Title
Essays in Environmental and Behavioral Economics

Author
Panassie, Yann

Publication Date
2018

UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays in Environmental and Behavioral Economics

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Economics

by

Yann Panassié

Committee in charge:

Professor Richard Carson, Chair
Professor Roger Bohn
Professor Roger Gordon
Professor Joshua Graff Zivin
Professor Mark Jacobsen

2018

The dissertation of Yann Panassié is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California San Diego

2018

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

Thanks to my advisor, Richard Carson, for the guidance, advice, and conversations that have led to this thesis. Thanks also to Josh Graff Zivin, Mark Jacobsen, Roger Gordon, and Roger Bohn for the various inputs which have greatly improved it.

Thanks to Yanjun Liao for being a great co-author, friend, and a pleasure to work with.

Thanks to my mom for her endless support.

Thanks to Jeffrey Perloff, Sofia Villas-Boas, Shachar Kariv, and Benjamin Wade for inspiring me to pursue a doctorate.

Finally, thanks to all the friends without whom my time in—and outside of—San Diego would not have been the same, whether at the office, in our homes, downtown, or playing sports on the field, court, or sand. You know who you are.

Chapter 1, in part, is currently being prepared for submission for publication of the material. Panassié, Yann. The dissertation author was the sole investigator and author of this material.

Chapter 3, in part, is currently being prepared for submission for publication of the material. Panassié, Yann; Liao, Yanjun. The dissertation author was the co-investigator and co-author of this material.

VITA

2012	B. S. in Environmental Economics and Policy, University of California Berkeley
2012-2018	Graduate Teaching Assistant, University of California San Diego
2016	C. Phil. in Economics, University of California San Diego
2018	Ph. D. in Economics, University of California San Diego

ABSTRACT OF THE DISSERTATION

Essays in Environmental and Behavioral Economics

by

Yann Panassié

Doctor of Philosophy in Economics

University of California San Diego, 2018

Professor Richard Carson, Chair

This dissertation studies questions in environmental economics by exploring the mechanisms through which government and private decisions interact in the transportation and housing markets. These have important environmental and distributional consequences in terms of mitigation of and adaptation to climate change. In Chapter 1, I compare new vehicle sales in the United States and Canada to determine whether updated EPA fuel economy labels introduced in 2012 succeeded in altering consumers' new vehicle purchase choices. I find small savings in gasoline consumption through a 1.5 percentage point increase in small car market shares, a corresponding decrease in SUV shares, and a 6% increase in the valuation of small SUVs' fuel economies. In Chapter 2, I study gasoline price volatility in California by estimating the gasoline price elasticity

of demand for driving, and show that this parameter is both highly inelastic and likely to vary over time. Chapter 3 focuses on the impacts of hurricanes on the Florida housing market. I show that hurricanes cause an equilibrium increase in home prices and a concurrent decrease in transaction probability, lasting up to three years. With supplementary evidence from demographic trends, I conclude that the main driver of these dynamics is a negative transitory shock to the housing supply in the aftermath of hurricanes as homes recover from physical damages. I further observe that new homeowners have higher incomes, resulting in a permanent shift in the demographic composition of disaster-prone areas, and suggesting important implications about the expected costs and distributional impacts of future federal disaster relief spending.

Chapter 1

The EPA Matters: Evidence from the 2013 Update to Fuel Economy Labels

1.1 Introduction

The US Environmental Protection Agency (EPA) redesigned the mandatory fuel economy labels which are affixed to a side window of all new vehicles for sale on dealer lots, starting with all 2013 model-year vehicles. In addition to the information already provided by the previous labels, the update included fuel cost saving or spending over five years relative to the average new vehicle,[1] and more prominently displayed the combined miles per gallon (mpg) figure. The redesign also added emissions ratings, as well as estimated fuel consumption in gallons per 100 miles because, as noted by the EPA, "unlike mpg, consumption relates directly to the amount of fuel used, and thus to fuel expenditures." Figures A1 and A2 respectively provide examples (from

[1]Under assumptions of 15,000 miles driven per year and a yearly-revised gasoline price ($3.70 in 2013).

hypothetical vehicles) of the labels used between 2008 and 2012, and the updated labels applied to all new vehicles since model-year 2013. The EPA cites multiple reasons for the redesign, including that "shoppers will have more information [...] to help save money on fuel and cut down on harmful pollution." It also sought to provide information specialized by vehicle types to reflect the growing prevalence of hybrid, fully electric, and other alternative fuel vehicles (see Figure A3), as well as to satisfy a new government requirement to include greenhouse gas emissions and smog pollution ratings. Some authors in the labeling literature, however, caution us about the potential drawbacks of cluttering information on consumer good labels. Chaffee and McLeod (1973), for example, finds that increasing the amount of information on a label may make processing any of it much more difficult to consumers, which, according to the nutrition labeling literature, can in turn result in people ignoring the labels altogether (Teisl and Roe 1998). It is also worth noting that in 2016, Canada's own fuel economy label (Figure A4) was updated to one that visually resembles the EPA's 2013 revision, but with gasoline consumption in liters per 100 km taking a more prominent position than the mpg fuel economy figure, and no relative savings estimate like the one found on the new US label.

My objective in this paper is to evaluate whether the new EPA labels have succeeded in altering consumers' behavior in the form of their aggregated purchase decisions, both through changes in their valuation of fuel economy within different segments of vehicles, and through the valuation of fuel economy implied by relative changes in the market shares of segments themselves. The rest of the paper is organized as follows. Section 1.2 briefly discusses the existing fuel economy literature, Section 1.3 describes the data I use in the analysis, Section 1.4 introduces the difference-in-differences identification strategy, Section 1.5 estimates the results and interprets the key findings, and Section 1.6 concludes.

1.2 The Fuel Economy Literature

A rich literature about consumer demand for fuel economy goes back at least four decades and has important implications regarding the effectiveness of gasoline taxes, CAFE standards, and other tools used to achieve more fuel efficient vehicle fleets. This literature can be divided into two broad categories, with one group of studies using vehicle market share data over time, and the other using individual choice data. Studies of the US market share of vehicles over time often model the supply side, usually as an oligopolistic market which must take into account the effects of both CAFE standards and the price of gasoline. Building on the conditional logit model in McFadden (1973), Berry et al. (1995) contributes an important instrumental variable approach to address the endogeneity issue by using characteristics of other vehicles from the same manufacturer or segment to predict prices, within-segment market shares, and other possibly endogenous characteristics. BLP explicitly aggregates consumer preferences into a parametric market demand system and combines this with cost function and pricing behavior assumptions to generate equilibrium prices and quantities. Klier and Linn (2012) further develops the instrumental variable methodology by using the characteristics of same-make vehicles with shared engine platforms across different segments as instruments for prices and other endogenous characteristics. On the less parametric side of the literature, Busse et al. (2013) looks for evidence of consumer myopia about future fuel costs using both individual level and aggregate vehicle choice data, and interpret the results of equilibrium prices and market shares responding to gasoline price changes as almost complete lack of myopia.

Yet as carefully documented in a literature review by Greene (2010), after four decades of research, there is still little consensus on whether consumers correctly, over-, or under-value fuel cost savings when making vehicle purchasing decisions. Studies are about evenly divided with no discernable pattern or trend and widely varying estimates,[2] and many authors conclude that consumers' preferences for fuel economy are heterogeneous. Central to the debate is the issue of

[2]See Figure A5.

the standard neoclassical assumption that consumers trade off the present cost of more expensive, more fuel efficient alternatives against discounted future fuel costs. One might expect this to be an unlikely calculation for all consumers to actually make because it requires specific assumptions about expected fuel prices and total miles driven over the vehicle's life, and a discount rate. Alternatively, some people may simply be unable (or unwilling to spend the time) to figure out how to compute fuel savings. In-depth interviews with representative California households leads Turrentine and Kurani (2007) to conclude that not only do households not systematically analyze fuel expenditures or track them over time, but that even the presumably most mathematically capable ones make large errors in estimating fuel costs over time despite being explicitly given all the necessary information to make the calculations.

Many consumers are perhaps unlikely to realize that fuel costs nonlinearly depend on the mpg level.[3] In fact, Larrick and Soll (2008) finds that people perceive fuel savings to increase linearly with miles per gallon, leading them to both under-value differences in mpg at low mpg levels *and* overvalue mpg at relatively higher levels (Figure A6). The scenario in the study's lab experiment features a hypothetical vehicle worth \$20,000 and rated at 15 mpg. Participants are asked for their willingness to pay (the mean of which is connected by the blue line in the figure) for the vehicle at different mpg levels assuming that they will drive 10,000 miles per year for 10 years, and that the price of gas is constant at \$2.80. The yellow curve shows the vehicle's actual value under the given assumptions after accounting for the fuel savings associated with the mpg level specified by the horizontal axis. The phenomenon of this linear valuation of mpg by consumers is termed "MPG Illusion" by the authors, and is further documented by Allcott (2013), which finds evidence of illusion in nationally representative vehicle ownership and fuel expenditures survey data. Allcott's simulations additionally imply that this may have an important effect on market shares.

[3]The dependence is inversely proportional to mpg, and directly proportional to fuel price.

1.3 Data

In this paper, I primarily rely on data from Ward's Auto Infobank between 2009 and 2015. In particular, Ward's reports total monthly US and Canadian sales of new cars and light trucks[4] by make and model (e.g. Volkswagen Golf), but not by trim level (e.g. Volkswagen Golf TSI S). Ward's does, however, collect an exhaustive set of specifications at the trim level, as well as track the distribution of engines installed by model-year. Because fuel efficiency, price, and other characteristics can vary considerably by trim, I use this engines installation data to create a more precise measure of average vehicle characteristics for each model whose engines data are available.[5] While this procedure does not yield an exact trim match because more than one trim level is often associated with each model-engine, fuel economy (and other characteristics of interest) tend to vary most substantially with different engine installations. A 2010 Honda Accord Sedan, for example, achieves 23/33 or 23/34 mpg (EPA city/highway estimates) in all its trims sharing a 2.4 liter 4-cylinder engine, while both iterations of the car with a 3.5 liter 6-cylinder are rated at 20/30 mpg. Horsepower and torque are essentially the same in all 4-cylinder trims, exactly the same in both 6-cylinder versions (and appreciably higher than the 4-cylinders), and the only significant differences across different trims sharing an engine are the base prices, which range from about $22,000 to $28,000 for the 4-cylinders, and $28,000 and $30,000 for the 6-cylinders. Such specification variation is representative of the typical car model, which tends to have one to three different engines for two to fewer than ten unique trims in the majority of vehicles, similar fuel economy and horsepower numbers across trims sharing an engine, but some price variation between trims with the same engine. In all cases, I assign the entire mass of each of a model's possible engine installations to the model's cheapest trim which features that engine.

For each year, I first match models' engine installations to the specifications data, and then take a sales-weighted average of the specifications from each matched trim of a model to

[4]Light trucks include all SUVs, vans, and trucks with a gross vehicle weight under 14,000 pounds.

[5]This covers almost all models, and the few that aren't available are still included in all analyses at their base trims.

construct average characteristics by model.[6] I then merge these data with the monthly sales data, and subsequently match the joined sales and specifications to files classifying models into twenty-seven classes (e.g. lower small car, large pickup, etc.), which I aggregate into eleven mutually exclusive segments as determined by size, price range, horsepower, and sometimes more subjective measures such as body style and functionality. Five segments belong to cars: small, midsize, large, luxury, and sport, while the remaining six are small, midsize, large, and luxury SUVs, trucks, and vans. Finally, I create a combined fuel economy measure for every model in each month using the same method as the EPA, which is given by the weighted harmonic mean of its estimated city (55% weight) and highway (45% weight) fuel economies in mpg.[7]

The following summary statistics briefly describe some important features of the US data. Small cars are unsurprisingly the cheapest, with average offerings ranging from about $16,000 in 2009 to $18,000 in 2015, while luxury cars are the most expensive and sell from an average of $47,500 in 2009 to about $55,000 in 2015. Because fuel costs are such a smaller percentage of total vehicle expenditures for luxury car and SUV buyers, who are also likely to be wealthier, we should expect this segment's sensitivity to fuel economy and gasoline prices to be lesser. Midsized cars have the most hybrid models until 2012 (then at 10 vehicles with hybrid engine variants), but are then surpassed by the number of hybrid or fully electric offerings of luxury cars, which reaches 12 by 2013. Midsized hybrid vehicles, however, are both notably more fuel efficient on average and have much higher sales volumes than their luxury cousins. In general, fuel economy is inversely related to vehicle weight and engine output, and mean horsepower does not rise much over the seven years of data in any segment except for trucks, and large and sports cars. As seen in Figures 1.1, 1.2, and A7, fuel economy is generally increasing over the data period. But as illustrated by Figure 1.10, any trend for fuel economy to be increasing relatively faster than horsepower over time has been of rare occurrence. Given their lower weights and

[6] Weights are given by the percentage of vehicles sold within a model featuring each possible engine.

[7] Averaging is always performed in terms of fuel consumption (i.e. gallons per mile) to obtain the correct sales-weighted fuel economies within models or segments.

typically smaller engines, we should on average expect smaller vehicles to be more fuel efficient than larger ones. Figures 1.1 and 1.2 below consider the sales-weighted average fuel economies of cars and other vehicles respectively. One of the most surprising features of this data is that

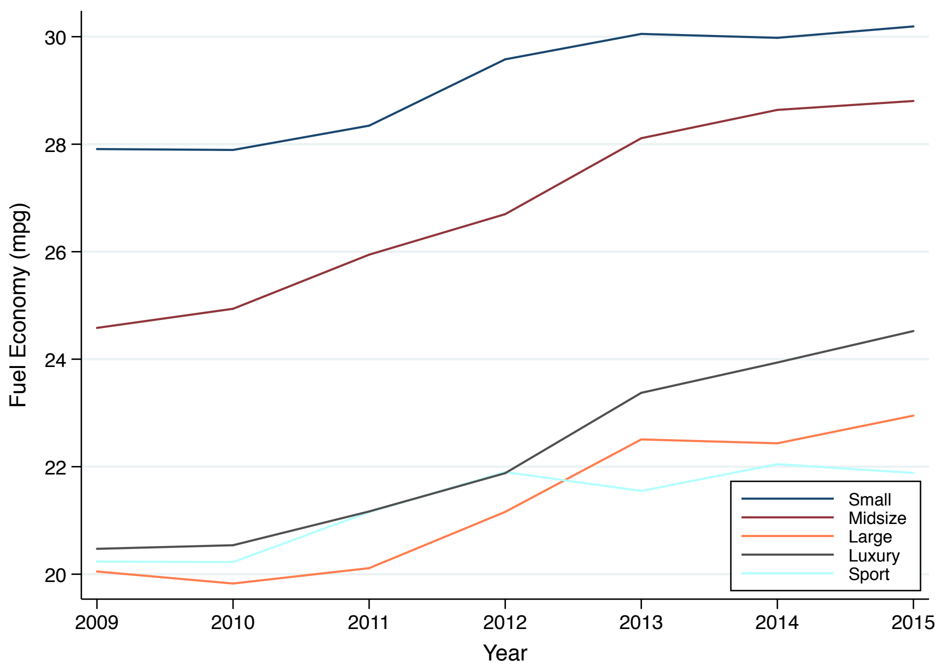

Figure 1.1: US sales-weighted fuel economy by car segment

midsized SUVs are more efficient than small ones in four out of seven years. Examining this puzzle more closely reveals an even greater one: their fuel consumptions are nearly even despite midsized SUVs on (weighted) average being about 5 to 15% heavier and having 10 to 20% more powerful engines. Digging even deeper, however, informs us that a much higher share of small SUV models are vehicles with emphasis on off-road capabilities (e.g. the Jeep Wrangler), and include characteristics, such as four-wheel drive, which often come at the expense of fuel economy. Furthermore, the share of midsized SUVs sold with hybrid drivetrains is over 3 times as large as that of small SUVs, but this contributes less to their remarkably close fuel economies because hybrids only make up 1% of midsized SUV sales. Finally, note the interesting divide around 25 mpg between small and midsized cars, and vehicles from all other segments. While

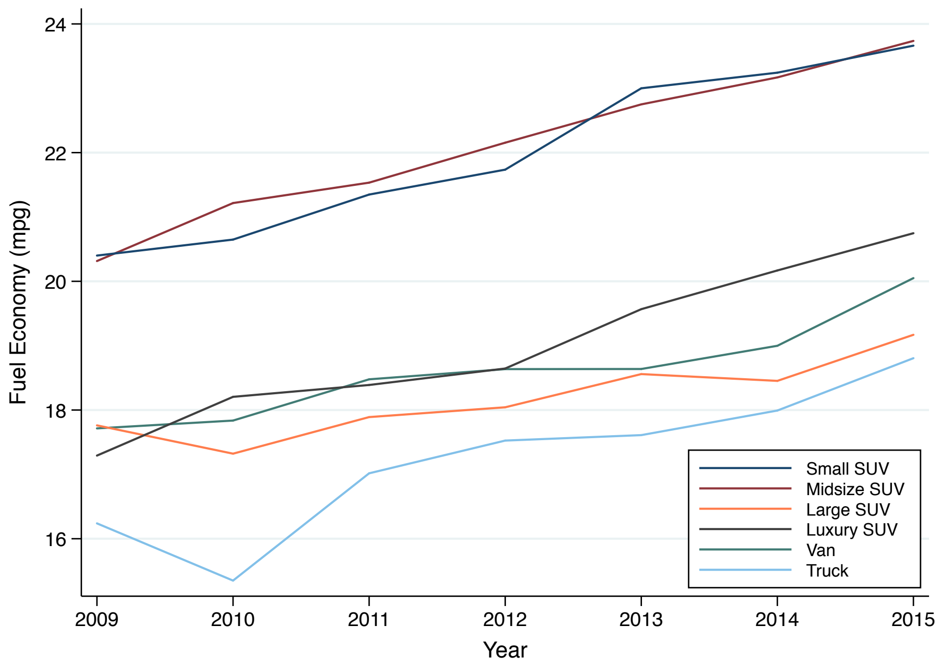

Figure 1.2: US sales-weighted fuel economy by light truck segment

this threshold should probably not be used as a precise behavioral parameter for fuel economy valuation in any econometric model, it can still be seen as a reference for predicting what we could expect to observe in the event that people generally behave in the way documented in Larrick and Soll (2008): as shown by relative slopes in Figure 1.3, a 1 to 10% range of discount rates proves inconclusive about whether small improvements in fuel economy are under- or over-valued between 19 and 25 mpg, but it does appear to be consistently over-valued for improvements above 25, and significantly so over 33 mpg. Therefore, labels correcting such perceptions could actually result in an unintended reduction in the efficiency valuation of small or midsized car buyers.

1.4 Identification

The identification strategy I rely on in this paper is a treatment-intensity difference-in-differences (DID) with Canada as the control group, and where segment market shares and

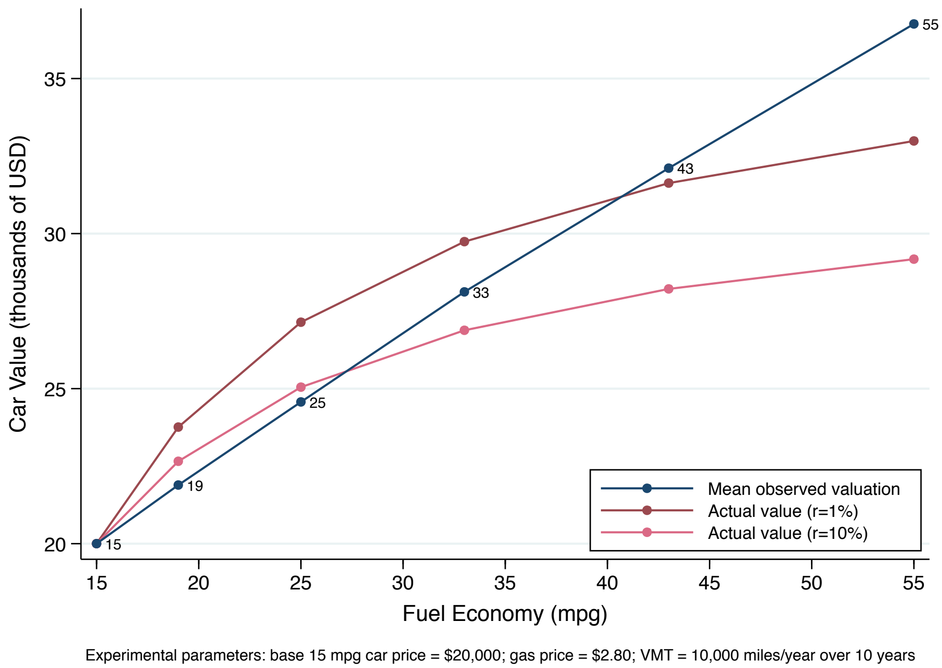

Figure 1.3: "MPG Illusion" revisited

segment sales-weighted average fuel consumption are the left hand side variables. The main advantage of this approach is that it allows me to nonparametrically estimate the effect of the labels on US consumers' purchase choices from two markets relying on nearly identical supply environments, with almost every model-engine pair available in the US also existing in Canada (sometimes under a different nameplate). In particular, this allows me to avoid the sometimes tenuous assumptions required in modeling the supply side of the market, as well as to ignore the well-studied but still debated impacts of CAFE standards.

Because I only observe a single treated and control group for each segment, it is especially important that there be no other significant policies during the same time period affecting the American and Canadian auto industries differently as this would confound my estimates. One of the most noteworthy of such possible policies from the past decade began in 2006, when the US started offering federal tax rebates of up to $7,500 for hybrids and other fuel efficient, low emissions vehicles. States, and even some cities have added their own incentives, up to a high of

$6,000 in Colorado in 2016, and the federal program was renewed in 2010 for plug-in hybrids and fully electric vehicles, with a gradual phaseout period per manufacturer after the sales quarter of its 200,000th credit-eligible vehicle. In turn, Canada introduced its Vehicle Efficiency Initiative in 2007, offering rebates of up to C$2,000—as well taxes of C$4,000 for a few of the most inefficient vehicles—and the most populous provinces (Ontario, Quebec, and British Columbia) also enacted their own programs, with incentives reaching as high as C$14,000 in Ontario for electric vehicles in 2017. Trucks have notably been exempt from Canada's federal inefficiency tax, which is believed to have induced some Canadians to substitute towards them and away from large SUVs. Since the extensive margin for these incentives occurred both before the period of analysis and around the same time in the two countries, any differences in how their respective auto industries may have been affected would have hopefully been largely set by 2009, such that whichever subsidies prevailed between then and 2015 would not induce any significant biases. Measuring their precise impact on alternative energy vehicle sales is an endeavor beyond the scope of this paper, but Section 1.5.3 provides suggestive evidence that these incentives have had very little differential effect on the market shares of American hybrid and electric vehicles versus that of their Canadian counterparts.

Moreover, the US enacted the Car Allowance Rebate System ("Cash for Clunkers") in the summer of 2009, and while the recession officially ended in the second quarter of that year in both the US and Canada, the latter may have recovered somewhat more quickly.[8] Analysis will therefore be performed both including and excluding the year 2009 from the sample. Finally, In 2015—the last year of the data period—Natural Resources Canada revised its fuel economy testing procedure to mirror the five-cycle test introduced by the EPA in 2008, producing perhaps the most serious threat to my identification strategy, and so results will also be presented omitting

[8]Figure A8 shows yearly total vehicle sales in the neighboring countries.

this year. The econometric specification I use in the first two results subsections is as follows:

$$y_{ist} = \gamma_{st} + \alpha_s D_i + \delta_s T_{ist} + \varepsilon_{ist} \qquad \text{(DID)}$$

where i are the two countries, s indexes over segments, and t over months of sample, γ_{st} are segment-specific time effects, and D_i is an indicator for the US. The dependent variable is either each segment's monthly share of total sales in its country or the country-segment-month's sales-weighted average new vehicle fuel consumption, and T_{ist} is a US treatment intensity index measuring the approximate fraction of vehicles sold in each segment-month which feature the revised EPA labels.

Additionally, the average monthly prices of regular unleaded gasoline in the US and Canada will be included in analogous models to account for their potential effects on market shares and segment-average fuel consumption. My preferred specifications, however, will be those estimated on the 2010-2014 trimmed samples that do *not* control for gas prices. This is in part because it is not necessarily clear how to allow Canada's gas price to influence its market for new vehicles relative to the effects of the US price in its own market. Some issues to consider, for example, include whether to use continuously-updated or period-averaged exchange rates, real or nominal prices, or whether to rely on growth rates, and if so, this raises yet another question about when prices should be normalized to each other. In light of the evidence suggesting that many consumers do not explicitly track or estimate their fuel expenditures, I will view the effect of fluctuating gas prices on automobile purchasing decisions not only as a real cost of ownership factor, but also as a behavioral parameter. As such, I will present results based on differences in the growth rates of nominal prices, and normalize them at the month which minimizes the sum of squared differences in the levels of nominal prices over the 2010 to 2014 period, October 2011 (see Figure A9). But it is important to note that results will differ very little by whether or not gas prices are controlled for regardless of which of version of the prices

I use. While differences in covariate levels between treatment and control groups matter when constraining a covariate to have the same effect in both groups, the Canadian time series are all highly correlated with their American counterparts irrespective of the assumptions under which they are respectively generated, and any normalization absorbs the majority of the difference in levels. Finally, allowing the price of gas to affect segment shares separately in the US and Canada sidesteps the normalization issue and point estimates again look very similar, but doing this comes at the expense of power: with only one country in treatment and control group each, estimation of the additional parameters reduces precision for most of the others in the model.

1.5 Results

1.5.1 Market Shares

I first explore the impacts of the labels on segment market shares. Table 1.1, provides some preliminary analysis by estimating the difference-in-differences coefficients δ_s on a broader definition of vehicle segmentation that aggregates all cars and all SUVs into two overall segments. We can immediately see that trimming the first and last years from the sample has a significant effect on the estimates. Unless otherwise specified, all parameter interpretation is henceforth done in terms of the third and fourth columns of the results tables in the text because of the identification issues outlined in the previous section. The SUV coefficient exhibits the best evidence of an effect, suggesting that overall SUV shares decreased by over 1.5 percentage points, and the corresponding increases in shares appear to be split between cars and trucks. Coefficients from flexible difference-in-differences specifications allowing each year to have its own intercept in segment market share are plotted below in panels (a), (b), (c), and (d) of Figure 1.4 respectively for small cars, small and large SUVs, and trucks[9] both to evaluate the parallel trends assumptions, and to examine any possible heterogeneity across segments in market share response timing. 2011

[9] See Figure A12 for other segments.

is the omitted year, and note that 2012 is always only partially treated and with different intensities depending on the segment in the sense that the earliest months of the year generally see no 2013 labels, the middle months begin having a few new model-year vehicles in most segments, and are therefore partially treated, while the final months experience the majority of the introduction of new model-years, and so approach being fully treated. Figure A13 depicts the distribution of the introduction of the new model-year 2013 vehicles and their accompanying labels across time.

Table 1.1: Gross segmentation market share difference-in-differences

	Full Sample	w/ Gas Price	2010-2014	w/ Gas Price
US × Car × Treat	1.70**	2.07***	0.75	0.81
	(2.43)	(3.01)	(0.95)	(0.99)
US × SUV × Treat	-2.09***	-2.27***	-1.73***	-1.81***
	(-4.25)	(-4.50)	(-2.99)	(-3.44)
US × Truck × Treat	0.38	0.11	0.86**	0.98*
	(0.96)	(0.23)	(2.32)	(1.94)
US × Van × Treat	0.02	0.08	0.10	0.04
	(0.08)	(0.22)	(0.44)	(0.17)
Observations	672	672	480	480
Within-R^2	0.82	0.83	0.87	0.87
SE Clusters	42	42	30	30

Notes: t statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Coefficient estimates in percentage points. Standard errors are clustered in country by 4-month blocks to allow for both within and across-segment serial correlation in errors within country-trimester. The effect of gas prices on segment shares is controlled for in the 2nd and 4th columns.

The parallel trends assumption not only holds for all four segments considered in the text, but for all other segments as well. Small cars begin to exhibit increases in share of over 1 percentage point by 2013, and eventually a much larger increase by 2015, confirming that it is prudent to focus on results omitting that year. Meanwhile small and large SUVs both experience statistically significant reductions in shares of around 1 and 0.5 percentage points respectively by 2012, which persist until 2015. Lastly, trucks do not show much evidence of an increase in shares following the introduction of the new labels, but rather that they experienced a relative dip in 2010, *before*

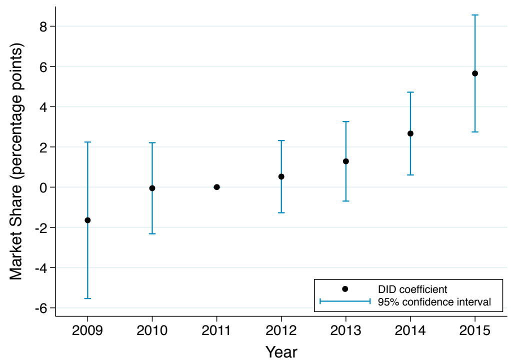

(a) Small cars

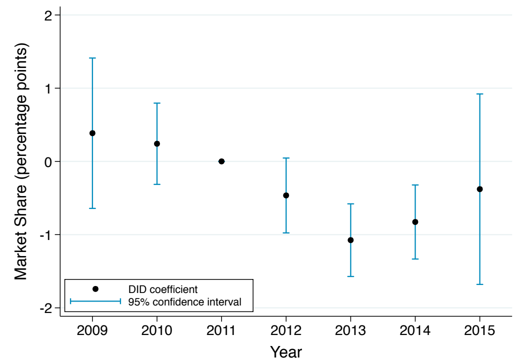

(b) Small SUVs

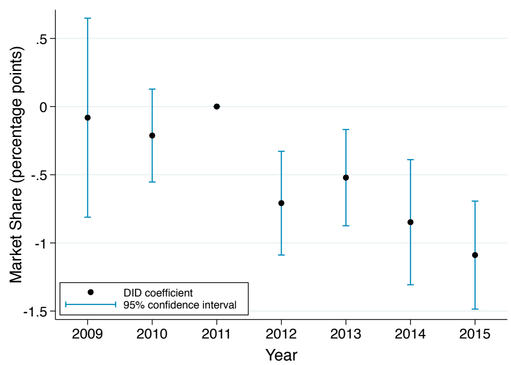

(c) Large SUVs

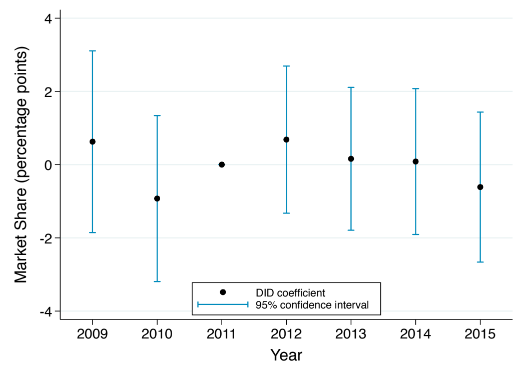

(d) Trucks

Figure 1.4: Market share flexible difference-in-differences

the update. Table 1.2 now returns to the treatment intensity difference-in-differences specification given in the previous section and presents estimates of the δ_s coefficients for all car and SUV segments.[10] The small cars estimate confirms that their share increased by nearly 1.5 percentage points due to the labels, small SUV shares decreased by about 1 percentage point, and large SUV shares dropped by a half percentage point. Other statistically significant coefficients suggest that large and luxury cars may have lost 0.3 percentage points of market share each because of the labels, while luxury SUVs may have lost as much as 0.4 percentage points of share, but statistical significance for these findings only appears when controlling for gasoline prices. The reported measure of fit is the R^2 within the segment-specific monthly mean shares (γ_{st}) because while

[10] The truck and van coefficients and their associated standard errors are mechanically identical to those from Table 1.1.

these alone allow for the model to fit very well, its performance in terms of the US-Canada share differential and the treatment effect is of particular interest.

