating the Gas Mileage of Vehicles." https://www.nytimes.com/2012/11/03/business/hyundai-and-kia-acknowledge-overstating-the-gas-mileage-of-vehicles.html (Accessed on August 29, 2020).

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.

R.L. Polk. 1993-2009. "Data on Vehicles in Operation." unpublished data (Accessed on July 30, 2019).

R.L. Polk. 2011-2014. "Data on All Dealer-Reported U.S. New Vehicle Transactions." unpublished data (Accessed on October 13, 2016).

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.

U.S. Energy Information Administration. 2011-2014. "Annual Nationwide Gasoline Prices." https://www.eia.gov/outlooks/aeo/data/browser/#/?id=3-AEO2018&sourcekey=0 (Accessed on October 1, 2019).

U.S. Energy Information Administration. 2020. "Energy and the Environment Explained." `https://www.eia.gov/energyexplained/index.php?page=environment_where_ghg_come_from` (Accessed on August 23, 2020).

Verboven, Frank. 2002. "Quality-Based Price Discrimination and Tax Incidence: Evidence from Gasoline and Diesel Cars." *RAND Journal of Economics*, 33(2): 275–297.