Given the evidence that parallel trends held for all segments, including in 2009, hypothetical differential effects of the recession on the US and Canada seem to be less of a source of potential bias than the latter's model-year 2015 revision of its fuel economy testing procedure. Table A1 reports results including 2009 and omitting only 2015 from the analysis, and finds that the effect of increased truck shares is both more than halved, and no longer statistically significant. Examining the evolution of Canadian and American monthly truck shares[11] reveals that their difference is especially volatile from early 2009 to mid 2010, with a global trough in August 2009—the most active of the two months of the Cash for Clunkers program—surounded by global peaks in the 2009 months immediately preceding and following it, as well as a series of local lows in the first half of 2010. Omitting the year 2009 thus resulted in artificially low US truck shares in the pre-2013 label period by ignoring the intratemporal substitution that is most likely occuring as a result of Cash for Clunkers, and the estimates from Table A1 confirm that this was driving at least half of the magnitude of the trucks estimates reported in the last two columns of Table 1.2. Nevertheless, in the spirit of providing the most conservative point estimate for the effect of the change in labeling regimes, the calculation of yearly fuel savings undertaken in the conclusion will assume the full relative increase in US truck shares found in the trimmed sample column of the table from the text.

The nature of my identification strategy allows me to estimate by how much segment shares grew or fell after all sorting, but not to empirically evaluate exactly how buyers are sorting into new segments because of the labels. Many paths could be consistent with the results presented in this section, but I will briefly describe what I believe to be the most likely one. Some would-be SUV buyers (midsized, large, and luxury) could have instead opted to purchase trucks. Other potential large SUV buyers switched to midsized SUVs, and some potential midsized SUV buyers

[11] See Figure A10.

Table 1.2: Market share difference-in-differences

	Full Sample	w/ Gas Price	2010-2014	w/ Gas Price
US × Small × Treat	3.07***	3.19***	1.44**	1.54**
	(4.11)	(4.01)	(2.33)	(2.33)
US × Midsize × Treat	-0.59*	-0.35	-0.42	-0.25
	(-1.73)	(-1.16)	(-1.01)	(-0.72)
US × Large × Treat	-0.45***	-0.50***	-0.19	-0.32**
	(-3.16)	(-3.56)	(-1.35)	(-2.27)
US × Luxury × Treat	-0.55***	-0.52**	-0.20	-0.31**
	(-2.86)	(-2.40)	(-1.51)	(-2.47)
US × Sport × Treat	0.16	0.18	0.08	0.07
	(1.33)	(1.33)	(0.62)	(0.49)
US × Small SUV × Treat	-0.89***	-1.11***	-1.00***	-1.16***
	(-4.81)	(-6.21)	(-5.72)	(-6.76)
US × Midsize SUV × Treat	-0.11	-0.06	-0.00	0.16
	(-0.33)	(-0.17)	(-0.00)	(0.42)
US × Large SUV × Treat	-0.64***	-0.64***	-0.50***	-0.51***
	(-4.49)	(-3.90)	(-3.45)	(-3.35)
US × Luxury SUV × Treat	-0.45***	-0.50***	-0.28	-0.38**
	(-2.98)	(-3.22)	(-1.48)	(-2.10)
US × Truck × Treat	0.38	0.11	0.86**	0.98*
	(0.96)	(0.23)	(2.31)	(1.94)
US × Van × Treat	0.02	0.08	0.10	0.04
	(0.08)	(0.22)	(0.44)	(0.17)
Observations	1848	1848	1320	1320
Within-R^2	0.96	0.96	0.97	0.97
SE Clusters	42	42	30	30

Notes: t statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Coefficient estimates in percentage points. Standard errors are clustered in country by 4-month blocks to allow for both within and across-segment serial correlation in errors within country-trimester. The effect of gas prices on segment shares is controlled for in the 2nd and 4th columns.

instead went with small SUVs, decreasing the shares of large SUVs while holding level that of midsized SUVs as the segment both gained and lost buyers. Finally, people who may have chosen small SUVs were probably the ones induced to switch to small cars instead after being nudged by the labels into realizing that the former are actually significantly more inefficient than the latter.

1.5.2 Fuel Consumption

I now turn to the results on fuel economy itself, beginning as in the previous section with coefficients from fully flexible difference-in-differences models plotted in panels (a), (b), (c), and (d) of Figure 1.5 for small cars, small and large SUVs, and trucks respectively.[12] Evidence about the parallel trends assumption is now mixed as I find that it holds in some segments, notably in midsized and large cars, and small and large SUVs, but not in others (small cars, luxury cars and SUVs, and trucks). This means that the segment-average fuel consumption estimates will only be able to be interpreted causally for the former group of segments, and all others should be treated with more caution.

The δ_s coefficients from the consumption (DID) specification are reported in Table 1.3 with estimates scaled to gallons consumed per 100 miles (gals/100 miles), and they indicate that the labels resulted in increases of 0.05 gals/100 miles, and 0.06 to 0.08 gals/100 miles for midsized cars and large SUVs respectively, whereas fuel consumption decreased by 0.05 gals/100 miles for large cars, and dropped as much as 0.25 gals/100 miles in small SUVs. Since midsized cars are quite efficient (on average only being bested by small cars—see Figure 1.1), a modest decline in efficiency is actually consistent with the mpg illusion prediction that fuel economy may be overvalued for some of the more efficient vehicles. The large and highly significant coefficient on small SUVs suggests that the labels succeeded in pushing buyers towards the more efficient models. Again, this is consistent with the illusion proposition that fuel economy is undervalued among relatively inefficient vehicles as small SUVs reach a sales-weighted average well below

[12] Estimates for other segments can be found in Figure A14.

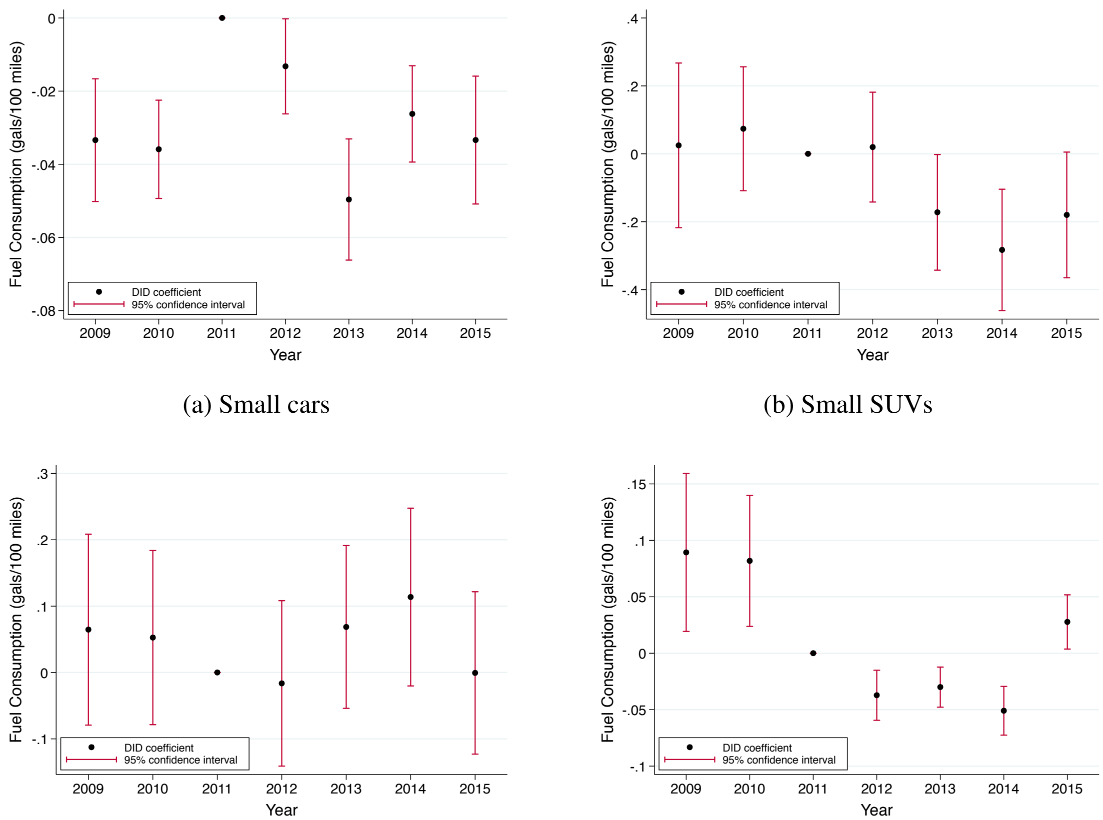

(a) Small cars

(b) Small SUVs

(c) Large SUVs

(d) Trucks

Figure 1.5: Fuel consumption flexible difference-in-differences

25 mpg in every year (Figure 1.2).

The only coefficient which, at first glance, may seem inconsistent with fuel economy being undervalued at low levels is that on large SUVs, which implies that they became nearly 0.1 gallon per 100 miles more inefficient because of the labels. Recalling the result from the previous section, however, that large SUVs lost about 0.5 percentage points of market share, this increase in average consumption is actually not surprising because people being induced to switch to more fuel efficient vehicles by the labels cared more about fuel economy than those who did not switch. They would have therefore been likely to choose some of the more fuel efficient large SUVs had they not instead purchased from another segment[13]—thus worsening observed

[13]Which almost necessarily has better fuel economy on average because large SUVs, being among the heaviest vehicles, are close to the most inefficient (second only to trucks, and about level with vans).

Table 1.3: Fuel consumption difference-in-differences

	Full Sample	w/ Gas Price	2010-2014	w/ Gas Price
US × Small × Treat	-0.01*	-0.01	-0.02*	-0.01
	(-1.71)	(-1.43)	(-2.03)	(-1.53)
US × Midsize × Treat	0.07***	0.06***	0.05**	0.05**
	(3.82)	(3.37)	(2.36)	(2.15)
US × Large × Treat	-0.05***	-0.05***	-0.05**	-0.05**
	(-3.27)	(-3.26)	(-2.07)	(-2.38)
US × Luxury × Treat	-0.03*	-0.04***	-0.05***	-0.07***
	(-1.97)	(-2.74)	(-3.26)	(-3.97)
US × Sport × Treat	-0.13*	-0.18***	-0.01	-0.04
	(-1.83)	(-2.71)	(-0.27)	(-0.69)
US × Small SUV × Treat	-0.23***	-0.24***	-0.25***	-0.26***
	(-6.13)	(-5.45)	(-5.05)	(-4.44)
US × Midsize SUV × Treat	-0.00	-0.00	0.01	0.01
	(-0.42)	(-0.38)	(1.04)	(0.64)
US × Large SUV × Treat	0.04	0.03	0.08**	0.06**
	(1.29)	(1.12)	(2.26)	(2.28)
US × Luxury SUV × Treat	0.04**	0.05***	0.05*	0.07***
	(2.55)	(2.99)	(1.97)	(3.73)
US × Truck × Treat	-0.06***	-0.09***	-0.07***	-0.08***
	(-3.12)	(-6.09)	(-3.35)	(-3.75)
US × Van × Treat	-0.09**	-0.06**	-0.01	-0.01
	(-2.38)	(-2.22)	(-0.62)	(-0.81)
Observations	1848	1848	1320	1320
Within-R^2	0.67	0.73	0.70	0.72
SE Clusters	42	42	30	30

Notes: t statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Coefficients in gallons per 100 miles. Standard errors are clustered in country by 4-month blocks to allow for both within and across-segment serial correlation in errors within country-trimester. The effect of gas prices on segment-average fuel consumption is controlled for in the 2nd and 4th columns and is found to be negative for most segments, or statistically indistinguishable from 0 otherwise.

large SUV average fuel consumption by changing segments. Interpreting the results relative to baseline utilizations, they correspond to around 1% increases in consumption each for midsized cars and large SUVs, and respectively 1 and 6% reductions for large cars and small SUVs. While statistically significant, the labels' impacts on large car and SUV fuel utilizations will actually contribute very little to the changes in economy-wide consumption I estimate in Section 1.6. This is both because these effects are quite small relative to utilizations, and because between 2012 and 2015, large cars and SUVs only made up about 2 and 5% of US vehicle sales respectively.

Of the segments for which the parallel trends assumption failed, all appear to have experienced small decreases in average fuel utilization post treatment, except for luxury SUVs whose consumption seems to have gone up slightly. Turning to Figure A14 in the Appendix, however, reveals that this increase occurred almost entirely in 2011, when no 2013 model-year luxury SUVs could have been sold—by law, new model-years can be introduced as early as January of the previous year, but no earlier—and the change can therefore not be attributed to the 2013 fuel economy labels. Finally, note that as in the shares specifications, the standard errors estimated in Table 1.3 are clustered by country by 4 months periods, but there is less reason to worry about the errors in the fuel consumption models being correlated across segments, and this fact can be used to extend the clustering's time dimension in order to allow for the potential serial correlation in the errors to last longer. This issue is explored in Appendix Table A2, which finds little evidence of changes in parameter statistical significance after expanding the clusters to last two years.

1.5.3 Hybrid and Electric Vehicle Incentives

In this section, I address the potentially confounding impacts of the different hybrid and electric vehicle subsidies introduced by US states and Canadian provinces over the sample period.

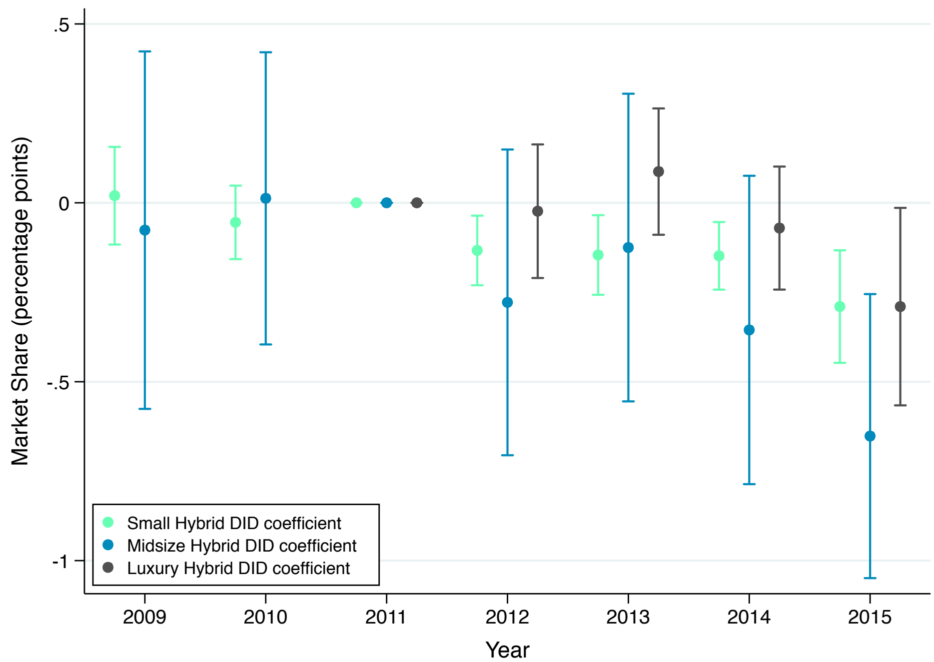

Figure 1.6: Flexible difference-in-differences: hybrid shares by segment

All results presented here are restricted to non-sports-car[14] models whose only engine offerings are hybrid or electric variants, which represent between 70 and 72% of all US hybrid and electric sales across the 7 years in the sample. These consist of the Honda Insight, Toyota Prius c, and Fiat 500e[15] for small cars, the Prius, Ford C-Max, and Nissan Leaf midsized cars, and the Chevrolet Volt, Lexus CT, and Tesla Model S luxury cars.[16] The restriction to these vehicles is necessary because for models with both gasoline-only and hybrid engine options, sales of the hybrid variants can only be inferred from the US engines production data, which would bias estimates of US-Canada differences in hybrid shares towards the null. As a first pass, Figure 1.6 plots yearly flexible difference-in-difference estimates of the sums of the hybrid models' market shares by segment. There are no discernable differences in market shares between 2009, 2010, and 2011 for

[14]Some sports cars like the Porsche 918 Spyder only feature a hybrid engine, but are prohibitively expensive performance-oriented vehicles with accordingly low sales, and little in common with mainstream hybrids.

[15]Ward's reports 500e sales independently from those of its Fiat 500 internal combustion engine cousin.

[16]The Fiat 500e, Nissan Leaf, and Tesla Model S are fully electric vehicles, but I sometimes use the term "hybrid" to refer to both hybrid and electric cars.

small and midsized hybrids, and no luxury hybrid-engine-only models in 2009 or 2010. Relative decreases in shares begin to appear by 2012 for the former two segments, but no discernable pattern emerges for the latter. Excluding 2015 from interpretation as before, the differences that do emerge relative to 2011 are very small, topping out at (a non-statistically-significant) -0.3 percentage points for midsized hybrids in 2014, hinting that the magnitude of changes in hybrid shares is likely of second order relative to the estimates from the previous sections. In order to more precisely bound the possible effects of incentives on shares, I now consider two types of difference-in-differences specifications.

The first is a placebo-style analysis where I separately focus on trimmed samples on both sides of the 2013 label introduction. I allow a possible treatment to begin—and persist for the remainder of the trimmed periods—at every month of 2010 on the pre-2013 label sample, and in the five months between September 2013 and January 2014 for the post label sample. The coefficients estimated from this exercise are plotted in the left and right panels of Figures 1.7 and 1.8 below for small and midsized hybrids respectively.[17] The reason why the left panels estimate more than twice as many coefficients as the right ones is purely empirical: we've seen plenty of evidence suggesting that Canada's revised testing methods have had significant impacts on its auto industry and so the year 2015 was dropped from the post sample, leaving me with less data than in the pre sample—which retained 2009 as parallel trends in market shares have held consistently.[18] The range of estimates indicates that any differences in tax incentives between the US and Canada can only account for a magnitude of about 0.07 percentage points for the sum of small hybrid shares, and at most 0.2 for midsized hybrids, albeit rarely with statistical significance for either segment. Note that the latter effect's direction is reversed across the two samples, which is by no means surprising as it is entirely plausible that on aggregate, new hybrid incentives were relatively more generous in the US until around 2012 before gradually falling behind the latest

[17] See Figure A15 in the Appendix for luxury hybrids.

[18] Specifically, I exclude every month after partial treatment begins in March 2012 on the left side; on the right side, every month before most labels have been introduced in September 2012, as well the year 2015 are dropped.

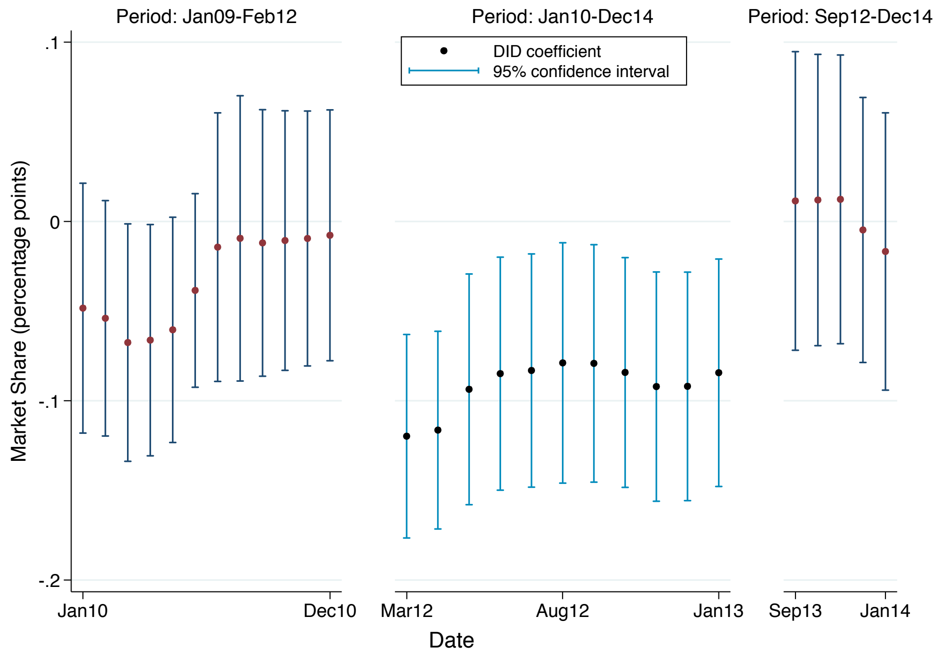

Figure 1.7: Small hybrid incentives difference-in-differences

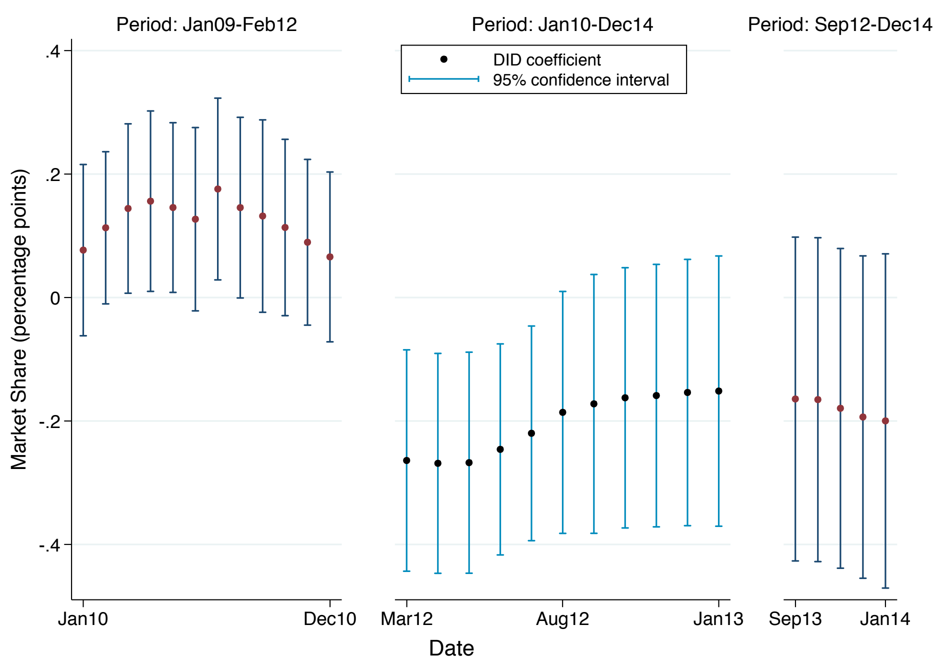

Figure 1.8: Midsize hybrid incentives difference-in-differences

Canadian subsidies sometime after that. The post sample placebo estimates showed hybrid shares from all three segments decreasing by an additional 0.1 percentage points had 2015 been included, but this bears no relevance to any of the results as none of the fuel savings calculations in the concluding section rely on specifications estimated on data that includes this year.

Armed with an idea of the magnitude of the effect of potentially differential incentives, we can now turn to the figures' middle panels: these plot coefficients from the second type of specification in which the sample restriction returns to the usual 2010 to 2014 range, providing estimates of the combined effects on segment summed hybrid shares of the 2013 labels *and* differences in the two countries' hybrid tax policies. Each of the 11 estimates is again based on an equation with a different treatment starting date corresponding to every month between March 2012 and January 2013,[19] and we can conclude that relative to their Canadian counterparts, US hybrids may have experienced small decreases in share, of about 0.08 to 0.12 percentage points and 0.15 to 0.27 percentage points for small and midsized hybrids respectively (while luxury hybrid shares remained even).

In summary, even if we believed that the entirety of the changes in hybrid and electric shares were attributable to differential incentives for these vehicles, the evidence presented in this subsection has shown that the bias this would induce on the market share estimates documented in Table 1.2 would likely be negligible. For small cars, the more consistently negative coefficients would even imply that their share could have increased slightly more without such differences. Furthemore, while Table 1.3 showed that midsized car fuel consumption increased by a statistically significant 0.05 miles per 100 gallons, the decline in midsized hybrids' shares that could be attributed to tax incentives can only account for around 5 to 10% of this rise—or have attenuated it by a similar amount if relative midsized hybrid shares instead rose by 0.2 percentage points, as the left and right panels of Figure 1.8 together proved inconclusive about the direction of the incentives' effect.

[19]See Figure A13 for the 2013 model-year diffusion rate informing my choice of treatment starting months.

1.5.4 An Alternative Modeling Detour

The wealth of vehicle characteristics available in the Ward's data enables me to experiment with some more parametric specifications that could theoretically allow for a more precise estimation of post-labels changes in within-segment fuel consumption preferences after controlling for other vehicle attributes. One of the simplest such specifications that can be looked at is a linear probability model of the form:

$$Share_{jt} = \alpha_s gpm_{jt} + X_{jt}'\beta_s + \tau_{st} + \varepsilon_{jt} \tag{Base}$$

$$+ \delta_s T_{jt} \times gpm_{jt} \tag{Full}$$

where j indexes over individual models of vehicles, T_{jt} is now a model-specific label introduction indicator, X_{jt} is a vector of control characteristics including manufacturer suggested retail price, the horsepower to curbweight ratio, and a model availability measure constructed from the data, and τ_{st} are segment-specific time effects. The time effects allow each segment by month of sample to have its own intercept, enabling mean segment shares to vary over time with any relevant factors (observed or not) like gasoline prices, the introduction of competing models, the impact of the labels on segment shares themselves, etc. All coefficients are segment-subscripted and thus interacted by segment indicators, allowing not only the control variables, but also both the baseline and post-treatment components of fuel economy to affect different segments' market shares differently. This is the critical aspect of this specification. First, it enables the identification of changes in average fuel economy valuation in every segment from the introduction of each vehicle-specific 2013 label. Second, it allows each segment to have its own valuation for all of the observed characteristics in order to reflect the heterogeneity of consumer preferences across segments; we might expect, for example, that increases in fuel economy or horsepower are valued asymmetrically in purchases of small versus those of sports cars, or that relative prices predict midsized and luxury car shares differently.

The equations are estimated separately on the US and Canadian data in order to consider an alternative against which to compare any changes in the evolution of US fuel economy preferences; note that this implies some parallel trends assumptions also need to be satisfied in this context. Estimates of the α_s and δ_s coefficients from the (Base) and (Full) specifications are reported in the first four columns of Appendix Table A3, but these results of course suffer from the well-known endogeneity issues addressed in different ways by the literature. One popular method used to tackle this endogeneity, vehicle fixed effects, could work well if vehicle fuel economy relative to monthly segment-means varied significantly across model-years (within models). But the baseline consumption statistical zeros in some of the largest segments by sales (small and midsized cars, and trucks) after their inclusion in the next four columns of Table A3 do not bode well. The zeros imply that a substantial amount of the predictive variation in fuel efficiency comes from its valuation across vehicle models, and too much variation is absorbed by these fixed effects to consider the strategy any further.[20]

Another possible approach introduced by Berry et al. (1995) consists in the two-stage least squares estimation of the following specification implied by a nested logit model of consumer choices (with "BLP" instruments as described further below):

$$ln(share_{jt}) - ln(share_{0t}) = \alpha_s gpm_{jt} + X'_{jt}\beta_s + \sigma ln\left(\frac{share_{jt}}{share_{st}}\right) + \varepsilon_{jt} \qquad \text{(Nested logit)}$$

where $share_{0t}$ is the market share of the outside good (used vehicles, motorcycles, etc.), $\frac{share_{jt}}{share_{st}}$ is the within-segment market share of model j, and σ is the within-segment correlation in preferences which relaxes (across, but not within segments) the independence of irrelevant alternatives assumption imposed by the standard logit model. I do not observe outside good sales, however, and I therefore modify this specification by including a month of sample fixed effect,

[20]Employing different functional forms, like taking the logs of variables, does not help alleviate the issue.

which will not only absorb this variable, but also allow for mean shares to vary over time:

$$ln(share_{jt}) = \tau_t + \alpha_s gpm_{jt} + X'_{jt}\beta_s + \sigma ln\left(\frac{share_{jt}}{share_{st}}\right) + \varepsilon_{jt} \qquad \text{(NL base)}$$

$$+ \delta_s T_{jt} \times gpm_{jt} \qquad \text{(NL full)}$$

In these models, the endogeneity will be explicitly addressed by instrumenting for fuel consumption, price, the horsepower to weight ratio, and within-segment market share with the segment by month of sample and make by month of sample means of all these variables but the latter; each vehicle's own characteristics are always excluded from the constructions of these means. Note that segment-specific time effects cannot be added as was done in the linear probability specifications because their inclusion would prevent the identification of the within-segment correlation parameter σ, reimposing the independence of irrelevant alternatives across segments—an assumption which is difficult to defend. Results are reported in Table A4, where the first four columns correspond to the (NL base) and (NL full) equations estimated first on the US, and then on Canada, and the last four substitute the month of sample fixed effects from these models with a more aggressive set of make by month of sample fixed effects which enable different companies to gain or lose average market share over time. These results are even more difficult to take seriously than those from the linear probability models: not only are most estimates quite sensitive to the choice of fixed effects, but many base coefficients are statistically positive despite the instrumental strategy, which would imply *negative* valuations of fuel economy if we actually believed that the endogeneity had properly been addressed.

In short, the considerable sensitivity of the results discussed in this subsection to modeling assumptions seems to suggest that if we hope to successfully use parametric methods in market share specifications to tackle further fuel economy questions, we may need to focus on developing new tools. These would need to be better equipped to handle the many endogenous characteristics that are often observable to researchers, but which, using current parametric approaches, cannot be relied on in this context to improve estimates' precision as proposed by theory.

1.6 Conclusion and the Role of CAFE Standards

Recall that new vehicle fuel economy was steadily increasing throughout the Ward's data period, and this is in large part attributable to CAFE standards. Following the National Highway Traffic Safety Administration's 2006 attempts at reform, the 2007 Energy Independence and Security Act signed by President George W. Bush required new CAFE standards beginning in 2011, with an increased fuel economy requirement by 2020 to at least 35 mpg for all passenger and non-passenger vehicles, and increases to the maximum feasible average fuel economy standard for each model-year fleet between 2021 and 2030. Figure 1.9 provides the history of CAFE standards versus achieved efficiency (not pictured are 2012 forecasts of the standards through 2025, which reach 55.3 and 39.3 mpg for cars and all other light duty vehicles respectively).[21] The time series portray that while fuel economy has increased significantly over the past decade, periods of steadily improving fuel economy have mostly been accompanied by quickly ramping up CAFE standards.[22] The twenty year period between the mid 80s and the mid 2000s saw both stagnant CAFE standards and new vehicle fuel economy, but this is not to say that the efficiency of engines has not steadily progressed throughout this time. Fuel economy is inversely related to engine output and vehicle weight, yet Figure 1.10 shows that the former has seen horsepower booms of 65 and over 95% in cars and light trucks respectively between 1985 and 2005, and the latter more modest increases of 13 and 25% respectively, all while fuel economy itself improved by only 8 and 4% during the same period of flat standards. To the extent that the consumer

[21] All CAFE fuel economies in this section are not only CAFE credit-adjusted, but also based on pre-2008 model-year EPA testing revisions for higher accelerations and speeds, air conditioning use, and cold-engine driving in stop-and-go traffic, hence the sizable differences in levels with fuel economies based on Ward's data (e.g. in Figure A7).

[22] Figure A16 adds the demand side factors affecting realized fleet fuel economy, vehicle miles traveled (VMT) and gasoline prices.

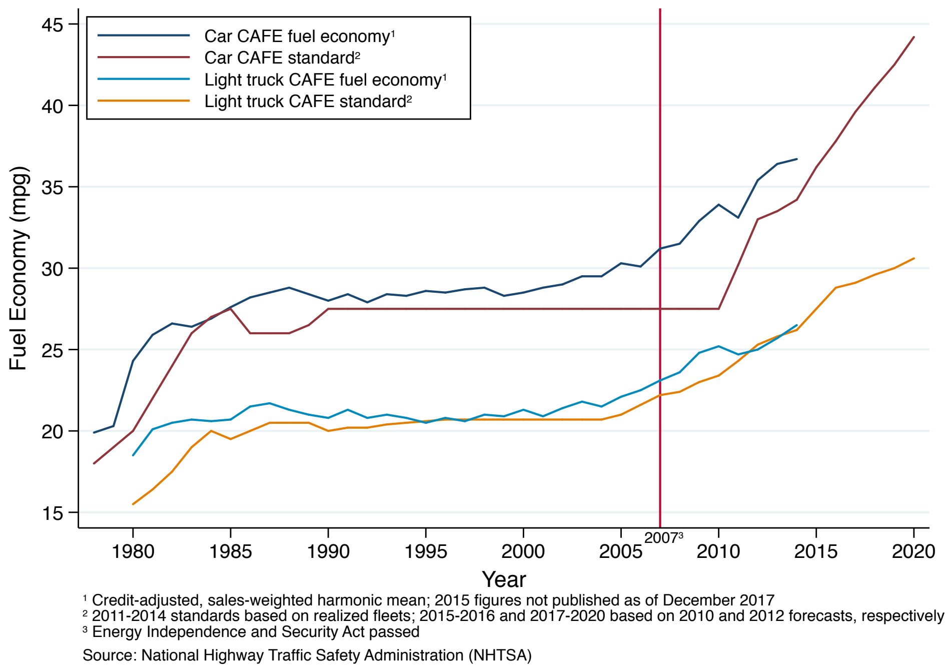

Figure 1.9: CAFE standards and fuel economy (1978-2020)

preferences[23] driving this allocation of technological progress persist today and that gas prices remain low, labels successfully conveying all the relevant information are not likely to be enough to generate sustained improvements in new fleet fuel economy.

I conclude by estimating the fuel savings engendered by the redesigned labels, relying on the statistically significant coefficients from my preferred specifications which are found in the third columns of Tables 1.2 and 1.3. Specifically, I consider that small car shares increased by 1.4 percentage points, small and large SUV shares fell by 1 and 0.5 percentage points respectively, truck shares increased by 0.8 percentage points,[24] large and luxury car shares each decreased by 0.2 percentage points, and luxury SUV shares dropped by 0.3 percentage points. Regarding the last three segments, the rationale is based on the evidence from most specifications that they experienced small but statistically significant reductions in shares. As for the consumption estimates, I use the ones from segments which respected the parallel trends assumption, namely

[23] And whatever they are perceived to be by manufacturers.
[24] See discussion about the sensitivity of the truck share estimates towards the end of Section 1.5.1.

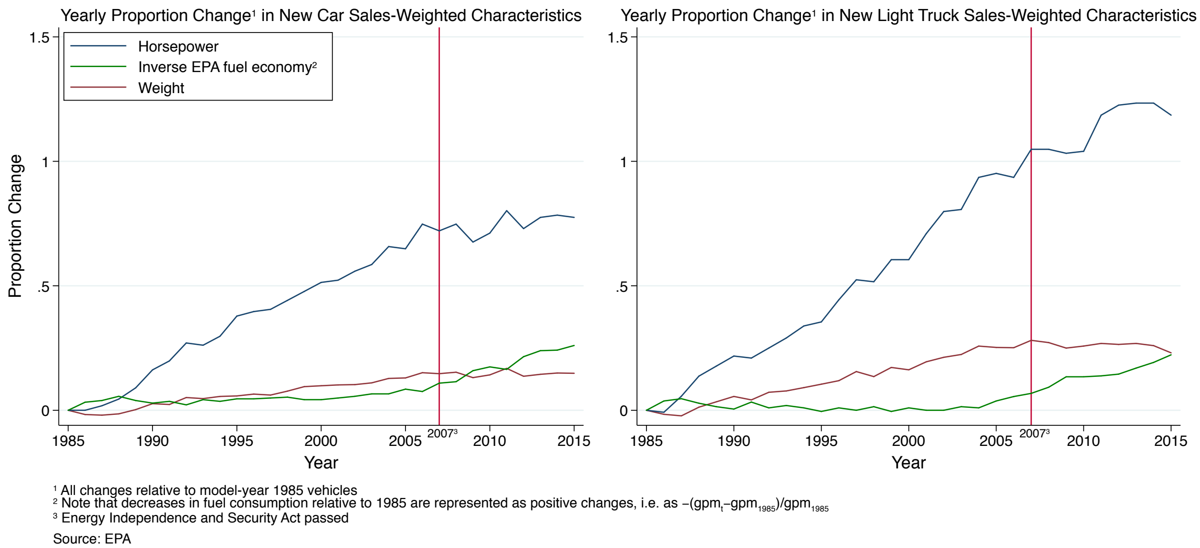

Figure 1.10: Efficiency gains allocation over time (EPA (2016))

that small SUV and large car fuel consumptions respectively improved by 0.25 and 0.05 gallons per 100 miles, while those of midsized cars and large SUVs worsened by 0.05 and 0.08 gallons per 100 miles.

An additional assumption about new vehicle miles traveled is required, and I rely on Department of Transportation estimates that American vehicles were driven an average of around 12 thousand miles a year between 2012 and 2015, imposing that new vehicles be used equivalently to the overall fleet average—an especially conservative assumption since they are in fact known to be driven more than older vehicles. Yearly fuel savings calculated under these assumptions, as well as the absence of a rebound effect, are produced in Figure 1.11 below. Because of the relatively large market share of midsized cars, the majority of the gains obtained by the improved small SUV fuel economy turn out to be offset by the much smaller loss in average midsized car efficiency, and foregone consumption thus almost entirely results from consumers switching across segments. Cumulative savings from the labels' gradual introduction in 2012 through 2015 add up to nearly 150 million gallons of gasoline, a quantity which, while an order of magnitude below the flow of fuel estimated to have been saved by CAFE standards (around a billion gallons

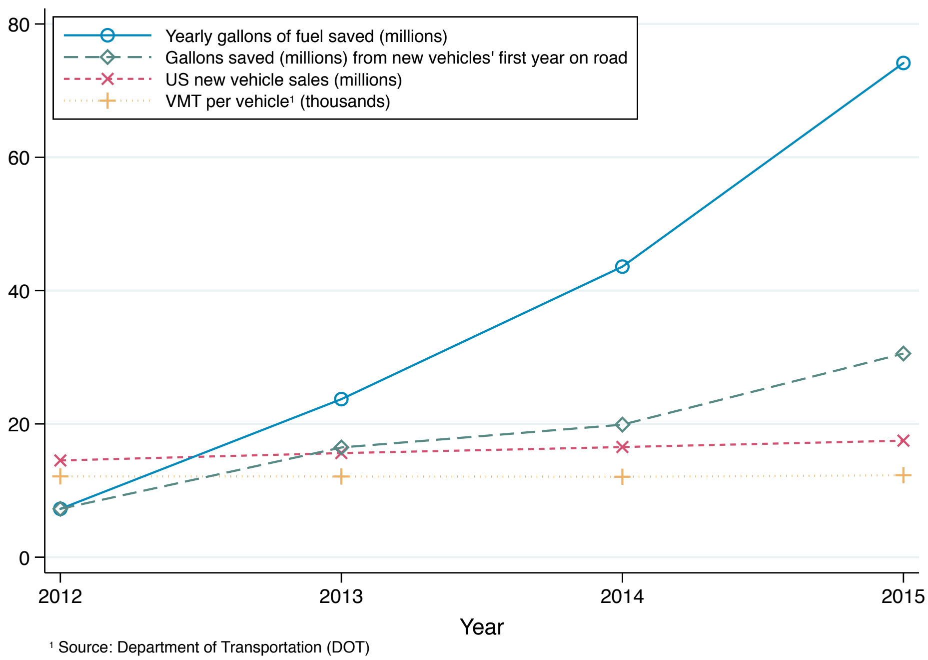

Figure 1.11: Estimated yearly gas consumption avoided by 2013 EPA labels

per year from each mpg increase in the standards[25]) is a pretty significant achievement for a virtually costless intervention. However, compared to the average of 391.4 million gallons of finished motor gasoline which were consumed daily in the United States in 2017, these savings are relatively small.[26]

The results presented in this paper have shown that the EPA's 2013 labels, redesigned to improve the delivery of relevant cost and environmental information, have helped produce a more fuel efficient fleet of new vehicles by affecting some consumers' purchasing decisions. But other tools likely need to continue being used to achieve further gains, and CAFE standards have been criticized by the literature for contributing to the growth of SUVs—these have always been subject to looser standards, which are now even further differentiated by vehicle "footprints", or wheelbase by track width. Most economists agree that raising gasoline taxes to correct pollution

[25] See for example Goldberg (1998), Kleit (2004), or Austin and Dinan (2005).
[26] Source: Energy Information Administration.

externalities is much closer than standards to being a first-best alternative,[27] and the literature should perhaps strive to more strongly emphasize that increased taxes can be made revenue neutral (via lump sum redistribution, for example) to help some politicians overcome their tax allergies.

[27] See van Benthem and Reynaert (2015): https://www.economist.com/blogs/freeexchange/2015/07/reducing-carbon-emissions.

Acknowledgments

This chapter, in part, is currently being prepared for submission for publication of the material. Panassié, Yann. The dissertation author was the sole investigator and author of this material.

Chapter 2

A Cautionary Tale on Estimating the Short Run Gasoline Price Elasticity of Demand for Driving

2.1 Introduction

A literature estimating the price elasticity of gasoline demand goes back to at least 1969, when Heien finds it to be -0.3 in the short run, and -0.7 in the long run in the United States. As data grew better and the econometric methods used to produce new estimates became more advanced, authors began testing more intricate hypotheses about its evolution and potential heterogeneities. Dahl (1982) was one of the first papers to look for evidence of differential gasoline demand elasticities, deriving a short-run estimate of around -0.2 across diverse countries. This figure was found to be stable between 1970 and 1978 across a wide range of incomes and gasoline prices,[1] and did not vary with the direction of price changes. More recently, Hughes et al. (2008) finds

[1] Including the price swings caused by the 1973 to 1974 oil embargo.

evidence of a significant shift in the short-run demand elasticity in the US, with estimates from monthly time series data between 1975 and 1980 ranging from -0.34 to -0.21, while the same models yield a much more inelastic -0.08 to -0.03 range from 2001 to 2006.

In a meta-analysis of the literature, Espey (1998) concludes that while the short-run elasticity appears to be declining over time, the long-run elasticity may have actually increased. The author notes that this result might at first seem contradictory, but posits that as gasoline prices rose during the 1970s and people made some initial adjustments in driving habits, there were fewer options for further short-run responses to price changes. However, as automobile fuel efficiency technology improved between the late 1970s and mid 1980s, long-run responses to price changes could now be achieved by purchasing more fuel-efficient vehicles. Espey observes that panel data tend to produce more elastic short-run estimates, but finds that within the US, data periodicity and geography do not have a unilateral effect on researchers' estimates. She infers that there are no significant differences between aggregate and per capita models. In contrast, Levin et al. (2017), relying on microdata from gas station sales, cautions that while monthly or state level aggregation seems to result in limited bias, country or yearly aggregation does produce significantly too inelastic estimates relative to what it finds based on the individual transactions.

The short-run gasoline demand elasticity can be approximated by a decomposition into a driving intensity component, and a within-household vehicle choice component for those with multiple vehicles to choose from.[2] Previous work has focused on driving demand as an entity of interest. Gillingham (2014), for example, uses a sample of vehicles registered in California between 2001 and 2003, and subsequently given a smog check between 2005 and 2009, and estimates a medium-run vehicle miles traveled (VMT) elasticity of -0.22. As part of an effort to calculate an optimal gasoline tax, Lin and Prince (2009) reports a California VMT elasticity of -0.07 in the short run between 1970 and 2007. These figures, however, are quite sensitive to sampling and modeling assumptions.

[2]Vehicle purchase choice is usually only considered relevant in the long run.

In this paper, we first seek to contribute to the literature by attempting to use shocks to California's supply of gasoline resulting from unplanned refinery shutdowns. Because California must sell reformulated gasoline different from that of most other states,[3] the idea is that the ensuing supply disruption can be used to identify a short-run gasoline demand elasticity. When this fails because of substitution towards other refineries' output, inventories, and imports of refined gasoline from other states and even countries,[4] we instead employ an econometric model isolating macroeconomic sources of shocks to VMT demand. Our results suffer from some of the same sensitivity to modeling choices as much of the previous literature, but we provide evidence that the short-run gasoline price elasticity of demand for VMT varies over time in both magnitude and volatility, possibly according to the level of gas prices. The remainder of the paper is organized as follows. In Section 2.2, we describe the data and present the identification strategy, Section 2.3 provides our models' estimates, and we conclude in Section 2.4.

2.2 Data and Econometric Framework

We rely on monthly data on vehicle miles of travel on California highways between 1996 and 2017 from Caltrans, monthly gas price data from the Energy Information Administration (EIA), and monthly California-level macroeconomic data on employment and wages from the Bureau of Labor Statistics. Figure 2.1 plots VMT and gas prices, and Figure 2.2 shows the relationship between VMT and the macroeconomic covariates in levels. Our strategy is to isolate the elasticity of demand by using the macroeconomic covariates to control for demand shocks, presumably constraining the remaining variation in prices to result from shocks to the supply of gasoline—and thus allowing the coefficient on price to capture the elasticity of demand. Because VMT exhibit a clear seasonal pattern from differential demand throughout the calendar year, our

[3] It must contain 10% ethanol as oxygenate.

[4] Critically, gasoline sales figures are known to be imprecise, contributing to the difficulty in measuring the magnitude of potential disruptions.

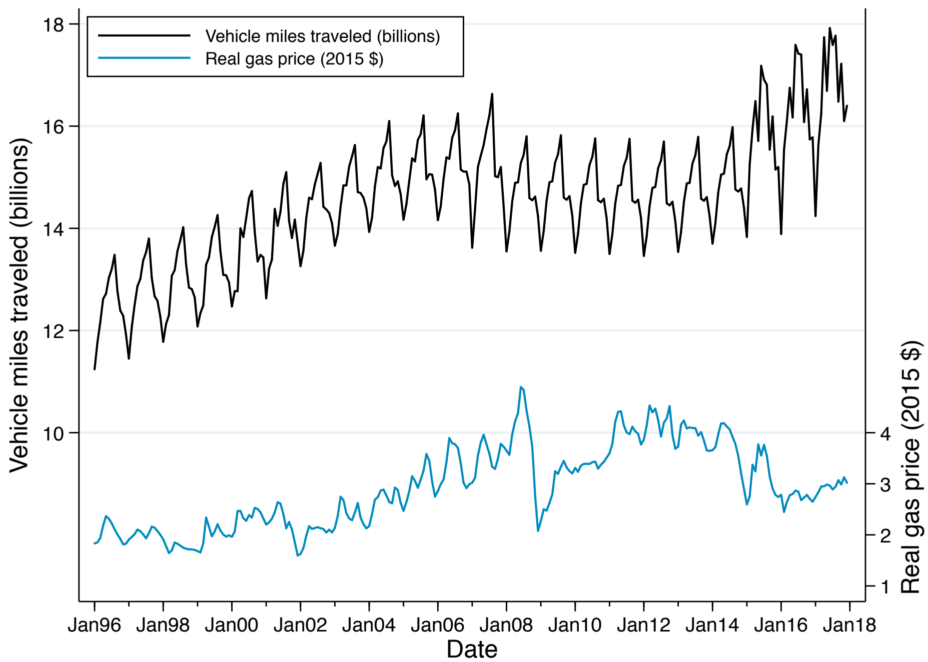

Figure 2.1: California VMT and real gasoline price. *Sources:* California Department of Transportation and Energy Information Administration

models include calendar-month fixed effects. We first estimate:

$$log(VMT_t) = \alpha_m + \eta log(Price_t) + \beta_1 log(Employed_t) + \beta_2 log(Unemployed_t) + \varepsilon_t \quad (2.1)$$

where VMT_t are vehicle miles traveled on California freeways and highways in month of sample t, $Price_t$ is the average price of gasoline in California in month t, $Employed_t$ and $Unemployed_t$ are respectively the employment and unemployment levels in the state in month t, α_m are the month fixed effects intended to capture the seasonal variation in demand for driving, and ε_t is the idiosyncratic error term. We then consider another set of models where we allow for the elasticity to vary over time (yearly, as indicated by the y subscript on the elasticity coefficient):

$$log(VMT_t) = \alpha_m + \eta_y log(Price_t) + \beta_1 log(Employed_t) + \beta_2 log(Unemployed_t) + \varepsilon_t \quad (2.2)$$

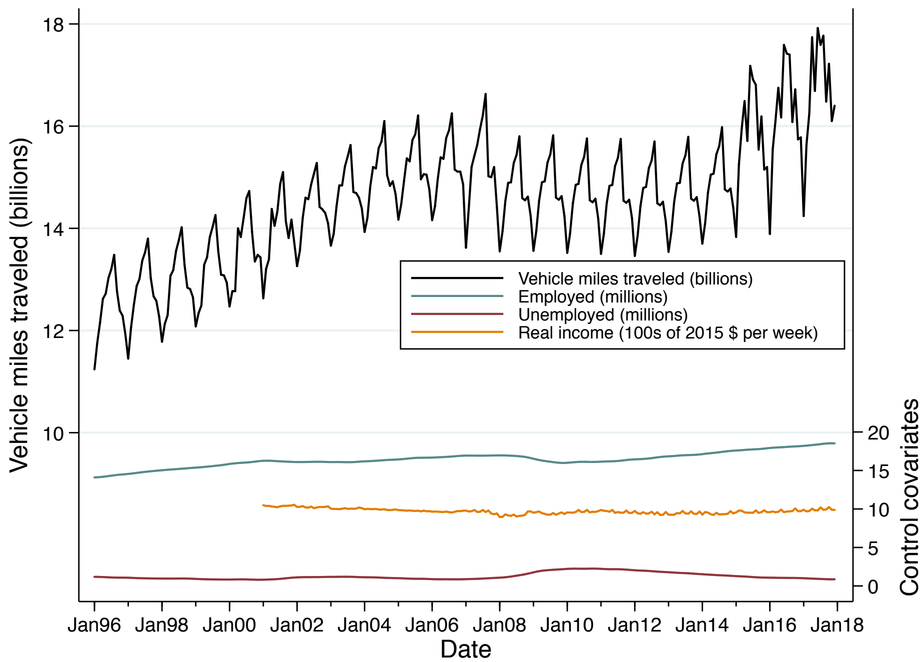

Figure 2.2: California VMT and macroeconomic covariates. *Sources:* California Department of Transportation and Bureau of Labor Statistics.

2.3 Results

Results from estimation of equation 2.1 on the full sample are reported in the first column of Table 2.1. The elasticity of demand estimate is -0.01. We also want to include average earnings as an additional control because of its potential to be another important source of demand shocks. However, since we only have access to average nonfarm real income in California between 2001 and 2017,[5] we first reestimate equation 2.1 on the subsample from these years to examine the possible difference introduced by omitting the first five years. Results are very comparable, as shown in the second column of Table 2.1 including an updated elasticity of -0.03. We then augment our model with the log of real average weekly earnings, whose little variation does not seem to contribute any information beyond that already contained in the other covariates.

[5]Data on total nonfarm labor force income between 2001 to 2006 are missing, but imputed based on this variable's relationship with incomes across a subset of industries (manufacturing, durable goods, nondurable goods, and film) for which data are available back to 2001. Between 2007 and 2017, the correlation these jointly exhibit with total nonfarm average earnings in a reduced form model is 0.96.

Table 2.1: Fixed elasticity model estimates

	Base model	2001-2017 period	Income control
log(Price)	-0.01	-0.03	-0.02
	(0.02)	(0.02)	(0.02)
log(Employed)	1.2***	0.9***	0.9***
	(0.06)	(0.1)	(0.1)
log(Unemployed)	0.03	-0.0003	0.00002
	(0.02)	(0.02)	(0.02)
log(Real income)			0.02
			(0.2)
Month FEs	Yes	Yes	Yes
Observations	264	204	204
R^2	0.89	0.84	0.84

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Dependent variable is log(VMT), and sixty-month lag Newey-West standard errors are reported in parentheses.

Estimates are reported in the third column, where we now observe an elasticity of -0.02. We account for heteroskedasticity and serial correlation throughout with Newey-West standard errors.[6] The elasticity coefficients are never statistically significantly different from 0, suggesting that these models are not capturing our coefficient of interest well enough.

Next, we estimate equation 2.2 and plot the resulting elasticity history along with 95% confidence intervals in Figure 2.3, while Figure 2.4 depicts the coefficients yielded by the model augmented with the income measure. These series indicate that the elasticity may indeed be varying over time, but the wide confidence intervals do not allow us to precisely test any of the specific hypotheses about its evolution we were initially interested in evaluating. We finally compare the residuals from the fixed and yearly-varying elasticity models in Figures 2.5 and 2.6 for the models without and with average earnings, respectively. The residuals from the equations

[6] With an aggressive lag length set to sixty months, although we find that there is almost no difference in standard error estimates when increasing lag length beyond twelve months.

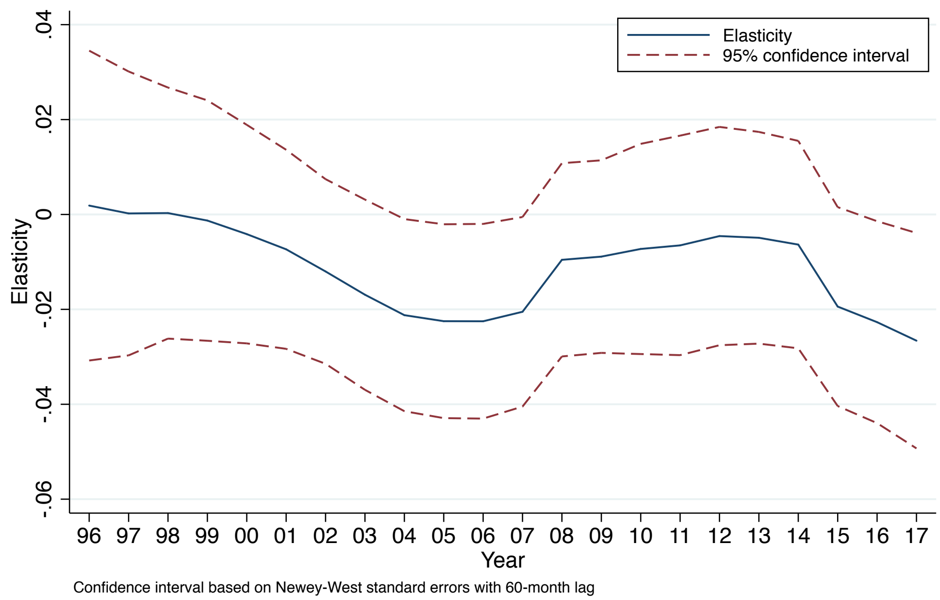

Figure 2.3: Yearly VMT demand elasticity estimates

where we allow the elasticity to vary yearly feature much less autocorrelation, highlighting the importance of accounting for the elasticity's dynamics when attempting to estimate it—or that some important variables are simply missing from the equation. Furthermore, these residuals series exhibit a couple of noteworthy patterns. First, they converge in the late 1990s as well as after 2015, when gasoline prices are relatively low. Second, their variances tend to increase around the same time, suggesting that lower gas prices may result in a wider range of short-run driving adjustments than higher prices, despite the use of a log-log specification intended to capture responses not to level, but to proportional changes in the price of gasoline.

2.4 Conclusion

In this paper, we contribute to a literature estimating the gasoline price elasticity of demand for driving. This parameter is of high interest to policy makers considering the effectiveness of taxes and other tools intended to curb both greenhouse gas emissions and traffic congestion.

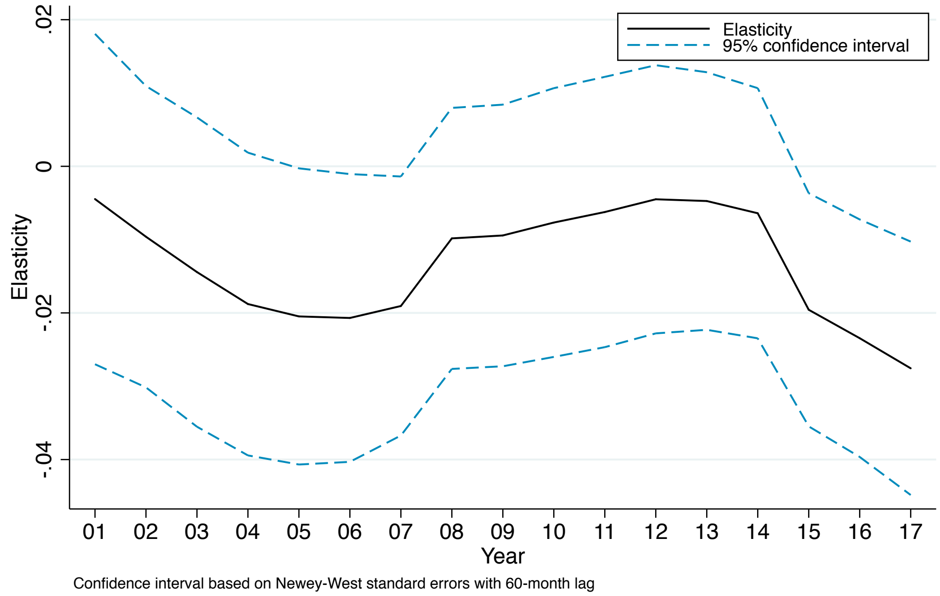

Figure 2.4: Yearly VMT demand elasticity estimates from model with income covariate

It is also an important component in attempting to understand how we might be able to limit large gasoline price spikes such as those experienced most dramatically in California, but also in the US in general. We find significant variation in California's VMT elasticity over time, but much like the existing literature, do not have direct evidence for which factors may underlie this variation. We conclude that more work is needed to improve our grasp of the highly studied but still not fully understood parameter, and the variables that might affect it.

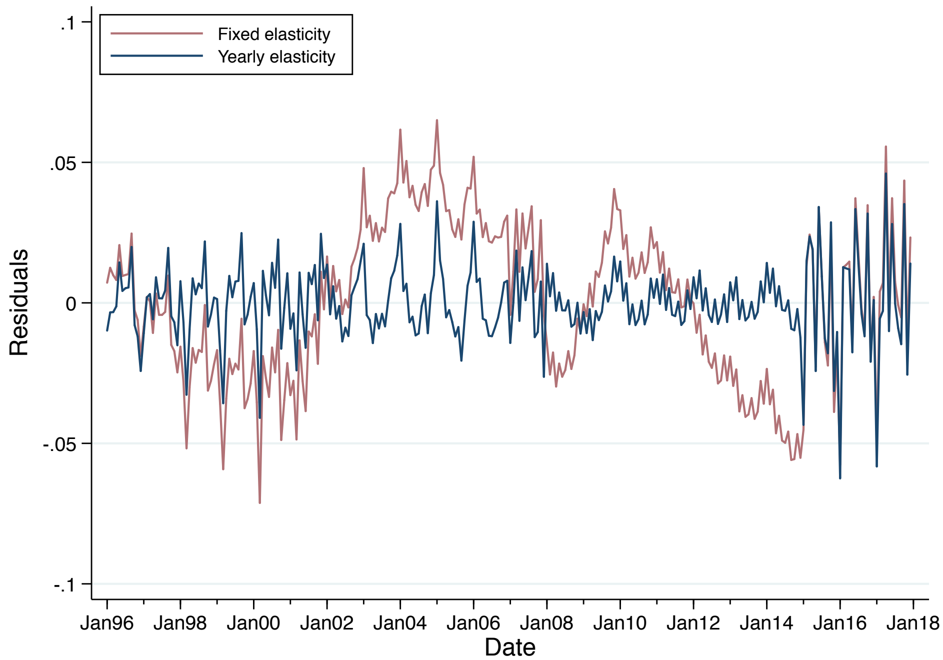

Figure 2.5: Residual comparison: fixed vs yearly elasticity

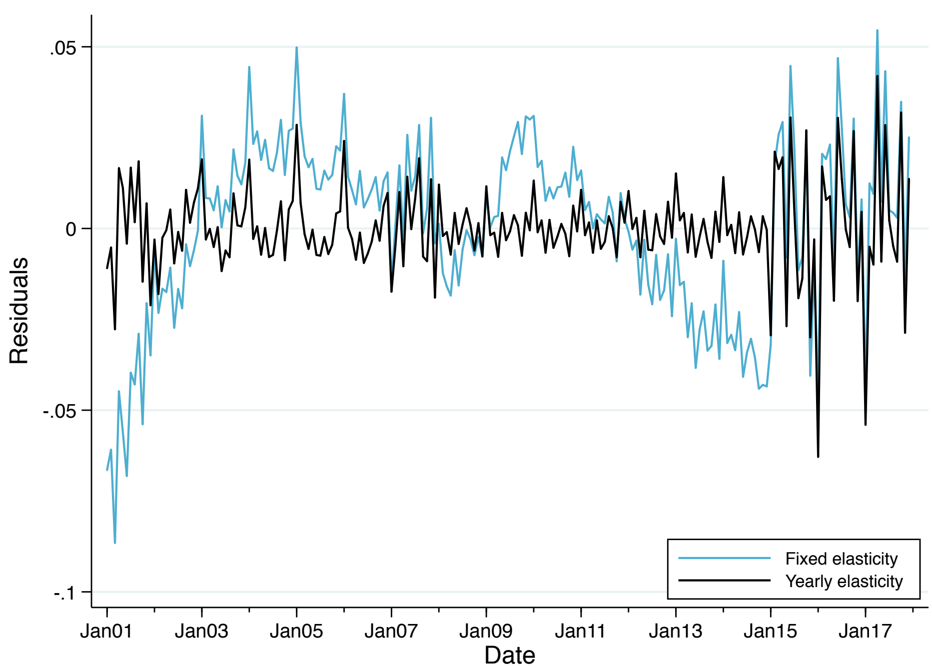

Figure 2.6: Residual comparison: fixed vs yearly elasticity from model with income covariate

Chapter 3

How Hurricanes Sweep Up Housing Markets: Evidence from Florida

3.1 Introduction

North Atlantic tropical cyclones, also known as hurricanes, are extreme weather events affecting populations living along the Atlantic seaboard and the Gulf of Mexico in the United States, in addition to populations in the Caribbean, Central America, and occasionally the northeastern Atlantic seaboard of South America. They generate extreme winds and flooding, which can impose devastating damages to homes and businesses in the vast areas their paths travel. Since the 1970s, climate scientists have observed an intensification of these cyclones, a pattern which is predicted to continue into the second half of this century (IPCC 2013). It is therefore important to study their economic impacts.

While figures depicting direct property losses and fatalities are readily available, it is much more difficult to estimate economy-wide impacts and make long-run projections. In order to credibly approach the latter, it is crucial to understand the mediating forces of market interactions and institutional responses in a large economic entity. In this paper, we study these dynamics in the housing market—often the sector most directly affected by hurricanes—using microeconomic data on Florida home transactions.

Following the seminal work by Rosen (1974), an extensive environmental economics literature has used hedonic regression coefficients to estimate the marginal willingness to pay for nonmarket goods.[1] Most previous studies on housing market responses to hurricanes follow this tradition and focus on their role in conveying risk information to households.[2] Using quasi-experimental designs, this literature has established a causal link between hurricanes and people's valuations for homes through changes in perceived flood risks. However, very little is known about the overall adjustments in the housing market in response to the physical damages of hurricanes and their long-term consequences.

In this paper, we demonstrate that a conjectured, but so far largely unexplored channel through which the housing market responds is a large supply adjustment, as indicated by changes in equilibrium prices and realized transactions. Our research design uses the exogenous variation generated by the randomness of hurricane paths, coupled with the fineness of our housing micro-dataset. We find two sets of main results. First, we establish that home prices spike in the thirty-six months following a hurricane event, in the entire market, as well as in the subset of homes which are sold both before and after the event. Second, we show that the transaction probability of homes in exposed areas falls by a comparable amount and over the same time period. These findings are consistent with the evidence[3] that hurricanes lead to significant physical damages to and destruction of part of the housing stock, thereby temporarily contracting available supply, reducing sales, and in turn driving up prices.

But what are the implications of this shift in the market's equilibrium? We answer this

[1] See Kuminoff et al. (2013) for a recent review.

[2] These studies typically examine differential price changes for houses located in or outside of a flood zone following a hurricane. Examples include Bin and Polasky (2004), Hallstrom and Smith (2005), Bin and Landry (2013), and Gibson et al. (2017).

[3] See, for example, https://www.wsj.com/articles/hurricane-irma-destroyed-25-of-homes-in-florida-keys.

question by matching our housing transactions dataset to data on mortgage applications, which proves to be representative of the entire market, and show that the incomes of buyers who purchase homes (with a mortgage) following a hurricane event increase nearly identically to the rise in home prices we document. Since on average, we observe no subsequent fall in either prices or buyer income below their pre-hurricane levels in any of the eleven years following a disaster, we argue that hurricane-exposed neighborhoods experience permanent increases in wealth, to the extent that entering higher income households are accompanied by relatively more valuable assets. As it has been well established that disaster damages rise with income levels, this has important implications for the costs and distributional impacts of future disaster relief spending.

This paper contributes to a recent literature that estimates the impacts of hurricanes on a variety of economic outcomes. These include flood insurance take-up (Gallagher 2014), demand for groceries (Gagnon and Lopez-Salido 2014), home improvement decisions (McCoy and Zhao 2018), household finance (Bleemer and Van der Klaauw 2017; Gallagher and Hartley 2017), location and employment choices (Deryugina et al. 2018), firm and labor market performance (Belasen and Polachek 2009; Seetharam 2018), migration patterns (Boustan et al. 2017), economic growth (Hsiang and Jina 2014; Strobl 2011), and government spending (Deryugina 2017). Our analysis of the housing market not only adds an important outcome to this existing set, but also complements previous results on mortgage payback and migration to shed light on how market and institutional processes contribute to post-hurricane demographic adjustments. Our research design is very similar to these studies in two major aspects. First, the identification primarily relies on the exogeneity of the location and timing of natural disasters. Second, we are able to consider the full dynamics of our outcome variables, providing evidence of the parallel trends assumption required by our identification strategy.

This paper is also closely related to the literature on the impact of natural disasters on the housing market. The repeated sales strategy behind our price results is particularly similar to that in Hallstrom and Smith (2005), Bin and Polasky (2004), Bin et al. (2008), and Gibson et al. (2017).

An important distinction, however, lies in the selection of the sample, and hence the interpretation of the results. By avoiding directly-affected locations, these studies make a deliberate effort to interpret the hedonic coefficients as flood risk premiums. We pursue the opposite approach by obtaining estimates that can be interpreted as changes in the market equilibrium. Our results are qualitatively similar to those of Murphy and Strobl (2009), which also finds a hurricane-induced upward shift in equilibrium prices. Our transaction-level data, however, enable us to provide crucial evidence regarding changes in equilibrium sales and buyer incomes, allowing us to better understand mechanisms and policy implications.

The remainder of the paper is organized as follows. Section 3.2 describes our data and the procedures we use to define hurricane exposure, Section 3.3 presents our identification strategy, Section 3.4 provides the estimates from our models, and Section 3.5 interprets our findings and concludes.

3.2 Data

3.2.1 Housing Transactions

We obtained Florida housing transaction data between 2000 and 2016 from Zillow, which reports around 95% of housing market transactions over this time period. We match the Zillow data to county tax assessments performed between 2013 and 2016, which Zillow also provided us and includes an essential set of hedonic characteristics for most homes, as well as richer collections of variables for smaller subsets of them. We only retain transactions for which the most important of these hedonic characteristics are available, including timing of the transaction and geographic coordinates,[4] and use census shape files to associate a census tract and county to each observation under both transaction time census county-tract regime and 2000 census county-tract regime. The latter will allow us to employ a set of time invariant fixed effects in our

[4]Respectively year and month, and latitude and longitude.

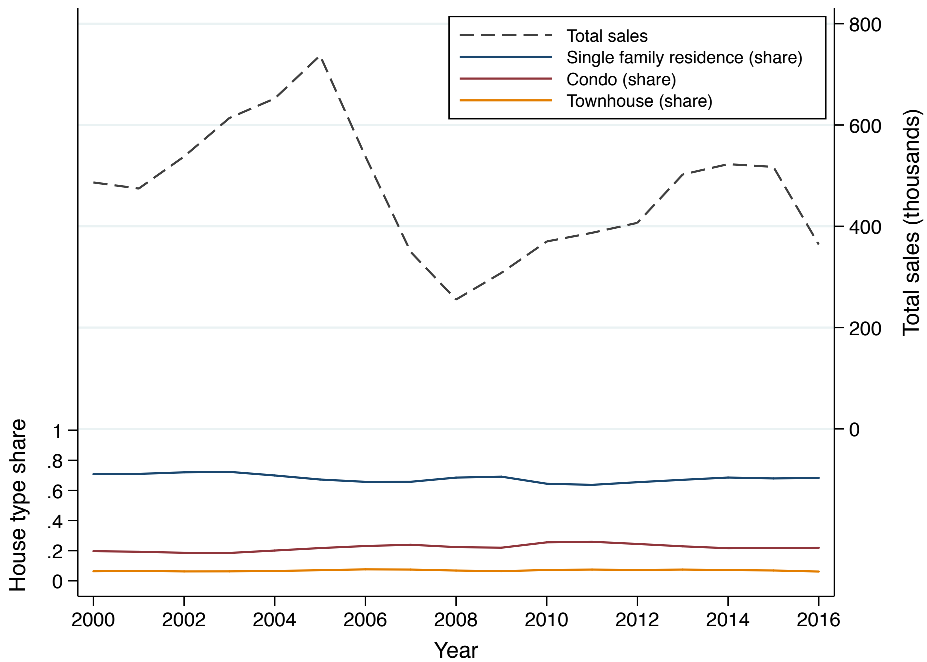

Figure 3.1: Florida housing market sales and composition

models, while we use the former to match the Zillow data for all home purchasers for whom a mortgage is reported by Zillow (i.e. borrowers) to data we obtained from the Home Mortgage Disclosure Act. Figure 3.1 plots yearly aggregate Florida sales, and sales shares by housing type. Figure B2 shows these series for the subset of all Zillow borrowers. Note that borrower home type shares are very similar to those of all buyers, with the exception of condo shares, which are inferior in the borrowers-only sample both in the early 2000s, and after 2007. This difference is accounted for by a higher single family residence share in the full sample for these periods. Next, Figure 3.2 presents monthly median prices by housing type for all transactions (see Figure B3 for borrowers only). Again, the trends and relationships in the two series are strikingly similar, with the exception of condo prices in early 2000 and post-2007, which appear to be relatively higher in the borrowers sample. Combined with both the differential condo sales shares and a dip in the share of homes sold with mortgages over these time periods, these patterns seem to suggest that loans for condos in the lower tail of the price distribution are the most likely types of loans to be foregone when relatively fewer mortgages are issued by financial institutions. Finally, the most dominant trends depicted in these figures is the housing market collapse beginning in 2006, and subsequent recovery from around 2010.

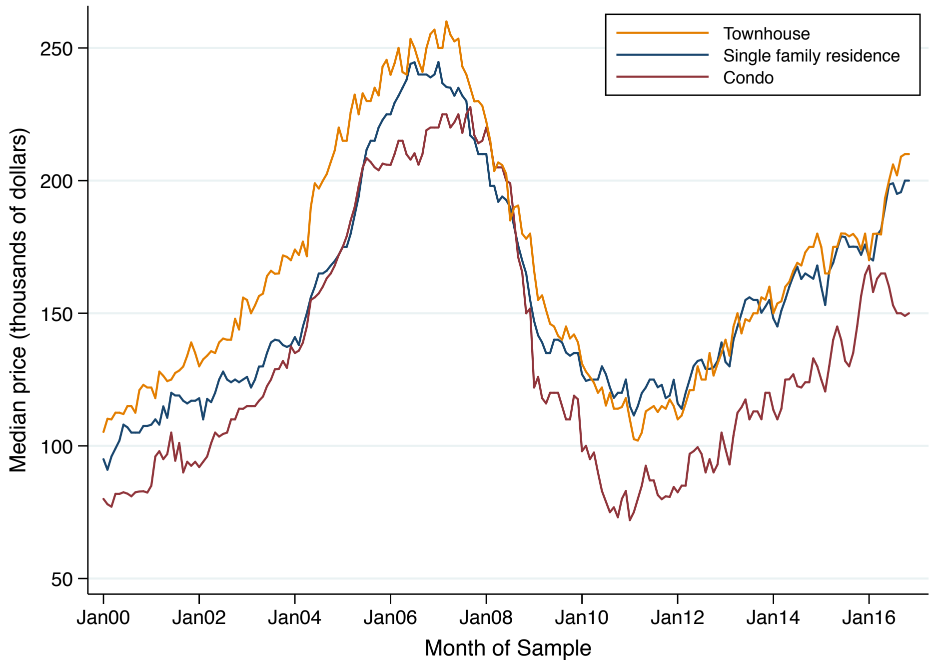

Figure 3.2: Florida housing market prices

3.2.2 Home Mortgage Disclosure Act

The Home Mortgage Disclosure Act (HMDA), enacted by Congress in late 1975, requires all large financial institutions[5] to submit all of their home lending activity every year to the Federal Financial Institutions Examination Council, which makes it publicly available in an effort to improve home financing market transparency. These large institutions must report the date, property location, and amount of each loan application, its purpose (purchase, improvement, or refinancing), and applicant demographic characteristics including annual income, gender, and ethnicity. We merge the subset of successful loan applications for purchases from HMDA to

[5] According to a yearly revised threshold, which is currently set at $45 million in assets.

the universe of Florida borrowers from Zillow by following the matching procedure described in Bayer et al. (2016). Matches are created based on the timing of each transaction and the location of the home,[6] in addition to the loan amount and lender name. The use of this latter variable implies an imperfect merge as some lenders make multiple loans of the same amount in a single census tract every year, and we drop all such matches, as we cannot infer which of the various buyer demographic characteristics should be joined to these multiple transactions. Furthermore, because lender names may be recorded differently in our two datasets, the quality of some pairings will be lower than others. We therefore only keep matches according to a threshold of the Jaccard similarity index[7] above which we observed it to produce the correct pairing in large random subsamples chosen from every year in our data, with special attention paid to matches just above candidate thresholds. A final issue arises when a single HMDA record can be joined to more than one transaction in Zillow, and we drop all such observations as well. The full procedure ultimately leaves us with just over half of the original Zillow borrowers data, with no significant yearly variation in pairing success, which will enable us to look at the evolution of borrower demographic characteristics following hurricane events, as well as to consider possibly differential borrower price sensitivity.

3.2.3 Hurricane History and Measurement

Most Florida census tracts experience at least one hurricane event between 1992 and 2017.[8] Around 90% of tracts are hit once or more, affecting a similar share of Florida's population. Figure 3.3 shows the distribution of the number of hurricane events each tract experienced in the 1992 to 2017 period, where a tract is defined as hit if its population-weighted centroid was ever within reach of the 64 nautical miles per hour (about 119 kilometers per hour) wind speed

[6] Respectively year and census tract, as these are the finest timing and geographic variables available in HMDA.

[7] Produced by an algorithm that compares the elements of two sets, in this case the strings containing the lenders' names from both datasets, and returns a similarity score indicating the extent of the overlap between their elements.

[8] We focus on hurricanes after 1990 because the previous ones to have crossed Florida, hurricanes Elena and Kate in 1985, occurred a distant sixteen years before the beginning of our housing transaction data period.

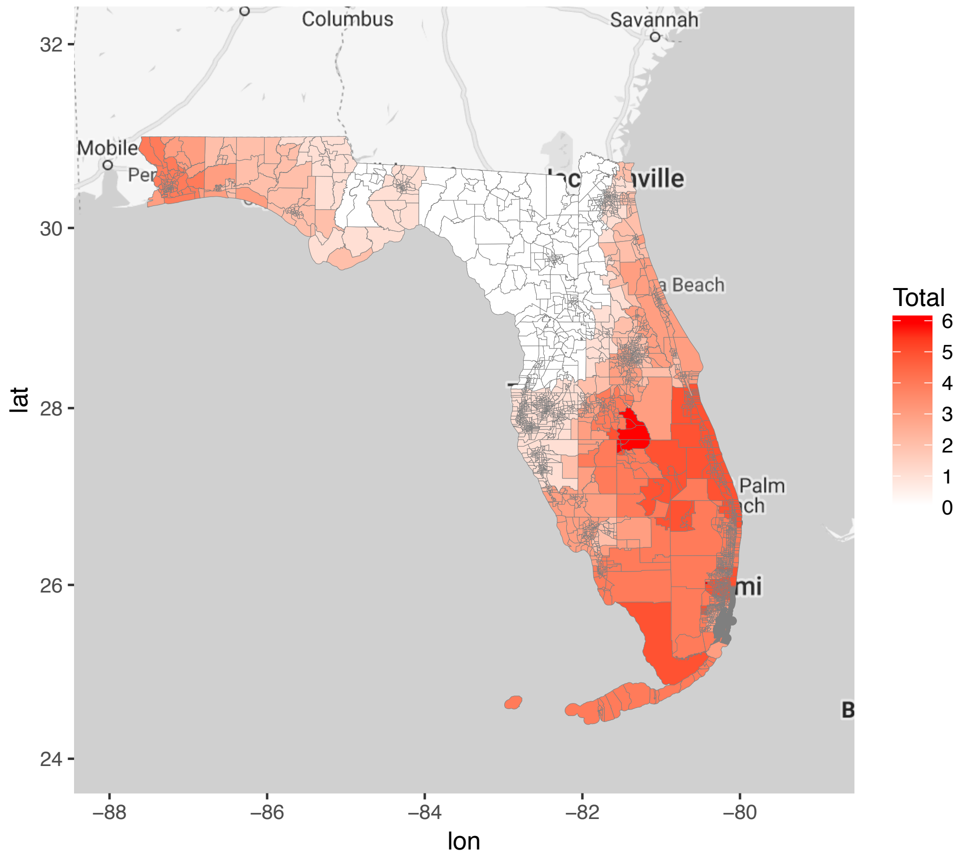

Figure 3.3: Florida exposure to hurricanes by census tract, 1992-2017

threshold that differentiates hurricanes from tropical storms. We obtained maximum sustained wind speed and 64 nautical miles per hour speed radii around each hurricane from the Regional and Mesoscale Meteorology Branch (RAMMB), where data for each hurricane are measured every six hours (at midnight, 6AM, noon, and 6PM). Because distance between measurements varies and can reach upwards of a hundred kilometers, we supplement the RAMMB estimates with our own calculations of the closest the center of each hurricane traveled to every tract centroid,[9] and then interpolate speed and radii at the coordinates where these minimum distances are found using linear weights from the two nearest observed measurements.[10] We further apply

[9]Assuming a linear path between hurricane-track-point observations.

[10]The smallest distances to coordinates along the hurricane's path turn out to be between the two closest track-points in most cases, but we compute minimum distances along paths between any neighboring coordinates of the five nearest observed track-points, as this minimum distance is not necessarily found between the nearest two observed

this procedure to obtain speed and radii interpolations for the nearest each hurricane traveled to every home transacted in our sample.

Next, we use an econometric model leveraging the nonlinear but strictly decreasing relationship between wind speed and its associated maximal reach radius, which we observe in the RAMMB data for 34, 50, and 64 nautical miles per hour (kn) wind speeds, to predict a radius for category 3 threshold wind speeds of 96kn. Our model could have also relied on minimum pressure and maximum wind speed, but we instead choose to employ a set of fixed effects, within which these two variables cannot be identified, as follows:

$$log(Maxradius_{sht}) = \alpha_{ht} + \beta_1 Speed + \beta_2 Speed^2 + \varepsilon_{sht} \tag{3.1}$$

where α_{ht} are hurricane-track-point fixed effects and *Speed* (s) takes one of the three speed values for which RAMMB reports maximal reach radii. The relationship between a wind speed threshold and its associated maximal radius is very well captured by this model as suggested by its estimation's R^2 of 0.93 (0.90 within track-point fixed effects). Full results are reported in Table B1.

Armed with estimates of this relationship, we extrapolate maximum radii predictions for 96kn wind speeds when such speeds are actually reached, and finally verify the consistency of our predictions by comparing them to the radius of the maximum speed reached in these hurricane-track-points, which the RAMMB data also contain. In particular, because radius is strictly decreasing in wind speed, our imputed radius should always be greater than (or equal to) the radius associated with the 96kn and above maximum speed reached in an observed measurement. Our model's prediction satisfies this condition for over 90% of imputations, and is within rounding error (all RAMMB radii measurements are rounded to the nearest 5 nautical miles) for 10 of the 13 extrapolations for which it fails. In these 13 cases, we replace the model's

coordinates—which may not even be neighbors. We then retain min{set of calculated minimum distances, minimum observed distance} since an observed reading itself could also be closest, because of hurricane track curvature.

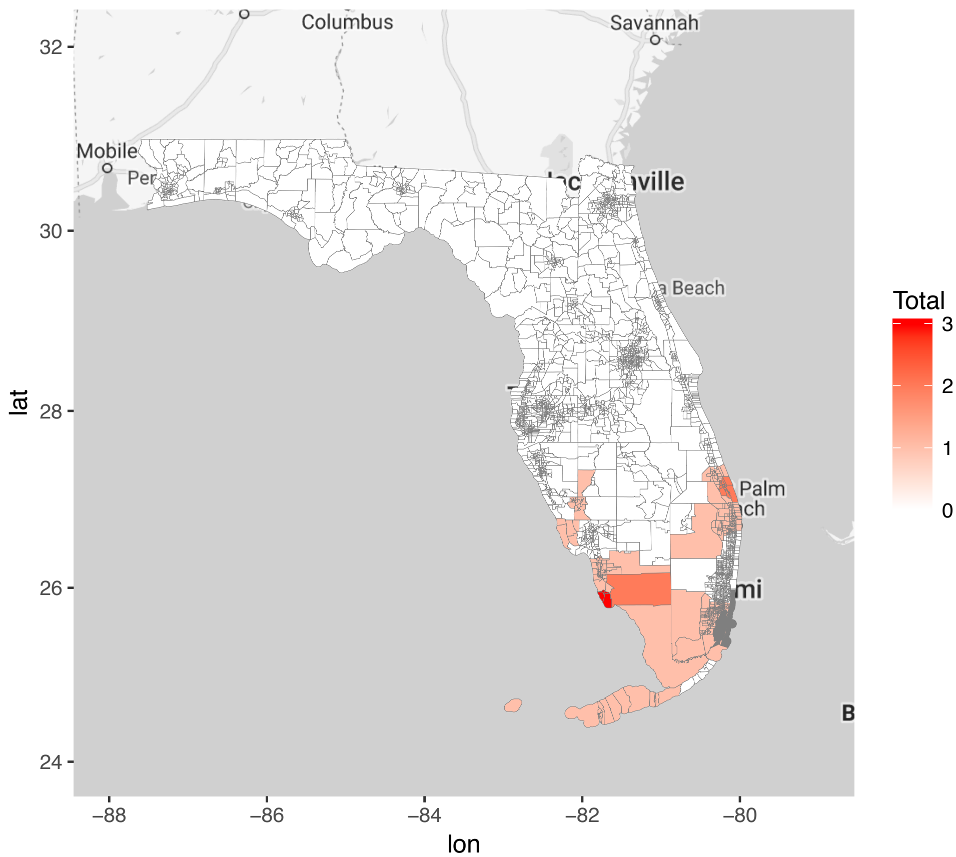

Figure 3.4: Florida exposure to category 3+ wind speeds by census tract, 1992-2017

predicted maximal radius for 96kn speeds with the wider observed radius of the maximum speed reached. We finally apply our interpolation scheme to predict 96kn maximal radii for the location nearest a tract centroid (or house in our sample) along a hurricane's path. Figure 3.4 shows the distribution of the number times census tracts were affected by category 3 wind speeds and above, according to whether a census tract population centroid was within reach of such speeds. Note that many fewer tracts (only about 15%) were ever exposed to 96kn wind speeds than to the 64kn hurricane speed threshold.

3.2.4 Hurricane Exposure

To determine whether a home in our sample was exposed to a hurricane between 1992 and 2017, we rely on the interpolated speed and radii estimates associated with the coordinates of the closest point along a hurricane's path, as described in Section 3.2.3. Similarly to the tract hit definition, we consider a home to have been exposed if it was within reach of a 64kn wind speed radius when the hurricane was nearest. Any home outside the reach of this radius, or for which the wind speed at its smallest distance to the hurricane's path was less than 64kn, is therefore considered unaffected. Tables B2 and B3 report the percentage of housing transactions in our data affected by each hurricane event in one year increments, respectively before and after each hurricane, where year -1 refers to twelve months before a hurricane, year 0 the first twelve months after, year 1 the next twelve months, and so forth. The last row reports cumulative treatment in each event time year, and note that this percentage is less than or equal to the sum of individual hurricane treatments because transacted homes may have been hit by multiple hurricanes in an event year, usually when such hurricanes occurred in the same calendar year. Because we're interested in exploring possible heterogeneity in housing market response to different hurricane intensities, we define indicators for exposure to category 3 hurricanes by considering exposure to our imputed 96kn wind speeds.[11] Tables B4 and B5 show transacted home percentage affected by category 3 or greater wind speeds. Tables B6 through B9 report parallel percentages using the census tract population-weighted centroid definition of treatment, which are similar, but mechanically cannot exhibit as much variation in exposure.

These tables all illustrate that the seventeen-year width of our repeated cross-section constrains the length of pre and post hurricane indicators we can include in our models while maintaining the interpretation of our findings to result from exposure to an average hurricane in our sample, as opposed to exposure to a small, specific subset of hurricanes. Including indicators

[11] We use the category 3 speed threshold both to be in line with previous literature, and because we do not believe the difference between the categories 1 and 2 thresholds is sufficient to produce measurable differences in outcomes. Category 4 wind speeds, on the other hand, almost never reach Florida shores over our hurricane time period.

for event years -10 to -7, for example, would coerce the identification of these pre-trends to come from (future) exposure to only two hurricanes, raising serious representativeness concerns, not to mention the added challenge of effectively controlling for unobservable time-varying differences between treated and untreated entities when treatment saturation is so low that few observations are driving the entirety of treatment effect identification.

3.3 Econometric Framework

We are primarily interested in how hurricanes affect the equilibrium prices in Florida's local housing markets. We model the dynamics of housing prices relative to hurricane events as follows:

$$log(Price_{ihmy}) = \sum_{\tau=-6}^{10} \beta_\tau Hurr_{imy}^{\tau} + HouseChar_{iy} + \delta_{ht} + \delta_{hm} + \delta_{hcy} + \varepsilon_{ihmy} \tag{3.2}$$

where i denotes an individual transaction, h denotes house type,[12] m is the month, and y the year of the transaction, and t and c are respectively the census tract and county in which the transaction occurred.[13] The unit of analysis is an individual transaction. $log(Price_{ihmy})$ is the log of the price of transaction i, taking place in month m of year y. $Hurr_{imy}^{\tau}$ is a set of indicators specifying whether the transacted house was affected by a hurricane τ years before (negative τs indicate a post-transaction hurricane, $\tau = 0$ refers to transactions in the first twelve months after a hurricane, $\tau = 1$ the next twelve months, and so on). $HouseChar_i$ is a set of house characteristics commonly used in hedonic models, including structural age, effective age[14], number of stories and number of bathrooms. We control for the latter two characteristics flexibly using a set of value bins.

[12]We define six main house types based on land use classification in county tax assessments. They include single family residential (68.3%), condominium (21.7%), townhouse (6.83%), residential-multifamily (2.61%), vacation home (0.14%) and miscellaneous (0.4%).

[13]Census tracts and counties are defined according to the 2000 census throughout to maintain geographic consistency across time.

[14]The time (in years) since the house last saw a major remodel.

We also account for both fixed and time-varying regional differences in housing attributes and local amenities through a rich set of geographic and temporal fixed effects. δ_{ht} are type-by-tract fixed effects, which eliminate cross-sectional correlations in the likelihood of being hit by a hurricane and time-invariant local amenities, such as being close to the coast. δ_{hm} are type-by-month fixed effects, controlling for the seasonality in both home prices and the timing of hurricanes. Since the financial crisis took place during our sample period and significantly impacted both the prices and the number of home sales, we also believe it is crucial to account for its heterogeneous effects across markets. We do this by including type-by-county-by-year fixed effects (δ_{hcy}) to control for changing macroeconomic conditions at the county level.

The key variables of interest are the hurricane indicators, whose construction was detailed in Section 3.2.4. The identification of the causal effect of hurricanes on the housing market relies on the exogeneity of storm paths and timing. Specifically, the identifying assumption is that, conditional on the set of controls, these indicators are orthogonal to any idiosyncratic shock (ε_{ihmy}) impacting our three transaction outcomes of interest (here prices, and later sale probabilities, and new homeowner incomes).

We estimate this model on the full sample, the borrower sample, and the HMDA sample. Their estimates represent the average effects on housing prices, which could result from two mechanisms:

1. (Supply) Hurricanes can disrupt the stock of housing in affected areas through destruction, and the potential associated time spent rebuilding.

2. (Demand) Hurricanes might directly change the valuation of houses in both impacted and neighboring areas.

To further our understanding of these mechanisms, we estimate several variants of equation (3.2). First, we replace the tract fixed effects with parcel-level ones. This approach restricts the identifying variation to price changes from repeated sales of the same home, allowing us to capture the equilibrium price effects holding *all* characteristics fixed, observable or not. Second, we will later change the outcome variable to estimate the effects of hurricanes on the incomes of average HMDA borrowers. This will shed light on any systematic shifts in new homeowners' economic profiles.

To investigate the extensive-margin responses directly, we also estimate the impact of hurricanes on the probability of transaction using the following model:

$$1(Transacted)_{py} = \sum_{\tau=-6}^{10} \beta_\tau Hurr_{py}^{\tau} + HouseChar_{py} + \delta_p + \delta_{hcy} + \varepsilon_{py}. \tag{3.3}$$

where p, h, y, and c denote parcel, housing type, year, and county as before. The unit of analysis, however, is now a parcel-year. $1(Transacted)_{py}$ is an indicator of whether a transaction record exists for parcel p in year y. The variables of interest are still the hurricane indicators, but they are now defined relative to the observation year. Time-varying characteristics of the house on the parcel, such as age and effective age, are also included. Because the threat to identification is similar to that in the previous model, we also include month and county-year fixed effects, in addition to the parcel fixed effects which will ensure that the price comparison is performed within-home. Furthermore, we redefine years to begin in August and end in July, because August is the earliest month in which any home from our sample could have been transacted immediately following a hurricane event. As such, we ensure that the year 0 hurricane event indicator for homes in our sample which were not actually transacted in a given year includes most homes which experienced hurricane winds preceding the roughly twelve-month period in which it was not transacted. The estimate of the event time 0 indicator would have otherwise been severely attenuated, as all pre-hurricane season non-transactions of the calendar year would have contributed to the time 0 indicator despite the hurricane occurring *after* these months.[15]

[15] Some hurricanes in our sample occur as late as October. This still implies that some non-transactions will occur in the same "year" as a hurricane which affected the home, and thus have their event time 0 indicator set to 1, despite the hurricane strike happening after the non-sale. Given our choice of August to begin the year, this mechanical issue can only affect transactions occurring in the three months preceding hurricane Wilma (October 2005) or the two

Additionally, only the remaining two- to four-month period of each calendar year following the hurricane season could have contributed to event year 0 treatment otherwise, while the subsequent months would have been included in the year 1 indicator despite these non-transactions occurring possibly less than three months after a hurricane.

We estimate this transaction probability model both overall and separately on the three main parcel types—single family residences, condominia, and townhouses. For each type, we construct a panel of all parcels that have been transacted during the sample period, and treat all parcel-year observations without a transaction record as having no transaction, unless the house had not yet been built, in which case we omit the observation. This approach introduces measurement error if there are unreported transactions. It can even lead to biased estimates if the missing pattern is endogenous to or correlated with hurricane events. In our analysis, we take into account detectable patterns of missing records in a few county-years by omitting these entire counties from our sample. Whether we omit the entire sales histories or only the missing years of data for these counties does not yield any distinguishable differences in results, as their combined sales account for less than 1% of total sales in the period in which no data are missing (2005-2016).

3.4 Results

3.4.1 Post-Hurricane Price Dynamics

The event time indicator coefficients from estimation of equation (3.2) are plotted along with their 95% confidence intervals in Figure 3.5, and full results are reported in the first column

months predecing Frances and Jeanne (both dissipating in September 2004). The ensuing bias towards 0 this implies on our estimate of the time 0 indicator is small and predictable, which is why we prefer it to the other alternative of beginning the year even later and contaminating *both* the year -1 event indicator with homes which were actually exposed to hurricanes prior to non-transaction *and* the year 0 indicator which should have instead been turned on, but was not. While our approach does remove a few observations from contributing to the identification of the -1 event time indicator, it is not contaminated since a future hurricane cannot have a causal effect on any of our outcomes in the year before it hits, thanks to the unpredictability of its path.

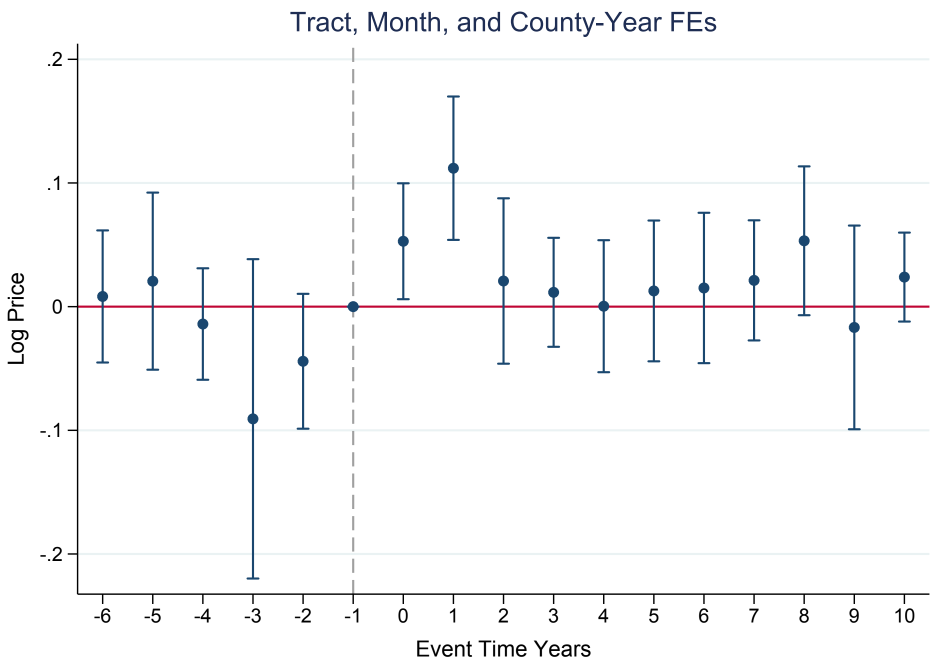

Figure 3.5: Hurricane effects on house prices – full sample

of Table B10.[16] The standard errors we rely on to construct the confidence intervals were clustered to allow for correlation in the idiosyncratic shocks to the prices of all transactions occurring in the same county, over the entire width of our repeated cross-section. None of the estimates for pre-hurricane transaction indicators are statistically different from 0, supporting our choice of the fixed effects we employ to control for preexisting differences in average census tract prices, pricing seasonality, as well as the county level time-varying differences in prices. Our event time 0 and 1 indicator estimates suggest that hurricanes result in increases in home prices of 5% in the first and 10% in the second twelve-month periods after the strike. The point estimate for the event time 2 indicator signals that the surge in prices ends sometime over the third post-hurricane twelve-month period, as the estimated increase relative to unaffected homes drops to 2%, and is no longer statistically significant. All later event time indicators are small, and not statistically

[16]Data provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at http://www.zillow.com/ztrax. The results and opinions are those of the authors and do not reflect the position of Zillow Group.

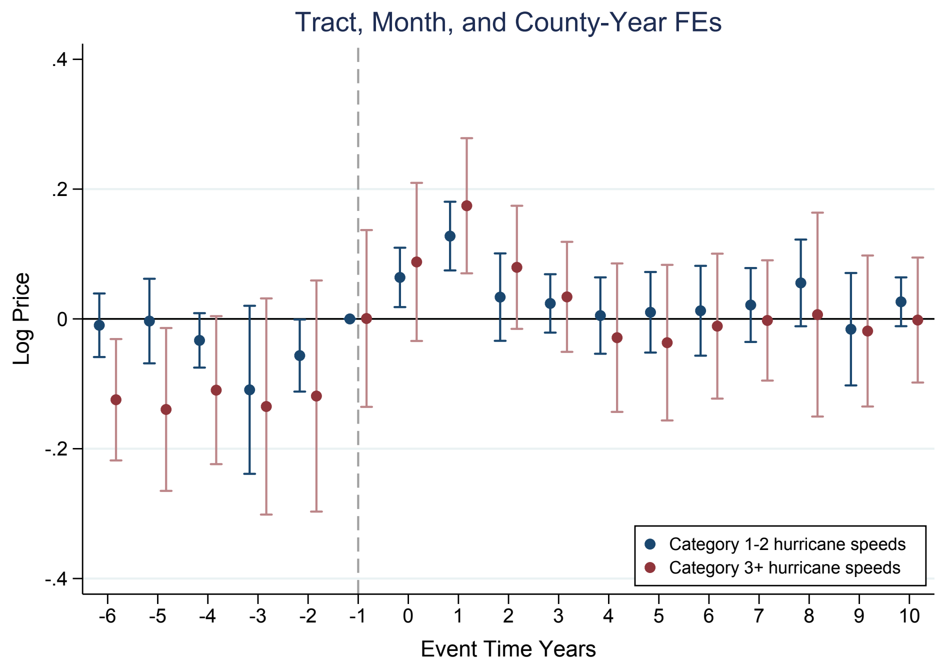

Figure 3.6: Heterogeneous effects of differential hurricane intensity – full sample

distinguishable from 0.[16]

Because we want to explore the heterogeneous effects of differential wind speed exposure, we next redefine our event time indicators in terms of the 64 to 95, and 96 and above thresholds differentiating categories 1 and 2 from category 3 and above hurricanes, and re-estimate equation (3.2). The results from this exercise are plotted in Figure 3.6. Note the similarity of the category 1 and 2 only estimates to those of the effects of any hurricane intensity winds from Figure 3.5, likely driven by the very large overlap between hit definitions. The category 3 and above coefficient estimates, on the other hand, rely on the variation depicted in Tables B4 and B5, and suffer from much wider confidence intervals as a result of the significantly lower treatment saturation. At most 1% of homes are ever sold from areas affected by category 3 and above wind speeds in any year before or after a hurricane, as such wind speeds rarely affected Florida over our

[16]Data provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at http://www.zillow.com/ztrax. The results and opinions are those of the authors and do not reflect the position of Zillow Group.

sample period (recall Figure 3.4, with the caveat that it depicts tract-wide and not individual home category 3 hurricane exposure). It will therefore be difficult for our model to identify differential responses by hurricane wind speed thresholds.

3.4.2 Dynamics of Repeated Transaction Sales Prices

One possible interpretation of the finding of increased sales prices in the immediate couple of years following a hurricane is that they resulted from a shift in the distribution of the hurricane resistance, quality or general desirability of transacted homes, without actually affecting any individual home's price. In order to investigate this possibility, we re-estimate our price models with the inclusion of parcel fixed effects, and restrict the sample to only include homes which were transacted at least once both before and after a hurricane. We find that the patterns documented in Section 3.4.1 are closely reproduced, as shown by the results in Figures 3.7, B4, and the second column of Table B10. These suggest that while there may indeed be a change in the characteristics of those homes which are transacted after a hurricane event (a shift towards more wind and flood resistant structures, for example), the actual homes sold in the first two to three years after being hit themselves appreciated relative to when they were sold outside of this post-hurricane window. Our point estimates imply that these homes on average sold at 5% higher prices in the twelve months immediately following the hurricane, as much as 14% in the next twelve months, and 8% in the third twelve months. In order to better understand the mechanisms behind this finding, we now turn towards potential hurricane-induced changes in the probability of home transaction.[16]

[16] Data provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at http://www.zillow.com/ztrax. The results and opinions are those of the authors and do not reflect the position of Zillow Group.

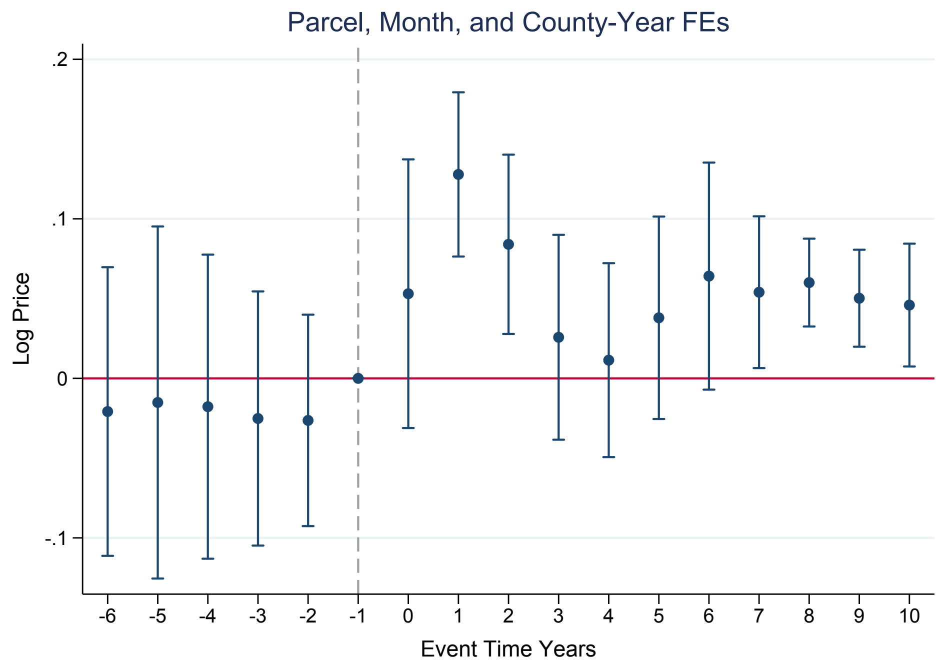

Figure 3.7: Hurricane effects on house prices – repeated sales

3.4.3 Dynamics of Transaction Probability

Hurricanes may influence the probability of home transaction in potentially competing ways. First, the extreme winds and severe flooding they generate can destroy or cause sufficient damages to buildings, making them very difficult to sell. This would result in a post-disaster inward supply shift, as it takes time for these would-be sellers to seek insurance payments or financial assistance from FEMA and rebuild their homes. Second, as a direct consequence of home destruction, affected homeowners may seek to temporarily rent other housing (or possibly buy) in unaffected neighboring communities, increasing demand in the rental market, which could in turn put upward pressure on demand in the purchasing market. Finally, hurricanes could affect demand for housing, not only in the immediately affected areas, but possibly at a broader geographic level, as a combination of information from directly affected friends and family members, and national media coverage update potential buyers' beliefs regarding either the probability of hurricane exposure in the region, or the effects of the realization of a hurricane event on impacted homes. We first attempt to rule out a migration-driven demand shock by examining population trends around hurricane years in affected and unaffected counties. Figure B1 shows that there is no evidence of migration into or out of Florida counties, whether or not they were recently impacted by a hurricane.[16]

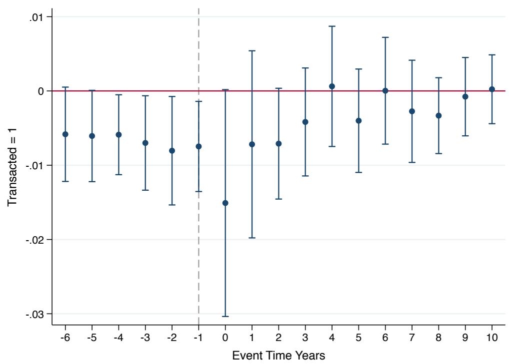

(a) All parcels

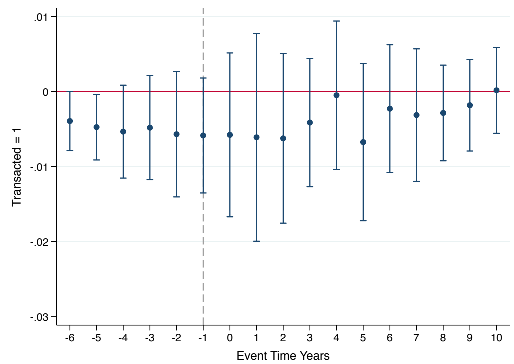

(b) Single family residences

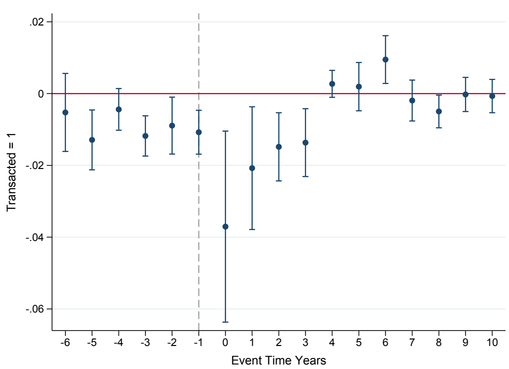

(c) Condominia

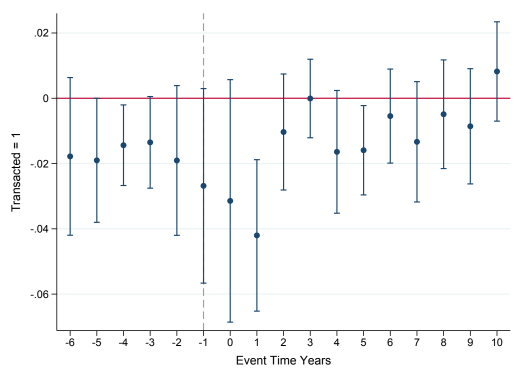

(d) Townhouses

Figure 3.8: Hurricane effects on transaction probability by parcel type

A large majority of affected homeowners seek post-hurricane financial relief in the form of insurance payment or FEMA assistance (see for example, Deryugina et al. 2018; Gallagher and Hartley 2017; Hoople 2013; Michel-Kerjan 2010; Kutz and Ryan 2006). This process requires otherwise potential sellers to wait for assessment of the damages to be performed, which could

[16]Data provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at http://www.zillow.com/ztrax. The results and opinions are those of the authors and do not reflect the position of Zillow Group.

take weeks, in addition to spending months rebuilding or repairing their homes, if they choose to do so, before attempting to sell it. In order to further explore this issue, we look for adjustments to the equilibrium transaction probability of homes which were recently affected by a hurricane event. To accomplish this, we reshape our full Florida dataset such that each observation is now a parcel-year with an associated indicator for whether or not a transaction occurred in that year, as described in Section 3.3.[16]

We first estimate our transaction probability model on the full sample of homes and provide results in panel (a) of Figure 3.8 (as well as in the third column of Table B10). Baseline probability of transaction is around 10%. Despite not being statistically different from 0 at the 5% significance level, our estimates do not allow us to reject outcomes ranging from a -3 percentage point decrease in transaction probability to a 1 percentage point increase across different post-hurricane years. Because such wide standard errors could suggest the presence of heterogeneity of the effect among affected homes, we proceed to estimating the model for single family residences, condominia, and townhouses separately. The null can still not be rejected in any event time period for single family residences (panel (b)), which make up around 70% of all housing market transactions. In contrast, we do observe a change in transaction probability for the latter two types of homes. Condominia (panel (c)), which usually represent 20% of transactions, are nearly 4 percentage points less likely to be transacted in the first twelve months post-hurricane, 2 percentage points in the following twelve months, and around 1 percentage point within both three and four years of the hurricane, while their sales may rebound by 1 percentage point for a couple of the ensuing years. Townhouses (panel (d)) make up the smallest share of the market at under 10%, and experience a possible drop of 3 percentage points in the first twelve months, and lose a statistically significant 4 percentage points in the next. Taken together, these results imply a 1 percentage point decrease in transaction probability in the entire market, or, given the 10%

[16]Data provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at http://www.zillow.com/ztrax. The results and opinions are those of the authors and do not reflect the position of Zillow Group.

baseline, a 10% fall in equilibrium sales in the first full year after a hurricane. After the second year, the implied drop in sales corresponds to around 8% of the market, and it is limited to 2% in the third and fourth. In conjunction with the increase in prices we documented in Sections 3.4.1 and 3.4.2, this fall in transactions is consistent with a supply shock, likely resulting from the damage to the housing stock caused by hurricanes.[16]

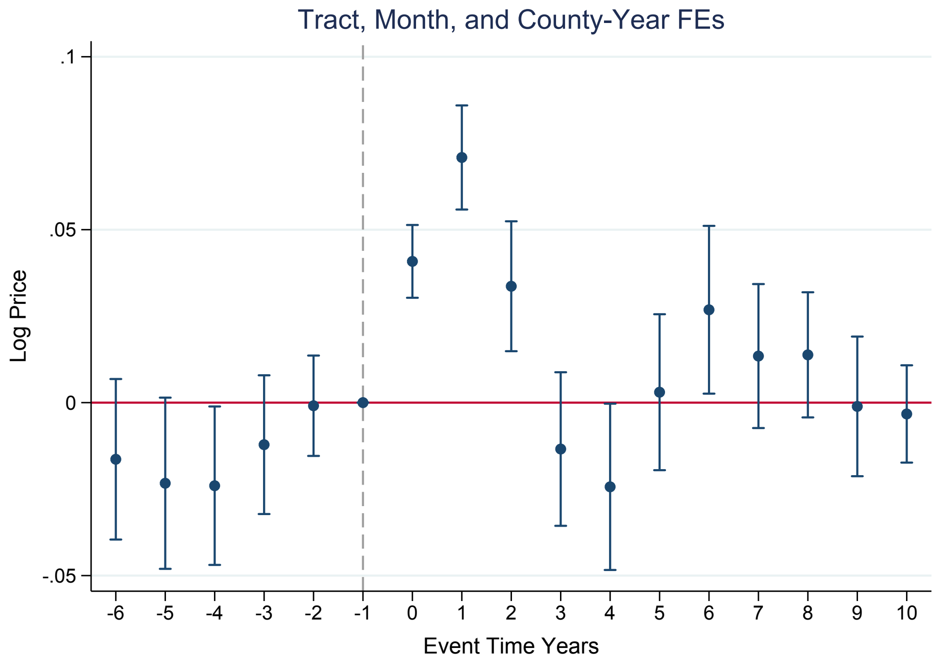

Figure 3.9: Hurricane effects on house prices – HMDA sample

3.4.4 Changes in New Homeowner Characteristics

Armed with a better understanding of the equilibrium shifts in transaction probability and prices occurring in the housing market following hurricane events, we now focus on the potential implications of these changes on hurricane-battered communities. We will concentrate on the evolution of income in the sample of borrowers obtained from the Zillow-HMDA matching

[16]Data provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at http://www.zillow.com/ztrax. The results and opinions are those of the authors and do not reflect the position of Zillow Group.

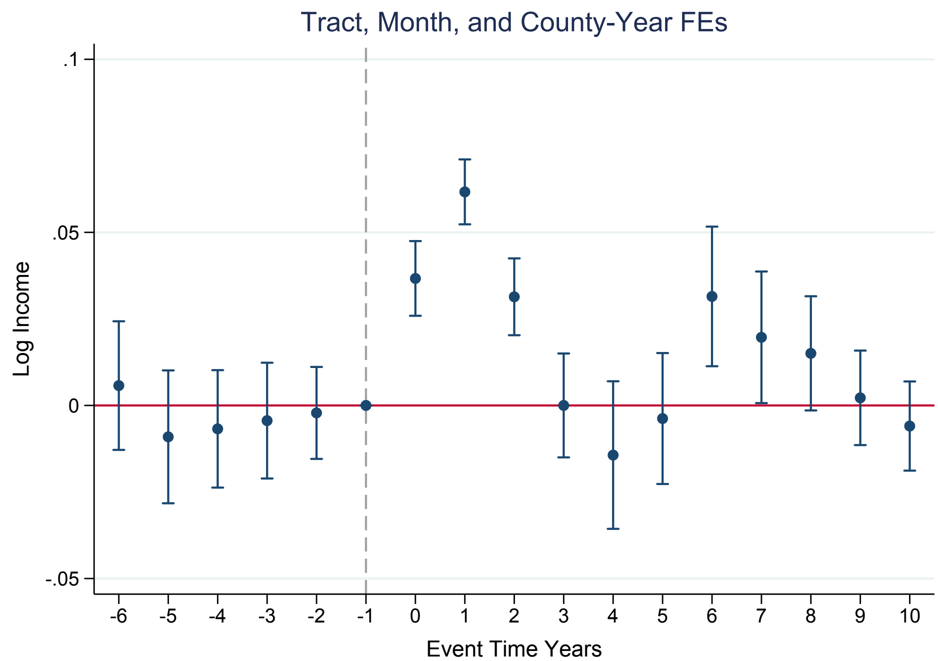

Figure 3.10: Hurricane effects on buyer income – HMDA sample

procedure, but we begin by estimating equation (3.2) on the subsample of all Zillow borrowers, and show, through the results plotted in Figures B5 and B6, that they exhibit a hurricane price response closely resembling that of our estimates on the full sample. Next, we repeat this exercise on the Zillow-HMDA matched sample of borrowers to demonstrate that these estimates, reproduced in Figures 3.9 and B7 are again very comparable to both sets of estimates derived from the borrowers-only and full Zillow samples. We finally turn to the effect of hurricanes on the composition of the income distribution of new buyers, and find, perhaps unsurprisingly, a similar pattern in post-hurricane average incomes as in prices: incomes increase by around 4% in the first twelve months, nearly 7% in the second twelve months, and revert to a 4% rise in the third (and seventh) twelve months before returning to their pre-hurricane means in later event time periods.[16] Because these point estimates are nearly identical to the price shifts we observe in

[16]Data provided by Zillow through the Zillow Transaction and Assessment Dataset (ZTRAX). More information on accessing the data can be found at http://www.zillow.com/ztrax. The results and opinions are those of the authors and do not reflect the position of Zillow Group.

the comparable HMDA-Zillow sample, it is natural to hypothesize that new buyer incomes move almost exactly in accordance with housing market prices. Using this hypothesis to extrapolate income estimates to the full Florida sample therefore suggests increases of 5, and over 10% in post-hurricane average buyer income, respectively for the first and second twelve months after exposure. Note that the income shift observed in these periods is permanent, in the sense that there is no later period income reversion below the average incomes of buyers from unaffected areas. This implies that the mean of the income distribution of hurricane-exposed neighborhoods is permanently higher post-disaster.

3.5 Conclusion

We find remarkably little evidence of any price adjustments to tropical cyclones beyond the response to the supply shock which our analysis of equilibrium transaction probabilities and prices carefully identifies. One potential explanation for this conclusion is that home buyers are very well informed about the risk of hurricane exposure and averse neither to the risk, nor to magnitude of damages inflicted in the realization of an event. Another possible interpretation is that while there is a behavioral response on the demand side of the market, it is dwarfed by the effect of the supply shock imposed by a hurricane event. This could result from a couple of potentially complimentary mechanisms. First, the stock of sellable houses on the market after a hurricane could drop so much that it results in an overwhelmingly large shock to supply, relative to any change in demand. Another channel could also be contributing to the dominance of this supply shock by increasing the value of the existing housing stock. Rebuilding after a hurricane might imply investing only in cosmetic improvements in an effort to signal that a home was not affected as severely as potential buyers might have expected. Reconstruction could also try to better account for the possibility of future events through investments in more wind and flooding resistant structures and materials.[17] Either or a combination of both types of improvements likely results in an increase in owners' expected valuations of their homes, raising prices at every quantity level along the supply curve: another contribution to the negative supply shock.

Finally, we return to the shift in average income we observed to discuss its potential implications. Recall that thanks to our within-parcel results, we know that not only do home prices increase in the years immediately following hurricane events, but also that the prices of the specific homes which were sold in their aftermaths appreciate relative to their non post-hurricane period sales. Coupled with the evidence we find that the average income of new buyers increases similarly over the same period, without later dropping below their original pre-hurricane averages, these findings depict a permanently richer demographic inhabiting hurricane-affected communities following a disaster. To the extent that such a demographic brings along more expensive assets, and spurs more economic development in these areas, this could result in more expensive future hurricane damage claims to disaster relief organizations, notably including publicly funded FEMA.

[17]This is very probable for Florida owners of older homes who rebuild them, as a 2018 report on new residential building codes and enforcement systems by the Insurance Institute for Business & Home Safety gave Florida a rating of 95/100, the highest score among all the states it evaluated.

Acknowledgments

This chapter, in part, is currently being prepared for submission for publication of the material. Panassié, Yann; Liao, Yanjun. The dissertation author was the co-investigator and co-author of this material.

Appendix A

Chapter 1 Figures and Tables

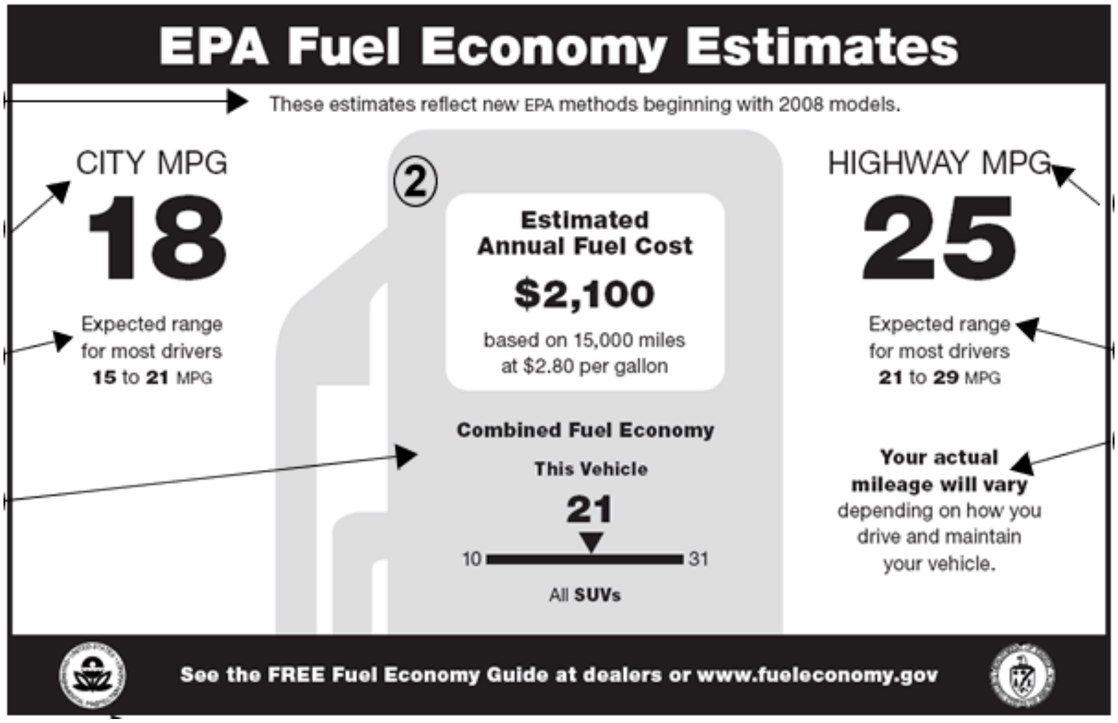

Figure A1: 2008 fuel economy label. Source: EPA

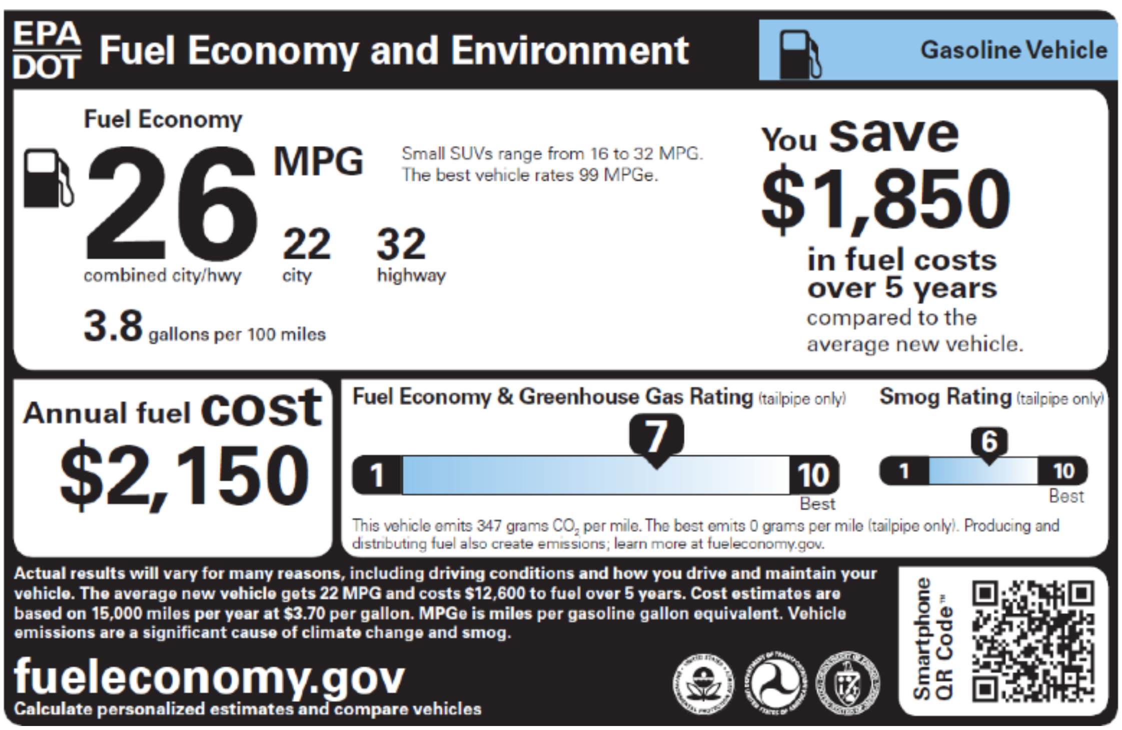

Figure A2: 2013 gasoline engine vehicle fuel economy label. Source: EPA

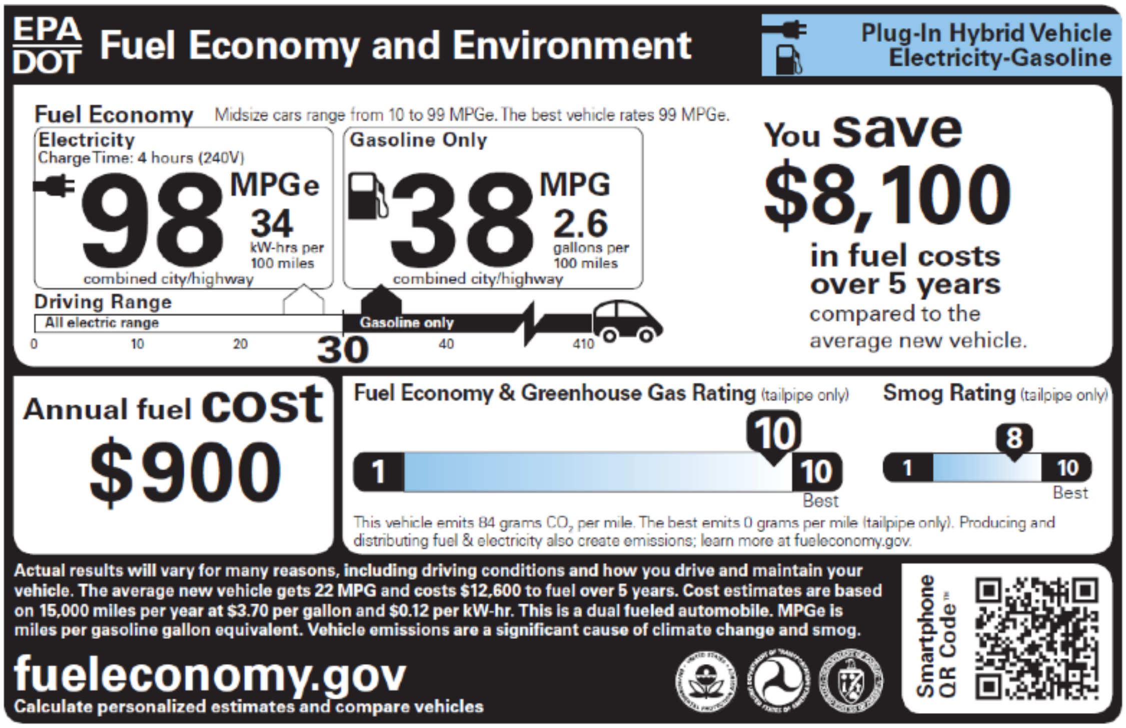

Figure A3: 2013 hybrid engine vehicle fuel economy label. Source: EPA

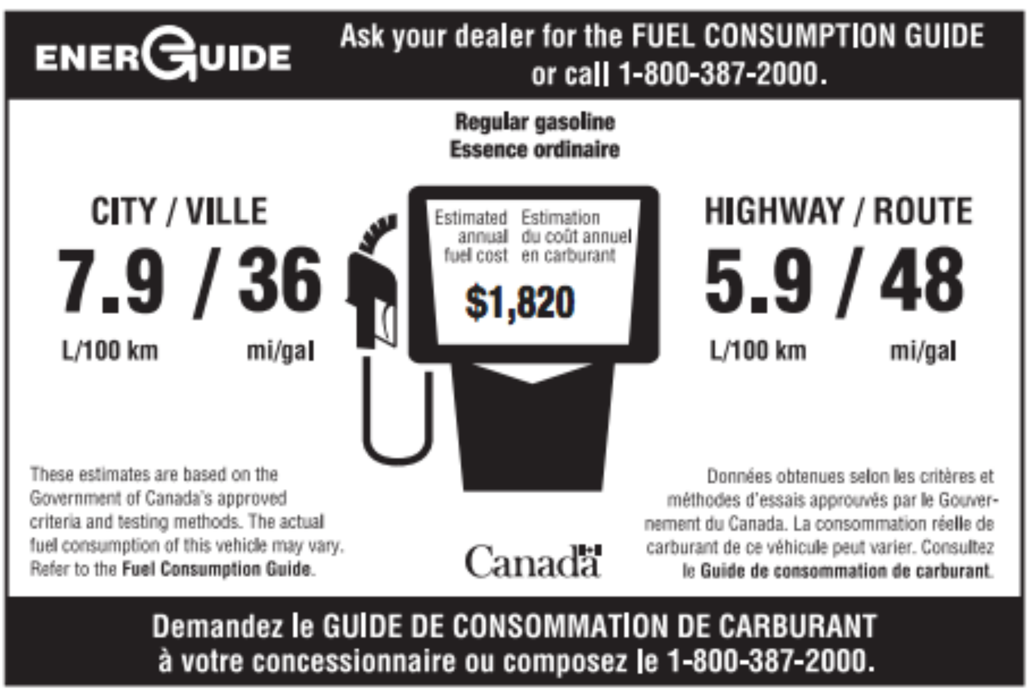

Figure A4: Canadian fuel economy label in use until 2015. Source: Natural Resources Canada

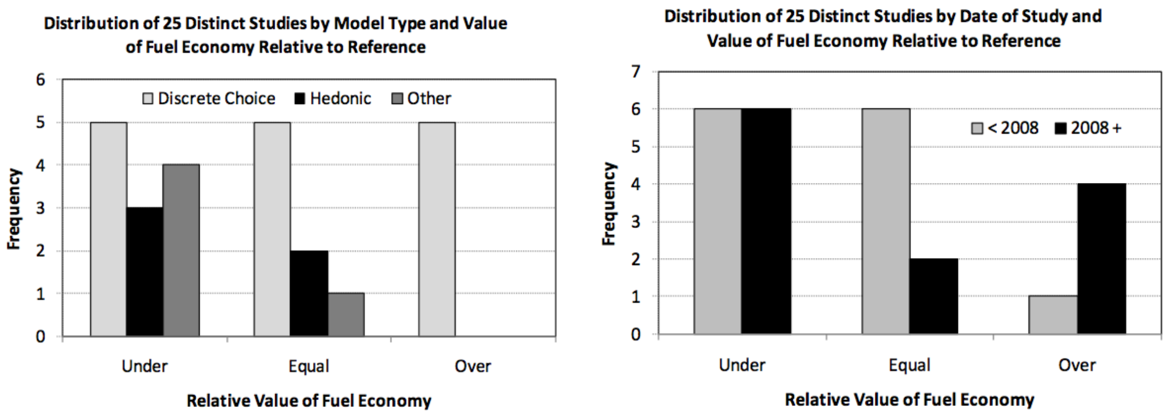

Figure A5: Distribution of valuations, Greene (2010). Reference is a 7% discount rate.

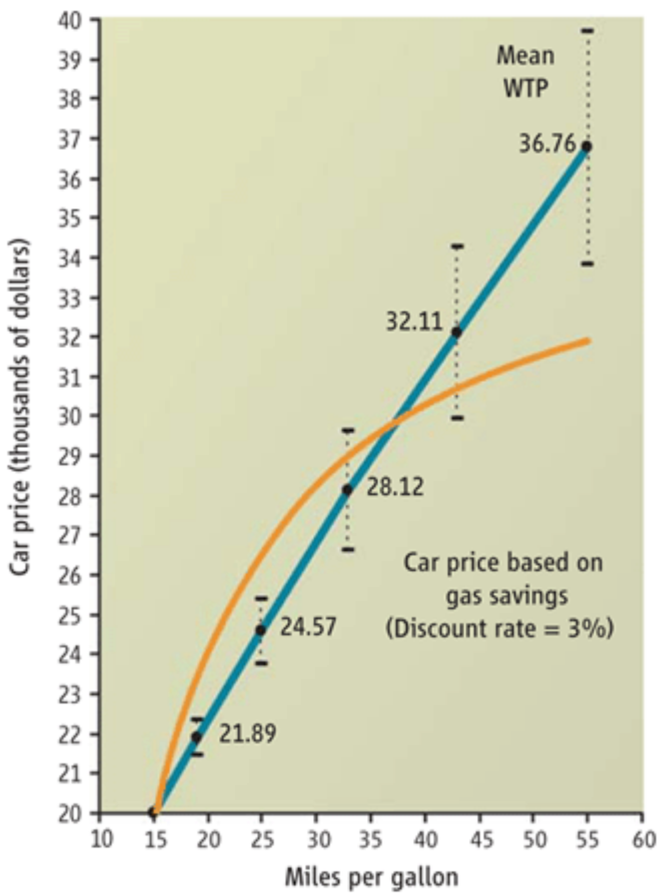

Figure A6: "MPG Illusion", from Larrick and Soll (2008). Reprinted with permission from AAAS.

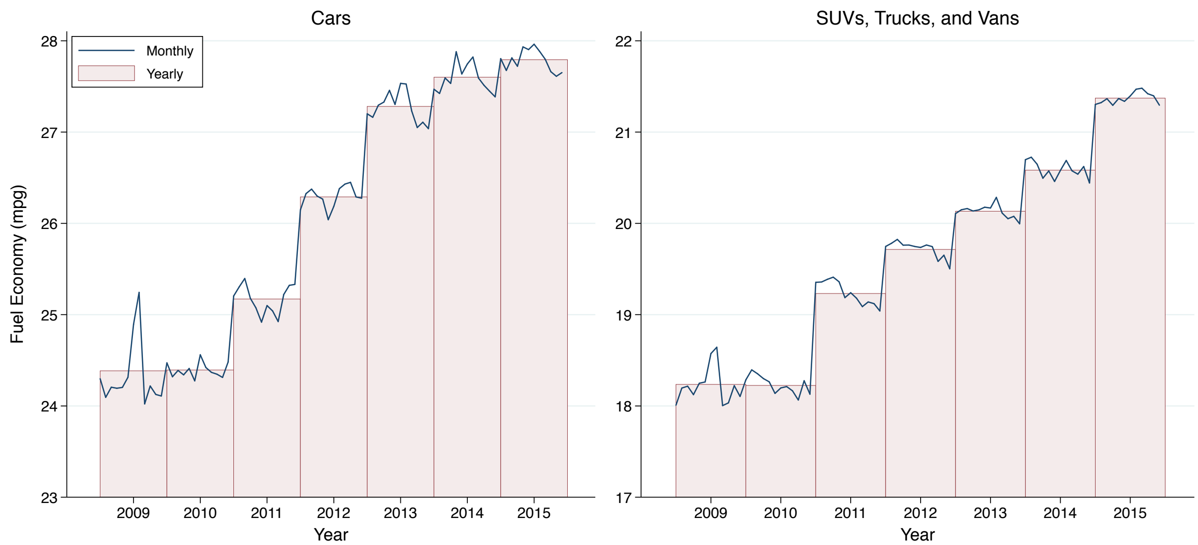

Figure A7: Sales-weighted fuel economy

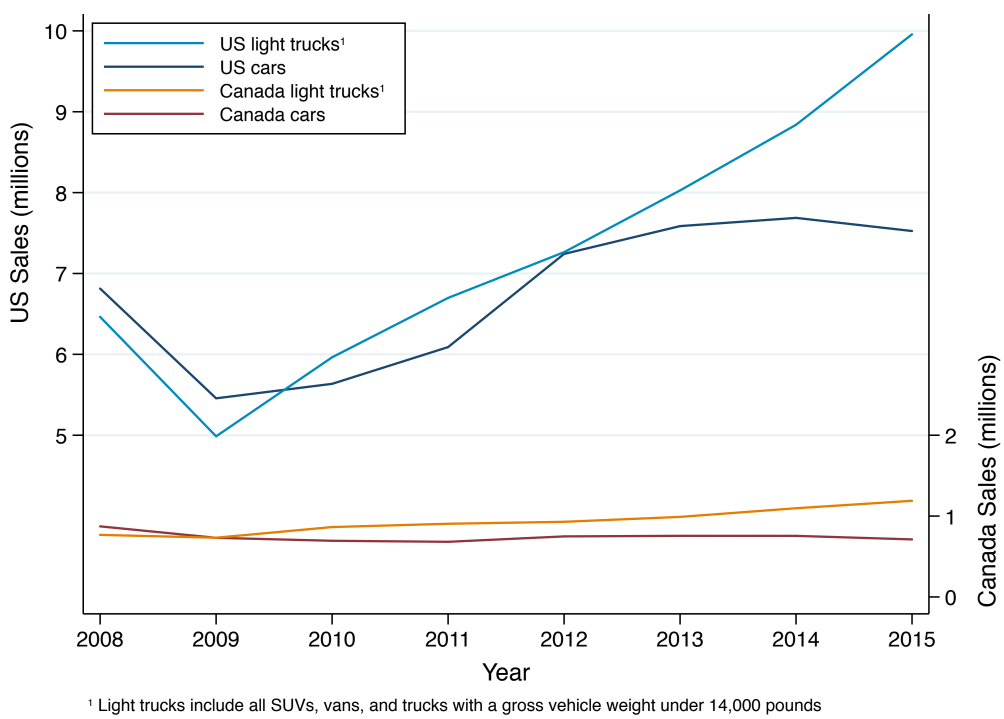

Figure A8: US and Canadian vehicle sales

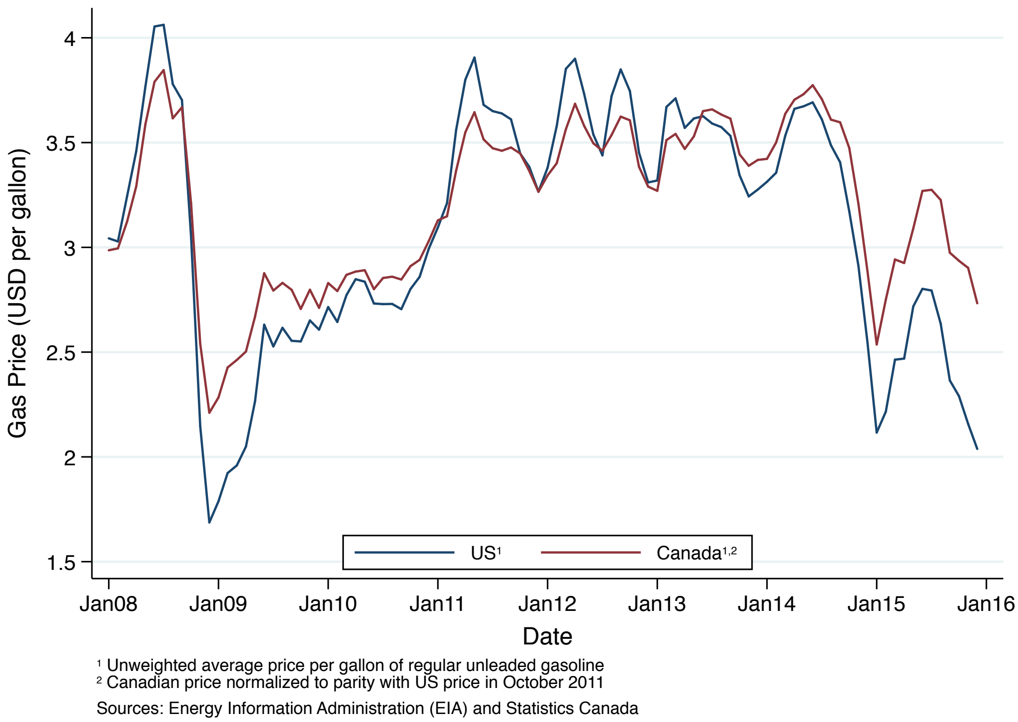

Figure A9: US and Canadian gas prices

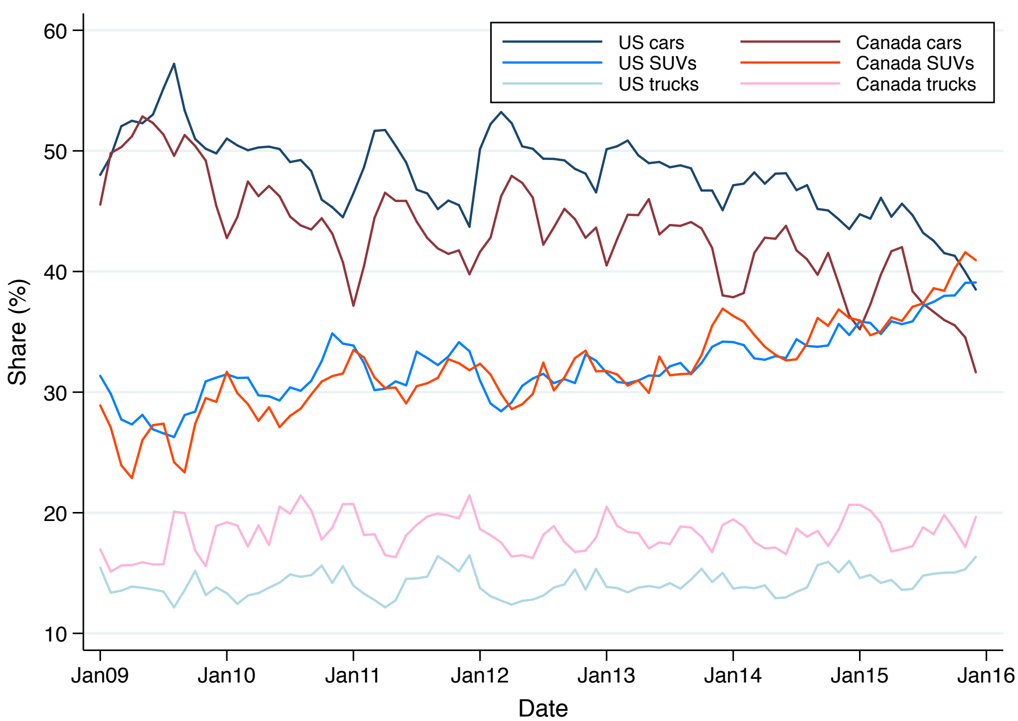

Figure A10: Monthly market share trends

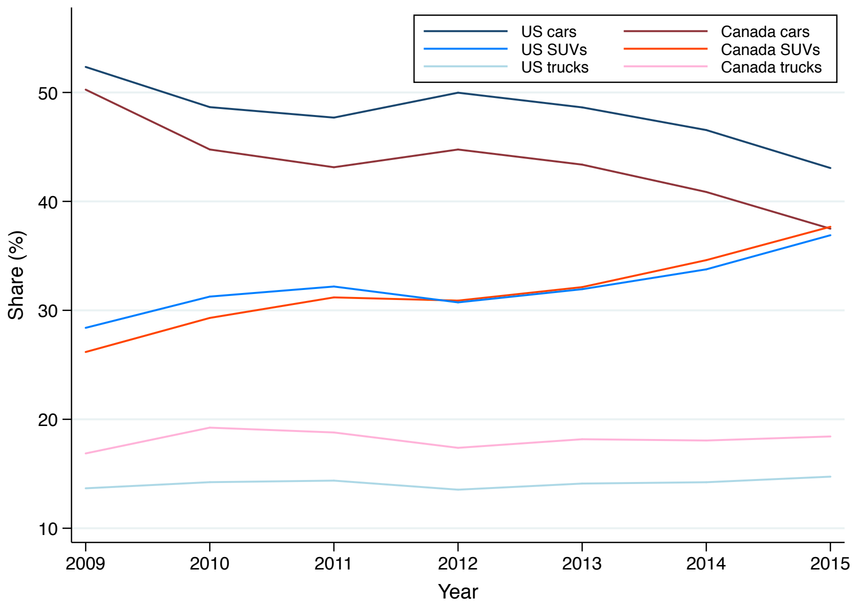

Figure A11: Yearly market share trends

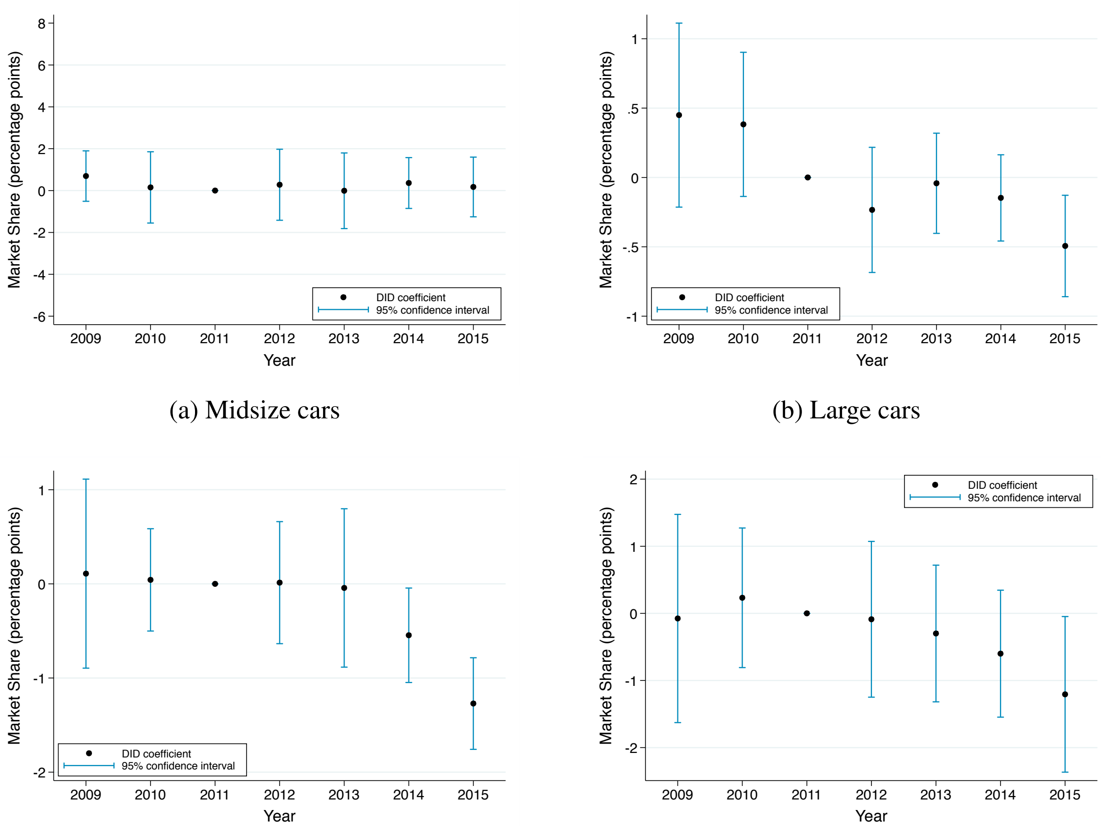

(a) Midsize cars

(b) Large cars

(c) Luxury cars

(d) Luxury SUVs

Figure A12: Market share flexible difference-in-differences

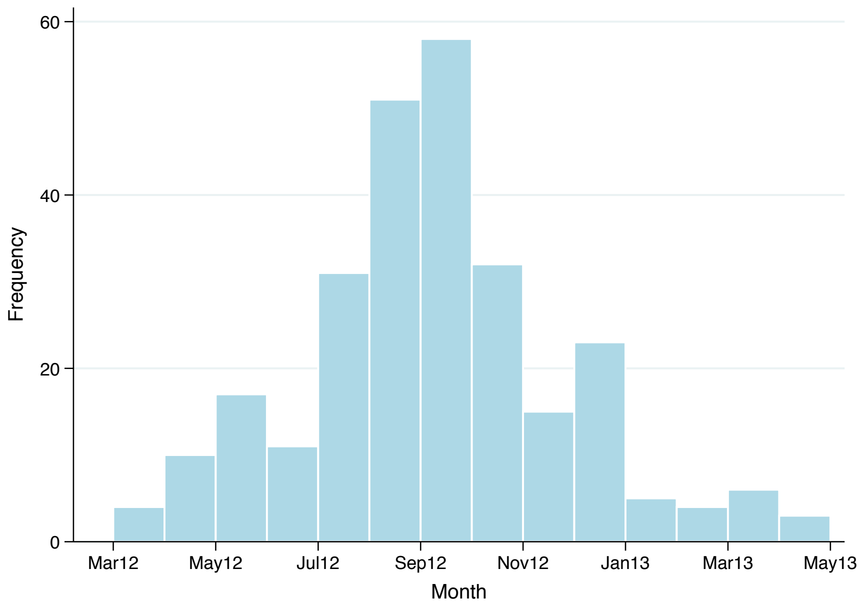

Figure A13: Monthly releases distribution of 2013 model-years

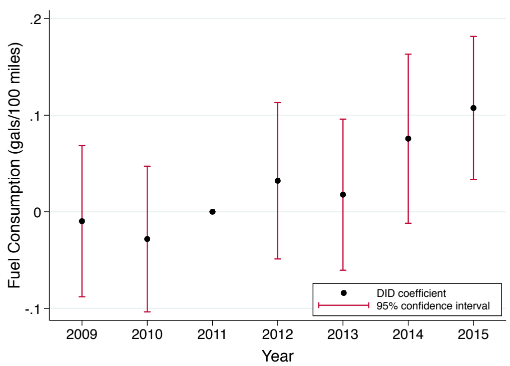

(a) Midsize cars

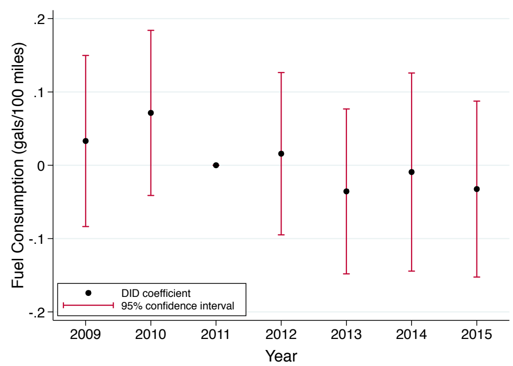

(b) Large cars

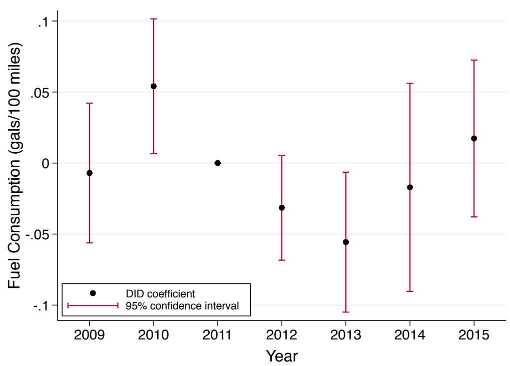

(c) Luxury cars

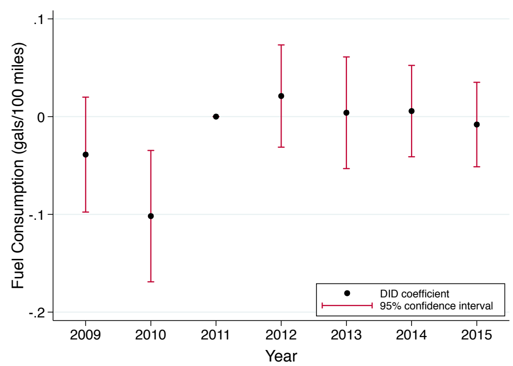

(d) Luxury SUVs

Figure A14: Fuel consumption flexible difference-in-differences

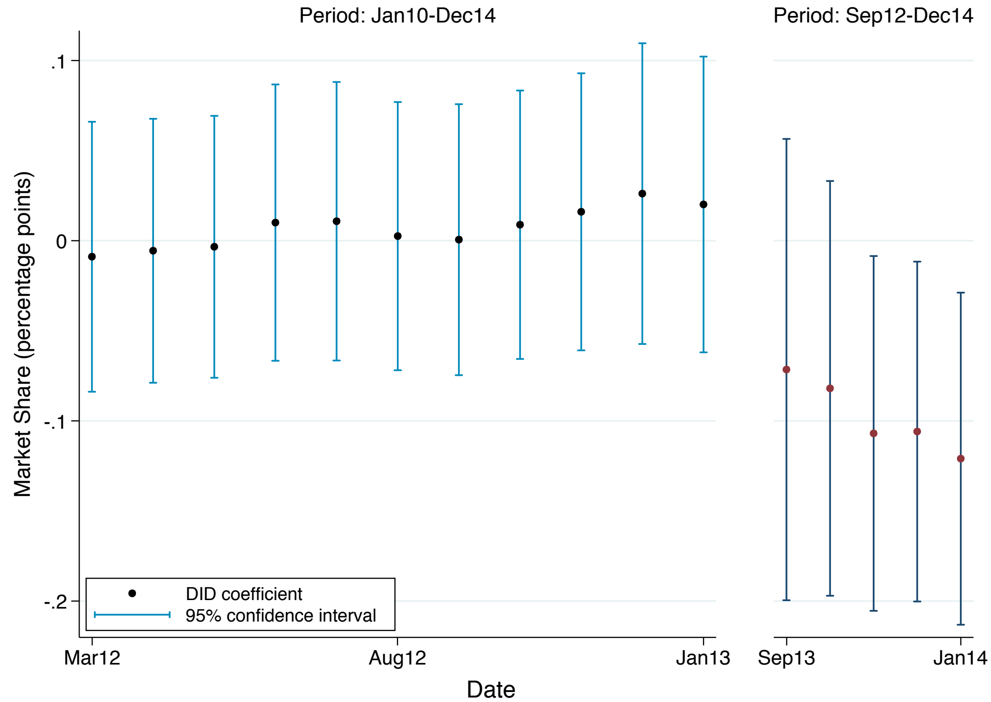

Figure A15: Luxury hybrid incentives difference-in-differences

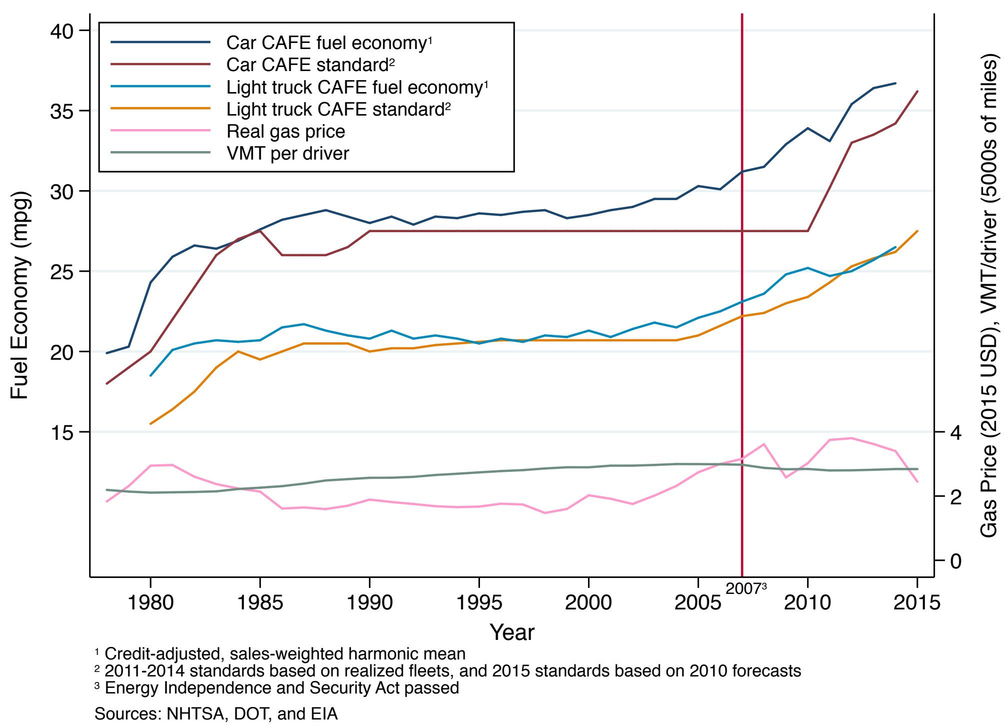

Figure A16: CAFE standards, fuel economy, and demand factors

Table A1: Market share difference-in-differences (2010-2014 and 2009-2014)

	2010-2014	w/ Gas Price	2009-2014	w/ Gas Price
US × Small × Treat	1.44**	1.54**	2.38***	2.41***
	(2.33)	(2.33)	(3.24)	(3.72)
US × Midsize × Treat	-0.42	-0.25	-0.36	-0.36
	(-1.01)	(-0.72)	(-0.95)	(-0.99)
US × Large × Treat	-0.19	-0.32**	-0.32**	-0.32***
	(-1.35)	(-2.27)	(-2.34)	(-2.73)
US × Luxury × Treat	-0.20	-0.31**	-0.29	-0.29
	(-1.51)	(-2.47)	(-1.63)	(-1.61)
US × Sport × Treat	0.08	0.07	0.15	0.15
	(0.62)	(0.49)	(1.07)	(1.08)
US × Small SUV × Treat	-1.00***	-1.16***	-1.11***	-1.14***
	(-5.72)	(-6.76)	(-5.76)	(-7.13)
US × Midsize SUV × Treat	-0.00	0.16	-0.08	-0.08
	(-0.00)	(0.42)	(-0.22)	(-0.22)
US × Large SUV × Treat	-0.50***	-0.51***	-0.58***	-0.57***
	(-3.45)	(-3.35)	(-3.75)	(-3.66)
US × Luxury SUV × Treat	-0.28	-0.38**	-0.38**	-0.38**
	(-1.48)	(-2.10)	(-2.19)	(-2.24)
US × Truck × Treat	0.86**	0.98*	0.38	0.33
	(2.31)	(1.94)	(0.83)	(0.88)
Observations	1320	1320	1584	1584
Within-R^2	0.97	0.97	0.96	0.96
SE Clusters	30	30	36	36

Notes: t statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Coefficient estimates in percentage points. Standard errors are clustered in country by 4-month blocks to allow for both within and across-segment serial correlation in errors within country-trimester. The effect of gas prices on segment shares is controlled for in the 2nd and 4th columns.

Table A2: Fuel consumption difference-in-differences, country×segment×2-year clustered SEs

	Full Sample	w/ Gas Price	2010-2014	w/ Gas Price
US × Small × Treat	-0.01	-0.01	-0.02*	-0.01
	(-1.19)	(-1.03)	(-1.93)	(-1.59)
US × Midsize × Treat	0.07***	0.06***	0.05***	0.05***
	(2.97)	(2.80)	(2.88)	(2.81)
US × Large × Treat	-0.05**	-0.05**	-0.05	-0.05*
	(-2.36)	(-2.56)	(-1.51)	(-1.77)
US × Luxury × Treat	-0.03	-0.04***	-0.05**	-0.07***
	(-1.53)	(-7.35)	(-2.46)	(-3.70)
US × Sport × Treat	-0.13	-0.18*	-0.01	-0.04
	(-1.12)	(-1.88)	(-0.20)	(-0.70)
US × Small SUV × Treat	-0.23***	-0.24***	-0.25***	-0.26***
	(-6.53)	(-8.43)	(-4.99)	(-7.15)
US × Midsize SUV × Treat	-0.00	-0.00	0.01	0.01
	(-0.30)	(-0.27)	(1.01)	(0.63)
US × Large SUV × Treat	0.04	0.03	0.08**	0.06***
	(0.95)	(0.84)	(2.63)	(3.95)
US × Luxury SUV × Treat	0.04*	0.05***	0.05	0.07**
	(1.88)	(2.92)	(1.32)	(2.41)
US × Truck × Treat	-0.06	-0.09***	-0.07**	-0.08***
	(-1.64)	(-4.91)	(-2.13)	(-2.83)
US × Van × Treat	-0.09	-0.06	-0.01	-0.01
	(-1.36)	(-1.17)	(-1.11)	(-0.96)
Observations	1848	1848	1320	1320
Within-R^2	0.67	0.73	0.70	0.72
SE Clusters	88	88	66	66

Notes: t statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Coefficient estimates in gallons per 100 miles. Standard errors are clustered in country by segment by 2-year blocks to allow for biennial serial correlation in errors within each country-segment. The effect of gas prices on segment-average fuel consumption is controlled for in the 2nd and 4th columns, and is found to be negative for most segments, or statistically indistinguishable from 0 otherwise.

Table A3: Market share linear probability models

	US Base	US Full	Can Base	Can Full	US Base	US Full	Can Base	Can Full
Small × Fuel Consumption	-0.09	-0.09	-0.53***	-0.48***	-0.11	-0.11	-0.09	-0.09
	(-0.95)	(-0.95)	(-2.98)	(-2.79)	(-0.88)	(-0.89)	(-0.91)	(-0.91)
Midsize × Fuel Consumption	-0.30**	-0.33***	-0.19***	-0.17***	-0.11	-0.11	0.07	0.07
	(-2.52)	(-2.88)	(-4.23)	(-3.94)	(-1.32)	(-1.32)	(1.31)	(1.50)
Large × Fuel Consumption	-0.03	-0.04	0.01	0.00	-0.36***	-0.40***	-0.11***	-0.14***
	(-0.47)	(-0.62)	(0.41)	(0.00)	(-4.29)	(-5.13)	(-2.96)	(-3.95)
Luxury × Fuel Consumption	-0.02	-0.02	0.00	0.01	-0.05**	-0.05**	-0.02	-0.02
	(-1.48)	(-1.50)	(0.35)	(0.59)	(-2.43)	(-2.46)	(-1.14)	(-1.18)
Sport × Fuel Consumption	0.05***	0.05**	0.01	0.01	-0.02**	-0.02**	-0.03***	-0.03***
	(2.62)	(2.33)	(0.85)	(0.57)	(-2.40)	(-2.57)	(-2.92)	(-2.98)
Small SUV × Fuel Consumption	0.16***	0.16***	0.08	0.07	-0.08	-0.09	-0.70***	-0.70***
	(3.06)	(3.09)	(1.24)	(1.03)	(-0.45)	(-0.50)	(-2.99)	(-2.96)
Midsize SUV × Fuel Consumption	-0.36***	-0.38***	-0.47***	-0.49***	-0.32***	-0.32***	-0.49***	-0.49***
	(-4.84)	(-5.20)	(-5.07)	(-5.18)	(-4.73)	(-4.95)	(-4.61)	(-4.66)
Large SUV × Fuel Consumption	-0.18***	-0.19***	-0.07***	-0.07***	-0.06**	-0.06**	-0.02	-0.03*
	(-6.72)	(-6.89)	(-4.37)	(-4.43)	(-2.36)	(-2.40)	(-1.53)	(-1.71)
Luxury SUV × Fuel Consumption	-0.08***	-0.08***	-0.05***	-0.05***	-0.02**	-0.02**	-0.03***	-0.03***
	(-4.20)	(-3.97)	(-3.15)	(-3.12)	(-2.20)	(-2.20)	(-3.52)	(-3.54)
Truck × Fuel Consumption	-0.19	-0.19	-0.20	-0.18	0.05	0.05	0.10*	0.11*
	(-1.57)	(-1.46)	(-1.30)	(-1.15)	(1.62)	(1.60)	(1.80)	(1.84)
Van × Fuel Consumption	-0.03	-0.05	0.09	0.06	0.07	0.07	0.01	0.03
	(-0.59)	(-0.88)	(1.28)	(0.74)	(1.16)	(1.23)	(0.08)	(0.35)
Small × Fuel Consumption × Post		0.00		-0.12		0.01		-0.01
		(0.06)		(-1.44)		(1.11)		(-0.49)
Midsize × Fuel Consumption × Post		0.05		-0.03*		0.00		-0.02**
		(0.83)		(-1.65)		(0.09)		(-2.12)
Large × Fuel Consumption × Post		0.02**		0.03***		0.03***		0.02**
		(2.16)		(2.61)		(4.13)		(2.07)
Luxury × Fuel Consumption × Post		0.00		-0.00		0.00		0.00
		(0.25)		(-0.65)		(0.98)		(0.77)
Sport × Fuel Consumption × Post		0.01		0.01		0.00		0.00*
		(1.22)		(1.47)		(1.17)		(1.68)
Small SUV × Fuel Consumption × Post		-0.01		0.04*		0.00		-0.00
		(-0.88)		(1.73)		(0.63)		(-0.65)
Midsize SUV × Fuel Consumption × Post		0.07**		0.05		0.02***		-0.00
		(2.21)		(1.10)		(2.97)		(-0.25)
Large SUV × Fuel Consumption × Post		0.01		0.00		0.00		0.01**
		(1.10)		(0.25)		(0.53)		(1.97)
Luxury SUV × Fuel Consumption × Post		-0.00		-0.00		0.00		-0.00
		(-0.03)		(-0.01)		(0.46)		(-0.61)
Truck × Fuel Consumption × Post		-0.02		-0.08		-0.00		0.01*
		(-0.31)		(-0.91)		(-0.36)		(1.68)
Van × Fuel Consumption × Post		0.04		0.08*		-0.00		-0.01
		(1.42)		(1.73)		(-0.57)		(-1.53)
Observations	14352	14352	14232	14232	14352	14352	14232	14232
Within-R^2	0.18	0.18	0.28	0.29	0.06	0.07	0.10	0.10
Vehicle Fixed Effects					Yes	Yes	Yes	Yes

Notes: t statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Dependent variable (share) in percentage points, consumption in gals/100 miles. All models include segment by month of sample fixed effects and control for vehicle availability and interactions of segment by price and power/weight. SEs clustered by segment-month and model-year.

Table A4: Market share nested logit models with BLP instruments

	US Base	US Full	Can Base	Can Full	US Base	US Full	Can Base	Can Full
Small × Fuel Consumption	0.09	0.07	0.39***	0.30***	-0.14	-0.14*	0.26*	0.08
	(1.60)	(1.41)	(3.90)	(3.05)	(-1.54)	(-1.71)	(1.93)	(0.69)
Midsize × Fuel Consumption	0.19***	0.08	0.06	0.04	0.03	-0.01	-0.02	-0.09
	(3.21)	(1.46)	(0.80)	(0.65)	(0.48)	(-0.18)	(-0.21)	(-1.21)
Large × Fuel Consumption	0.21**	0.14**	0.25*	0.11	0.19**	0.10	0.51**	0.15
	(2.34)	(2.16)	(1.79)	(1.40)	(2.10)	(1.35)	(2.44)	(0.96)
Luxury × Fuel Consumption	0.11***	0.07***	0.13***	0.11***	0.07	-0.03	0.23***	0.00
	(2.84)	(3.55)	(2.90)	(3.40)	(0.95)	(-0.42)	(2.65)	(0.05)
Sport × Fuel Consumption	-0.23***	-0.20***	-0.31***	-0.27***	-0.26***	-0.26***	-0.22***	-0.29***
	(-6.25)	(-7.46)	(-5.32)	(-7.00)	(-4.32)	(-4.36)	(-2.70)	(-3.76)
Small SUV × Fuel Consumption	-0.06	-0.02	-0.13**	-0.06*	-0.17**	-0.15*	-0.13*	-0.13**
	(-1.11)	(-0.39)	(-2.50)	(-1.78)	(-2.30)	(-1.79)	(-1.83)	(-2.01)
Midsize SUV × Fuel Consumption	-0.17***	-0.13*	-0.08	-0.02	-0.30***	-0.33***	-0.12	-0.12
	(-2.71)	(-1.94)	(-0.94)	(-0.24)	(-3.70)	(-3.15)	(-1.19)	(-1.04)
Large SUV × Fuel Consumption	0.20**	0.19**	-0.02	0.14	0.08	0.03	-0.04	0.08
	(2.40)	(2.43)	(-0.17)	(1.33)	(1.00)	(0.38)	(-0.34)	(0.71)
Luxury SUV × Fuel Consumption	-0.07***	0.00	-0.12***	-0.01	-0.21***	-0.19***	-0.13*	-0.05
	(-3.13)	(0.05)	(-3.30)	(-0.25)	(-3.56)	(-2.96)	(-1.73)	(-0.64)
Truck × Fuel Consumption	0.14***	0.17***	0.14**	0.21***	0.05	-0.00	0.11	0.14*
	(4.00)	(4.63)	(2.25)	(3.49)	(0.77)	(-0.00)	(1.41)	(1.83)
Van × Fuel Consumption	0.00	0.02	-0.04	-0.02	-0.05	-0.10**	0.02	-0.02
	(0.10)	(1.05)	(-0.94)	(-0.56)	(-1.41)	(-2.13)	(0.36)	(-0.51)
Small × Fuel Consumption × Post		-0.02**		-0.05***		0.00		-0.05*
		(-2.01)		(-3.16)		(0.09)		(-1.90)
Midsize × Fuel Consumption × Post		-0.05***		-0.04**		-0.01		-0.03
		(-3.89)		(-2.36)		(-0.40)		(-1.07)
Large × Fuel Consumption × Post		-0.05***		-0.09***		-0.03		-0.09***
		(-4.36)		(-4.28)		(-1.51)		(-2.96)
Luxury × Fuel Consumption × Post		-0.04***		-0.05***		-0.02		-0.04*
		(-4.79)		(-4.10)		(-1.31)		(-1.86)
Sport × Fuel Consumption × Post		-0.01		-0.02		0.00		-0.01
		(-1.64)		(-1.14)		(0.23)		(-0.54)
Small SUV × Fuel Consumption × Post		-0.02*		-0.01		0.01		0.00
		(-1.72)		(-0.78)		(0.29)		(0.03)
Midsize SUV × Fuel Consumption × Post		-0.02*		-0.02		0.00		-0.01
		(-1.74)		(-1.35)		(0.02)		(-0.77)
Large SUV × Fuel Consumption × Post		-0.02**		-0.00		-0.01		-0.00
		(-2.39)		(-0.10)		(-0.80)		(-0.13)
Luxury SUV × Fuel Consumption × Post		0.00		0.01		0.01		0.03
		(0.03)		(1.22)		(0.76)		(1.33)
Truck × Fuel Consumption × Post		-0.01		-0.02		0.01		-0.01
		(-1.56)		(-1.46)		(0.67)		(-0.56)
Van × Fuel Consumption × Post		-0.03***		-0.04***		-0.01		-0.03*
		(-5.19)		(-3.72)		(-0.96)		(-1.75)
Within-R^2	0.99	0.99	0.98	0.99	0.99	0.99	0.98	0.98
Month of Sample Fixed Effects	Yes	Yes	Yes	Yes				
Make by Month of Sample Fixed Effects					Yes	Yes	Yes	Yes

Notes: t statistics in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Numbers of observations suppressed for space but similar to Table A3. Dependent variable is log(share), and consumption in gals/100 miles. Specifications instrument for within-segment share and interactions of segment by consumption, price, and power/weight with segment-year and make-year means. SEs clustered by segment-month and model-year.

Appendix B

Chapter 3 Figures and Tables

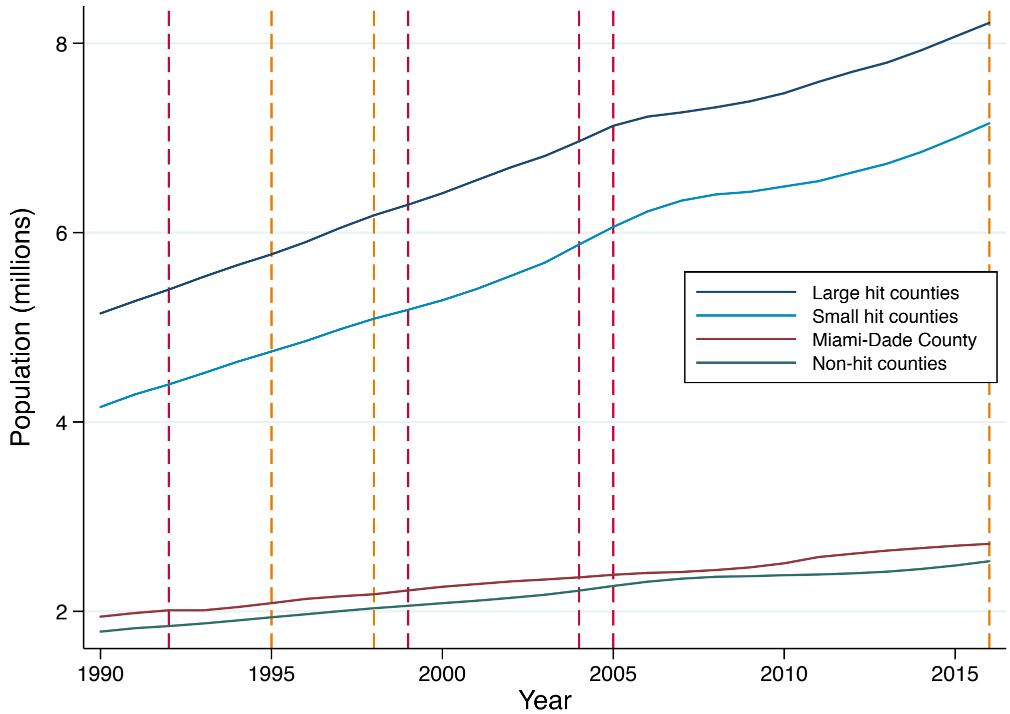

Figure B1: Florida population by county aggregation. Red dashed lines denote hurricane years in which at least 1/3 of hit county populations were affected (including Miami-Dade) according to the census tract hit definition, with the exception of Andrew in 1992, which affected all of Miami-Dade, but respectively only 17% and 5% of the other large and small county populations. Orange dashed lines represent all other hurricane years. The 1995 and 1998 hurricanes only affected populations in the group of small counties (around 18% in both years). *Source:* Authors' calculations based on data from the US Census Bureau.

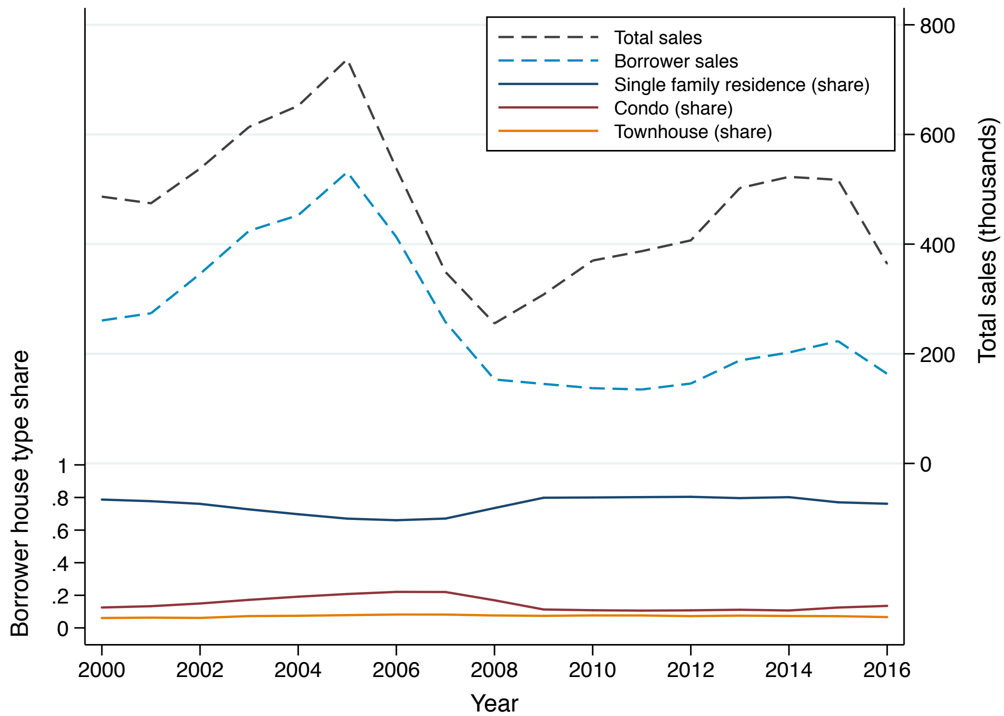

Figure B2: Florida borrowers market sales and composition

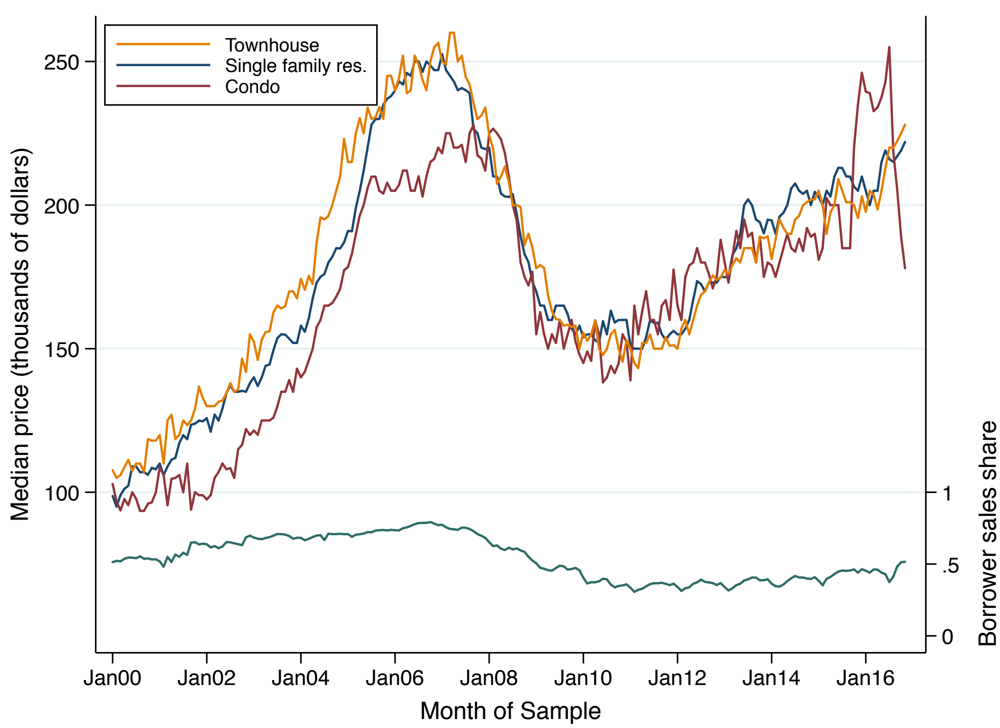

Figure B3: Florida borrowers market prices and sales share

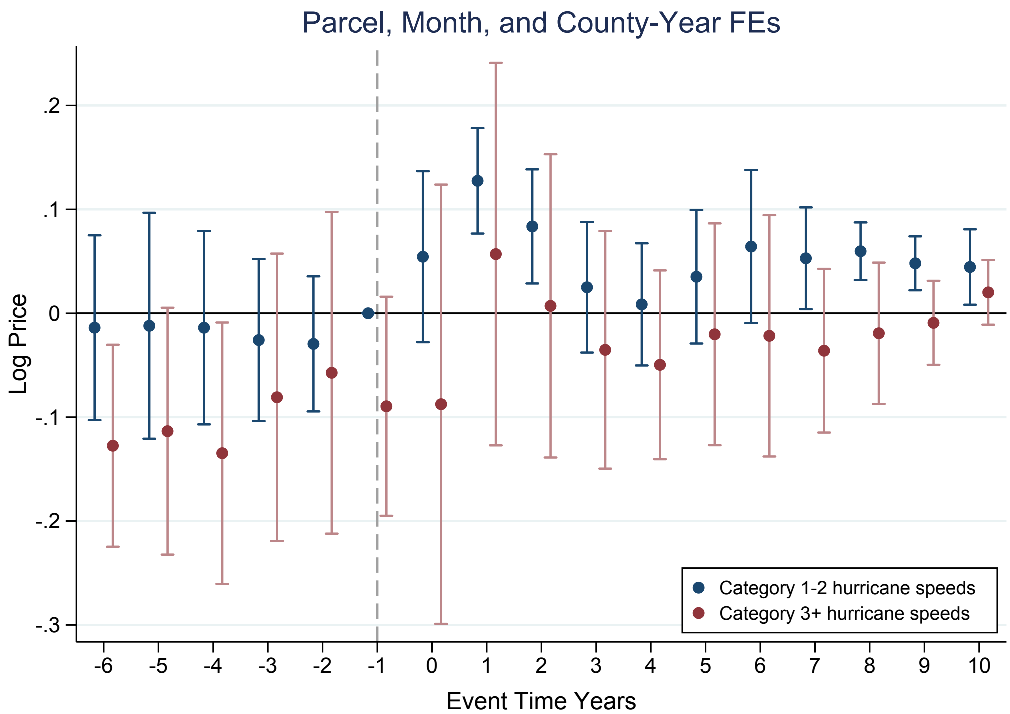

Figure B4: Heterogeneous effects of differential hurricane intensity – repeated sales

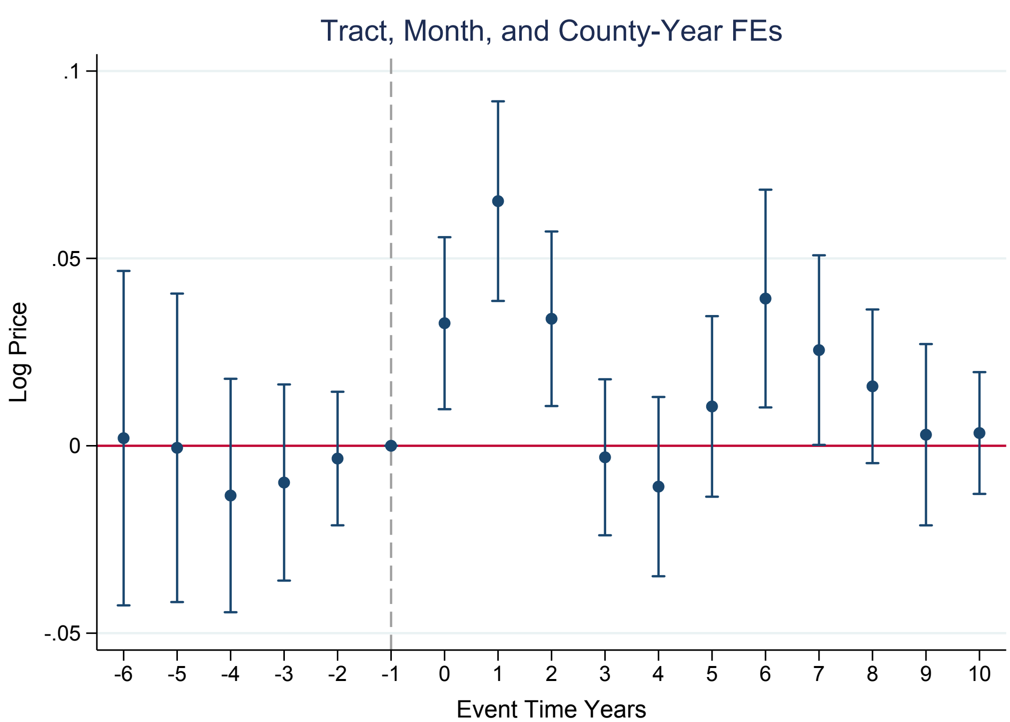

Figure B5: hurricane effects on house prices – borrower sample

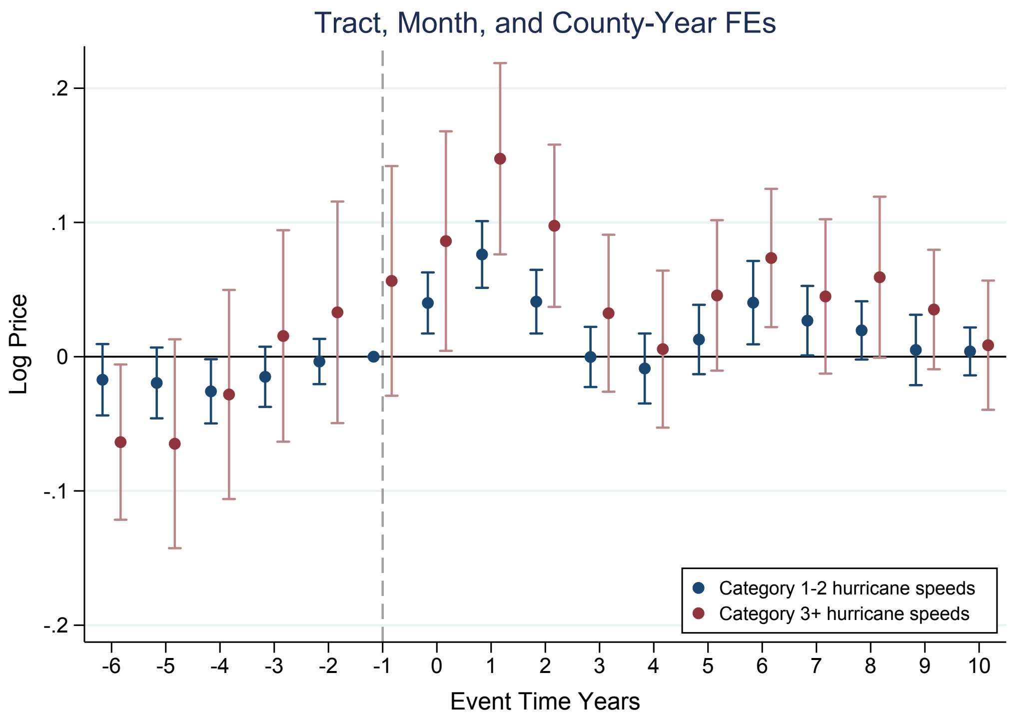

Figure B6: Heterogeneous effects of differential hurricane intensity – borrower sample

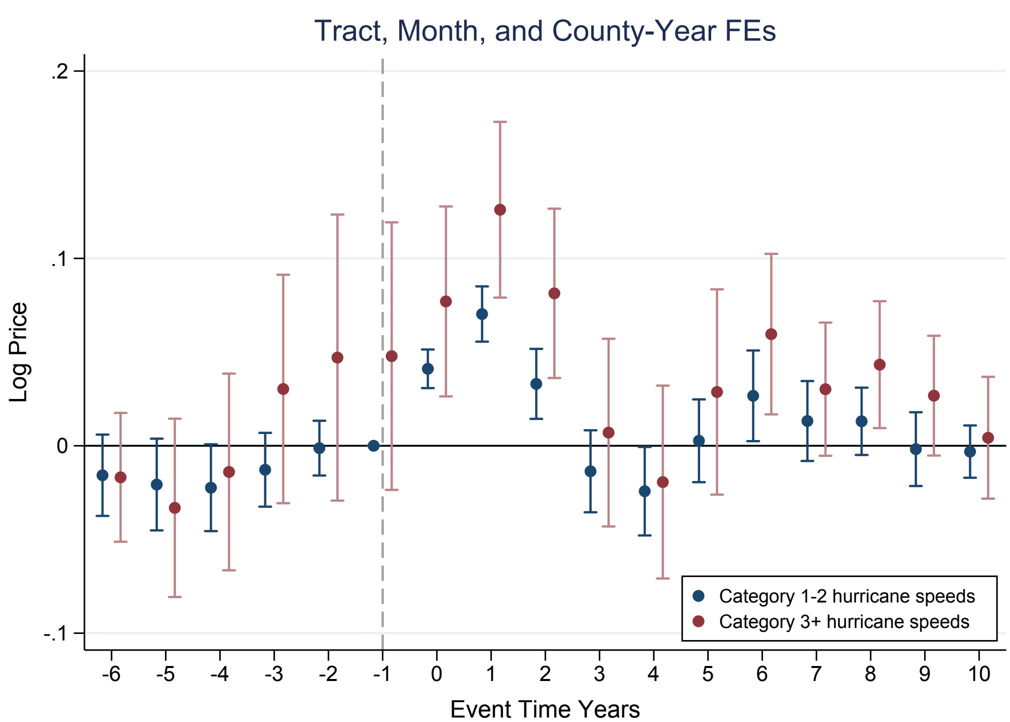

Figure B7: Heterogeneous price effects of differential hurricane intensity – HMDA sample

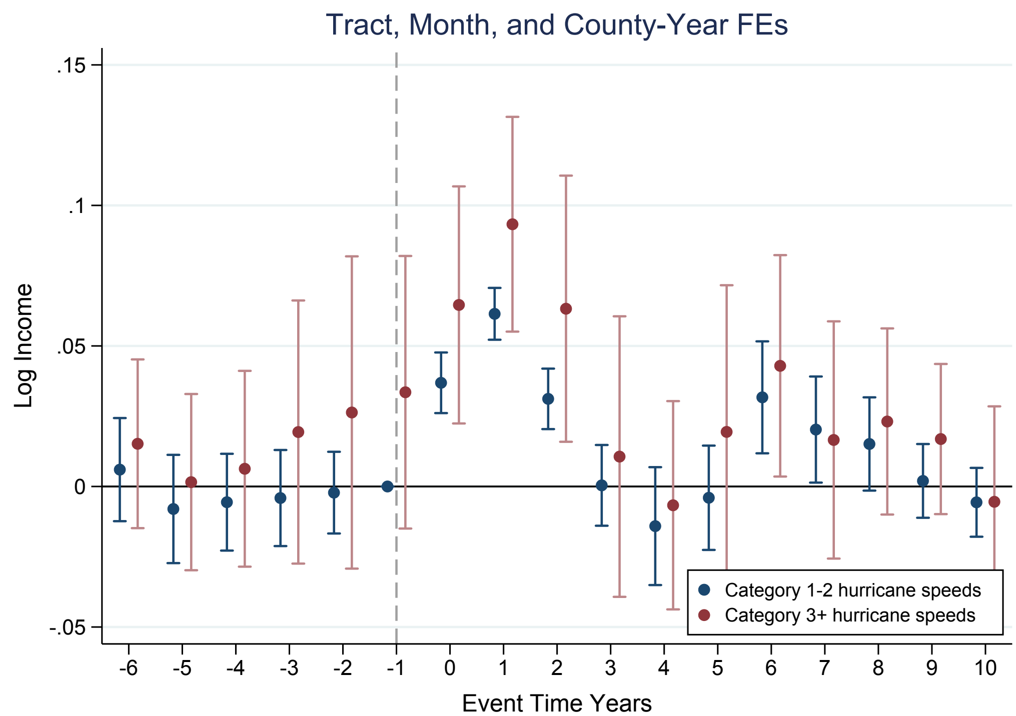

Figure B8: Heterogeneous income effects of differential hurricane intensity – HMDA sample

Table B1: Wind speed and maximal reach radius model

	log(Maxradius)
Speed	-0.0224***
	(0.004)
Speed^2	-0.0002***
	(0.00004)
Hurricane-track-point FEs	Yes
N	1188
R^2	0.93
Within-R^2	0.90

Notes: Standard errors in parentheses (clustered at the hurricane-track-point level).
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table B2: Percentage of houses sold from hit areas by hurricane, years pre

	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1
Andrew92	.	.	.	.	.	.	.	.	.	.
Opal95	.	.	.	.	.	.	.	.	.	.
Earl98	.	.	.	.	.	.	.	.	.	.
Georges98	.	.	.	.	.	.	.	.	.	.
Irene99	.	.	.	.	.	.	.	.	.	.
Charley04	.	.	.	.	.	1.3	2.0	2.1	2.3	2.5
Frances04	.	.	.	.	.	1.0	1.4	1.6	1.8	2.1
Jeanne04	.	.	.	.	.	1.1	1.4	1.6	1.8	1.9
Ivan04	.	.	.	.	.	0.1	0.1	0.2	0.2	0.2
Dennis05	.	.	.	.	0.1	0.1	0.1	0.1	0.2	0.2
Katrina05	.	.	.	.	0.5	0.8	0.8	0.8	0.9	1.0
Wilma05	.	.	.	.	2.6	3.1	3.4	3.8	4.1	4.6
Hermine16	0.1	0.1	0.0	0.0	0.0	0.0	0.1	0.1	0.1	0.1
Matthew16	0.6	0.5	0.5	0.6	0.6	0.7	0.9	1.0	1.0	0.7
Irene17	1.1	1.3	1.6	1.7	1.7	2.0	2.1	2.3	1.9	0.3
Total	1.8	1.8	2.2	2.3	5.0	8.5	10.2	11.3	11.8	10.9

Notes: House hit definition described in Section 3.2.4. Periods denote hurricane event years not covered by the 2000-2016 sample.

Table B3: Percentage of houses sold from hit areas by hurricane, years post

	0	1	2	3	4	5	6	7	8	9	10
Andrew92	.	.	.	.	.	.	.	1.0	1.5	1.4	1.4
Opal95	.	.	.	.	0.2	0.3	0.3	0.4	0.4	0.4	0.3
Earl98	.	0.2	0.3	0.3	0.4	0.4	0.4	0.3	0.2	0.2	0.2
Georges98	.	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Irene99	2.8	3.4	3.6	4.0	4.2	4.6	3.7	2.4	1.6	1.8	2.5
Charley04	2.5	2.1	1.4	0.9	0.9	1.4	1.5	1.5	1.8	1.8	1.9
Frances04	2.3	2.0	1.3	0.8	0.9	1.2	1.1	1.2	1.5	1.6	1.7
Jeanne04	2.0	1.5	1.0	0.7	0.7	1.0	1.1	1.1	1.3	1.4	1.4
Ivan04	0.2	0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Dennis05	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Katrina05	0.9	0.6	0.4	0.3	0.6	0.6	0.6	0.7	0.7	0.7	0.6
Wilma05	3.4	2.2	1.5	1.8	2.5	2.6	2.6	3.0	3.1	3.1	2.3
Hermine16	0.0	.	.	.	.	.	.	.	.	.	.
Matthew16	0.1	.	.	.	.	.	.	.	.	.	.
Irene17	.	.	.	.	.	.	.	.	.	.	.
Total	11.7	10.3	8.4	8.1	9.3	10.7	9.9	10.2	10.5	10.8	10.7

Notes: House hit definition defined in Section 3.2.4. Periods denote hurricane event years not covered by the 2000-2016 sample.

Table B4: Percentage of houses sold from category-3-speed hit areas by hurricane, years pre

	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1
Andrew92	.	.	.	.	.	.	.	.	.	.
Opal95	.	.	.	.	.	.	.	.	.	.
Earl98	.	.	.	.	.	.	.	.	.	.
Georges98	.	.	.	.	.	.	.	.	.	.
Irene99	.	.	.	.	.	.	.	.	.	.
Charley04	.	.	.	.	.	.	.	.	.	.
Frances04	.	.	.	.	.	0.0	0.1	0.1	0.1	0.1
Jeanne04	.	.	.	.	.	0.1	0.1	0.2	0.3	0.2
Ivan04	.	.	.	.	.	0.0	0.0	0.0	0.0	0.0
Dennis05	.	.	.	.	.	.	.	.	.	.
Katrina05	.	.	.	.	.	.	.	.	.	.
Wilma05	.	.	.	.	0.3	0.3	0.4	0.4	0.4	0.4
Hermine16	.	.	.	.	.	.	.	.	.	.
Matthew16	.	.	.	.	.	.	.	.	.	.
Irene17	0.1	0.1	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.0
Total	0.1	0.1	0.2	0.2	0.5	0.6	0.8	0.9	0.9	0.8

Notes: House hit definition described in Section 3.2.4. Periods denote hurricanes not reaching category 3 speeds in Florida or event years not covered by sample.

Table B5: Percentage of houses sold from category-3-speed hit areas by hurricane, years post

	0	1	2	3	4	5	6	7	8	9	10
Andrew92	.	.	.	.	.	.	.	0.3	0.5	0.4	0.4
Opal95	.	.	.	.	.	.	.	.	.	.	.
Earl98	.	.	.	.	.	.	.	.	.	.	.
Georges98	.	.	.	.	.	.	.	.	.	.	.
Irene99	.	.	.	.	.	.	.	.	.	.	.
Charley04	.	.	.	.	.	.	.	.	.	.	.
Frances04	0.1	0.1	0.0	0.0	0.0	0.1	0.1	0.1	0.1	0.1	0.1
Jeanne04	0.2	0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.2	0.1
Ivan04	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Dennis05	.	.	.	.	.	.	.	.	.	.	.
Katrina05	.	.	.	.	.	.	.	.	.	.	.
Wilma05	0.3	0.2	0.2	0.2	0.2	0.3	0.3	0.3	0.3	0.3	0.2
Hermine16	.	.	.	.	.	.	.	.	.	.	.
Matthew16	.	.	.	.	.	.	.	.	.	.	.
Irene17	.	.	.	.	.	.	.	.	.	.	.
Total	0.7	0.5	0.3	0.3	0.4	0.4	0.4	0.8	1.0	1.0	0.9

Notes: House hit definition described in Section 3.2.4. Periods denote hurricanes not reaching category 3 speeds in Florida or event years not covered by sample.

Table B6: Percentage of hit tracts by hurricane, years pre

	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1
Andrew92	.	.	.	.	.	.	.	.	.	.
Opal95	.	.	.	.	.	.	.	.	.	.
Earl98	.	.	.	.	.	.	.	.	.	.
Georges98	.	.	.	.	.	.	.	.	.	.
Irene99	.	.	.	.	.	.	.	.	.	.
Charley04	.	.	.	.	.	1.0	1.5	1.5	1.5	1.5
Frances04	.	.	.	.	.	1.0	1.5	1.5	1.5	1.5
Jeanne04	.	.	.	.	.	0.9	1.2	1.2	1.2	1.2
Ivan04	.	.	.	.	.	0.1	0.2	0.2	0.2	0.2
Dennis05	.	.	.	.	0.1	0.1	0.1	0.1	0.1	0.1
Katrina05	.	.	.	.	0.4	0.7	0.7	0.7	0.7	0.7
Wilma05	.	.	.	.	2.1	2.5	2.5	2.5	2.5	2.5
Hermine16	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Matthew16	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7
Irene17	2.2	2.2	2.2	2.2	2.2	2.2	2.2	2.2	2.2	0.7
Total	3.0	3.0	3.0	3.0	5.2	8.0	9.0	9.0	9.0	7.5

Notes: Census tract hit definition characterized in Section 3.2.3. Periods denote hurricane event years not covered by the 2000-2016 sample.

Table B7: Percentage of hit tracts by hurricane, years post

	0	1	2	3	4	5	6	7	8	9	10
Andrew92	.	.	.	.	.	.	.	0.8	1.1	1.1	1.1
Opal95	.	.	.	.	0.2	0.3	0.3	0.3	0.3	0.3	0.3
Earl98	.	0.2	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3
Georges98	.	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Irene99	2.5	3.0	3.0	3.0	3.0	3.0	3.0	3.0	3.0	3.0	3.0
Charley04	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5
Frances04	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5
Jeanne04	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.2
Ivan04	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2	0.2
Dennis05	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Katrina05	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7
Wilma05	2.5	2.5	2.5	2.5	2.5	2.5	2.5	2.5	2.5	2.5	2.5
Hermine16	0.0	.	.	.	.	.	.	.	.	.	.
Matthew16	0.2	.	.	.	.	.	.	.	.	.	.
Irene17	.	.	.	.	.	.	.	.	.	.	.
Total	8.7	9.2	9.3	9.3	9.5	9.6	9.6	10.3	10.7	10.7	10.7

Notes: Census tract hit definition characterized in Section 3.2.3. Periods denote hurricane event years not covered by the 2000-2016 sample.

Table B8: Percentage of category-3-speed hit tracts by hurricane, years pre

	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1
Andrew92	.	.	.	.	.	.	.	.	.	.
Opal95	.	.	.	.	.	.	.	.	.	.
Earl98	.	.	.	.	.	.	.	.	.	.
Georges98	.	.	.	.	.	.	.	.	.	.
Irene99	.	.	.	.	.	.	.	.	.	.
Charley04	.	.	.	.	.	.	.	.	.	.
Frances04	.	.	.	.	.	0.0	0.0	0.0	0.0	0.0
Jeanne04	.	.	.	.	.	0.1	0.1	0.1	0.1	0.1
Ivan04	.	.	.	.	.	0.0	0.0	0.0	0.0	0.0
Dennis05	.	.	.	.	.	.	.	.	.	.
Katrina05	.	.	.	.	.	.	.	.	.	.
Wilma05	.	.	.	.	0.2	0.3	0.3	0.3	0.3	0.3
Hermine16	.	.	.	.	.	.	.	.	.	.
Matthew16	.	.	.	.	.	.	.	.	.	.
Irene17	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.0
Total	0.1	0.1	0.1	0.1	0.4	0.5	0.5	0.5	0.5	0.4

Notes: Tract hit definition described in Section 3.2.3. Periods denote hurricanes not reaching category 3 speeds in Florida or event years not covered by sample.

Table B9: Percentage of category-3-speed hit tracts by hurricane, years post

	0	1	2	3	4	5	6	7	8	9	10
Andrew92	.	.	.	.	.	.	.	0.3	0.4	0.4	0.4
Opal95	.	.	.	.	.	.	.	.	.	.	.
Earl98	.	.	.	.	.	.	.	.	.	.	.
Georges98	.	.	.	.	.	.	.	.	.	.	.
Irene99	.	.	.	.	.	.	.	.	.	.	.
Charley04	.	.	.	.	.	.	.	.	.	.	.
Frances04	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Jeanne04	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Ivan04	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Dennis05	.	.	.	.	.	.	.	.	.	.	.
Katrina05	.	.	.	.	.	.	.	.	.	.	.
Wilma05	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3
Hermine16	.	.	.	.	.	.	.	.	.	.	.
Matthew16	.	.	.	.	.	.	.	.	.	.	.
Irene17	.	.	.	.	.	.	.	.	.	.	.
Total	0.4	0.4	0.4	0.4	0.4	0.4	0.4	0.7	0.8	0.8	0.8

Notes: Census tract hit definition described in Section 3.2.3. Periods denote hurricanes not not reaching category 3 speeds or event years not covered by sample.

Table B10: Main price and transaction probability model results

	Full sample prices	Repeated sales prices	Transaction probability
Age	0.0002	-13.8	156788.2
	(0.0005)	(31302.4)	(260689.2)
Effective age	-0.006***	-0.002**	-0.001***
	(0.0010)	(0.001)	(0.0004)
Event time year =			
-6	0.008	-0.02	-0.006*
	(0.03)	(0.04)	(0.003)
-5	0.02	-0.02	-0.006*
	(0.04)	(0.05)	(0.003)
-4	-0.01	-0.02	-0.006**
	(0.02)	(0.05)	(0.003)
-3	-0.09	-0.03	-0.007**
	(0.06)	(0.04)	(0.003)
-2	-0.04	-0.03	-0.008**
	(0.03)	(0.03)	(0.004)
0	0.05**	0.05	-0.02*
	(0.02)	(0.04)	(0.008)
1	0.1***	0.1***	-0.007
	(0.03)	(0.03)	(0.006)
2	0.02	0.08***	-0.007*
	(0.03)	(0.03)	(0.004)
3	0.01	0.03	-0.004
	(0.02)	(0.03)	(0.004)
4	0.0004	0.01	0.0006
	(0.03)	(0.03)	(0.004)
5	0.01	0.04	-0.004
	(0.03)	(0.03)	(0.003)
6	0.02	0.06*	0.00003
	(0.03)	(0.04)	(0.004)
7	0.02	0.05**	-0.003
	(0.02)	(0.02)	(0.003)
8	0.05*	0.06***	-0.003
	(0.03)	(0.01)	(0.003)
9	-0.02	0.05***	-0.0008
	(0.04)	(0.02)	(0.003)
10	0.02	0.05**	0.0002
	(0.02)	(0.02)	(0.002)
County-year-type and month-type FEs	Yes	Yes	Yes
Parcel FEs		Yes	Yes
Tract FEs	Yes		
Floors and bathrooms bins	Yes		
N	7918062	1338384	31941428
R^2	0.53	0.78	0.09

Notes: Standard errors in parentheses (clustered at the county level). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Dependent variable is log(price) in the first two columns, and a (yearly) transaction indicator in the third.

Bibliography

Allcott, H. (2013). The welfare effects of misperceived product costs: Data and calibrations from the automobile market. *American Economic Journal: Economic Policy*, 5(3):30–66.

Austin, D. and Dinan, T. (2005). Clearing the air: The costs and consequences of higher cafe standards and increased gasoline taxes. *Journal of Environmental Economics and management*, 50(3):562–582.

Bayer, P., McMillan, R., Murphy, A., and Timmins, C. (2016). A dynamic model of demand for houses and neighborhoods. *Econometrica*, 84(3):893–942.

Belasen, A. R. and Polachek, S. W. (2009). How disasters affect local labor markets the effects of hurricanes in florida. *Journal of Human Resources*, 44(1):251–276.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890.

Bin, O., Kruse, J. B., and Landry, C. E. (2008). Flood hazards, insurance rates, and amenities: Evidence from the coastal housing market. *Journal of Risk and Insurance*, 75(1):63–82.

Bin, O. and Landry, C. E. (2013). Changes in implicit flood risk premiums: Empirical evidence from the housing market. *Journal of Environmental Economics and management*, 65(3):361–376.

Bin, O. and Polasky, S. (2004). Effects of flood hazards on property values: evidence before and after hurricane floyd. *Land Economics*, 80(4):490–500.

Bleemer, Z. and Van der Klaauw, W. (2017). Disaster (over-) insurance: the long-term financial and socioeconomic consequences of hurricane katrina. *Working Paper*.

Boustan, L. P., Kahn, M. E., Rhode, P. W., and Yanguas, M. L. (2017). The effect of natural disasters on economic activity in us counties: A century of data. Technical report, National Bureau of Economic Research.

Busse, M. R., Knittel, C. R., and Zettelmeyer, F. (2013). Are consumers myopic? evidence from new and used car purchases. *The American Economic Review*, 103(1):220–256.

Chaffee, S. H. and McLeod, J. M. (1973). Consumer decisions and information use. *Consumer behavior: Theoretical sources*, pages 385–415.

Dahl, C. A. (1982). Do gasoline demand elasticities vary? *Land Economics*, 58(3):373–382.

Deryugina, T. (2017). The fiscal cost of hurricanes: disaster aid versus social insurance. *American Economic Journal: Economic Policy*, 9(3):168–98.

Deryugina, T., Kawano, L., and Levitt, S. (2018). The economic impact of hurricane katrina on its victims: evidence from individual tax returns. *American Economic Journal: Applied Economics*, 10(2):202–33.

EPA (2016). Light-duty automotive technology, carbon dioxide emission, and fuel economy trends: 1975 through 2016. Technical report.

Espey, M. (1998). Gasoline demand revisited: an international meta-analysis of elasticities. *Energy Economics*, 20(3):273–295.

Gagnon, E. and Lopez-Salido, D. (2014). Small price responses to large demand shocks. *Working Paper*.

Gallagher, J. (2014). Learning about an infrequent event: evidence from flood insurance take-up in the united states. *American Economic Journal: Applied Economics*, 6(3):206–33.

Gallagher, J. and Hartley, D. (2017). Household finance after a natural disaster: The case of hurricane katrina. *American Economic Journal: Economic Policy*, 9(3):199–228.

Gibson, M., Mullins, J. T., and Hill, A. (2017). Climate change and flood beliefs: Evidence from new york real estate. *Working Paper*.

Gillingham, K. (2014). Identifying the elasticity of driving: evidence from a gasoline price shock in california. *Regional Science and Urban Economics*, 47:13–24.

Goldberg, P. K. (1998). The effects of the corporate average fuel efficiency standards in the us. *The Journal of Industrial Economics*, 46(1):1–33.

Greene, D. L. (2010). How consumers value fuel economy: A literature review. Technical report.

Hallstrom, D. G. and Smith, V. K. (2005). Market responses to hurricanes. *Journal of Environmental Economics and Management*, 50(3):541–561.

Hoople, D. (2013). The budgetary impact of the federal government's response to disasters. *Congressional Budget Office*.

Hsiang, S. M. and Jina, A. S. (2014). The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones. Technical report, National Bureau of Economic Research.

Hughes, J., Knittel, C., and Sperling, D. (2008). Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal*, 29(1).

IPCC (2013). *Summary for Policymakers*, book section SPM, page 130. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Kleit, A. N. (2004). Impacts of long-range increases in the fuel economy (cafe) standard. *Economic Inquiry*, 42(2):279–294.

Klier, T. and Linn, J. (2012). New-vehicle characteristics and the cost of the corporate average fuel economy standard. *The RAND Journal of Economics*, 43(1):186–213.

Kuminoff, N. V., Smith, V. K., and Timmins, C. (2013). The new economics of equilibrium sorting and policy evaluation using housing markets. *Journal of Economic Literature*, 51(4):1007–62.

Kutz, G. D. and Ryan, J. J. (2006). *Hurricanes Katrina and Rita Disaster Relief: Improper and Potentially Fraudulent Individual Assistance Payments Estimated to be Between 600 Million and 1.4 Billion: Testimony Before the Subcommittee on Investigations, Committee on Homeland Security, House of Representatives*. US Government Accountability Office.

Larrick, R. P. and Soll, J. B. (2008). The mpg illusion. *SCIENCE-NEW YORK THEN WASHINGTON-*, 320(5883):1593.

Levin, L., Lewis, M. S., and Wolak, F. A. (2017). High frequency evidence on the demand for gasoline. *American Economic Journal: Economic Policy*, 9(3):314–47.

Lin, C.-Y. C. and Prince, L. (2009). The optimal gas tax for california. *Energy Policy*, 37(12):5173–5183.

McCoy, S. J. and Zhao, X. (2018). A city under water: A geospatial analysis of storm damage, changing risk perceptions, and investment in residential housing. *Journal of the Association of Environmental and Resource Economists*, 5(2):301–330.

McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior.

Michel-Kerjan, E. O. (2010). Catastrophe economics: the national flood insurance program. *Journal of Economic Perspectives*, 24(4):165–86.

Murphy, A. and Strobl, E. (2009). The impact of hurricanes on housing prices: evidence from us coastal cities. *Working Paper*.

Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of political economy*, 82(1):34–55.

Seetharam, I. (2018). The indirect effects of hurricanes: Evidence from firm internal networks. *Working Paper*.

Strobl, E. (2011). The economic growth impact of hurricanes: evidence from us coastal counties. *Review of Economics and Statistics*, 93(2):575–589.

Teisl, M. F. and Roe, B. (1998). The economics of labeling: An overview of issues for health and environmental disclosure. *Agricultural and Resource Economics Review*, 27:140–150.

Turrentine, T. S. and Kurani, K. S. (2007). Car buyers and fuel economy? *Energy Policy*, 35(2):1213–1223.

van Benthem, A. and Reynaert, M. (2015). Can fuel-economy standards save the climate? *The Economist*.