The 2008 U.S. Auto Market Collapse*

Bill Dupor[†], Rong Li[‡], Saif Mehkari[§], and Yi-Chan Tsai[¶]

May 28, 2024

Abstract

New vehicle sales in the U.S. fell nearly 40 percent during the 2007-2009 recession, causing significant job losses and unprecedented government interventions in the auto industry. This paper explores three potential explanations for this decline: increasing oil prices, falling home values, and falling household income expectations. First, we use the historical macroeconomic relationship between oil prices and vehicle sales to show that the oil price spike explains roughly 15 percent of the auto sales decline between 2007 and 2009. Second, we establish that declining home values explain only a small portion of the observed reduction in household new vehicle sales. Using a county-level panel from the episode, we find (1) a one-dollar fall in home values reduced household new vehicle spending by 0.5 to 0.7 cents and overall new vehicle spending by 0.9 to 1.2 cents and (2) falling home values explain between 16 and 19 percent of the overall new vehicle spending decline. Next, examining state-level data for 1997-2016, we find (3) the short-run responses of new vehicle consumption to home value changes are larger in the 2005-2011 period relative to other years, but at longer horizons (e.g. 5 years), the responses are similar across the two sub-periods and (4) the service flow from vehicles, as measured by miles traveled, responds very little to house price shocks. Third, we establish that declining current and expected future income expectations potentially played an important role in the auto market's collapse. We build a permanent income model augmented to include infrequent repeated car buying. Our calibrated model matches the pre-recession distribution of auto vintages and the liquid-wealth-to-income ratio, and exhibits a large vehicle sales decline in response to a mild decline in expected permanent income due to a transitory slowdown in income growth. In response to the shock, households delay replacing existing vehicles, allowing them to smooth the effects of the income shock without significantly adjusting the service flow from their vehicles. Augmenting our model with a richer set of household expectations allows us to match 65 percent of the overall new vehicle spending decline (i.e. roughly the portion of the decline not explained by oil prices and falling home values). Combining our negative results regarding housing wealth and oil prices with our positive model-based findings, we interpret the auto market collapse as consistent with existing permanent income based approaches to durable goods purchases (e.g., Leahy and Zeira (2005)).

*Thanks to Evan Carson, James Eubanks and Rodrigo Guerrero for useful research assistance as well as Adrienne Brennecke for help acquiring the auto count data. The analysis set forth does not reflect the views of the Federal Reserve Bank of St. Louis or the Federal Reserve System. We thank Chris Foote, Kevin Kliesen, John Leahy, Viacheslav Sheremirov and Paul Willen for very useful comments, as well as conference and seminar participants at the Federal Reserve Bank of Boston, the Federal Reserve Bank of St. Louis, the NBER EFG Meeting and the University of Illinois.

[†] Federal Reserve Bank of St. Louis, william.d.dupor@stls.frb.org, billdupor@gmail.com.

[‡] School of Finance, Renmin University of China, lirong.sf@ruc.edu.cn.

[§] University of Richmond, smehkari@richmond.edu.

[¶] National Taiwan University, yichantsai@ntu.edu.tw.

1 Introduction

The fall in total new autos purchased played a large role in the aggregate personal consumption decline during the 2007-2009 recession. Figure 1 plots the accumulated change in motor vehicle consumption relative to 2007.[1] It drops dramatically, reaching negative $200 billion by 2010, and recovers very slowly. In contrast, as seen in the figure, the corresponding variable for total consumption (excluding vehicles) never becomes negative and recovers very quickly.

Figure 1: Cumulative change in components of personal consumption expenditure since 2007

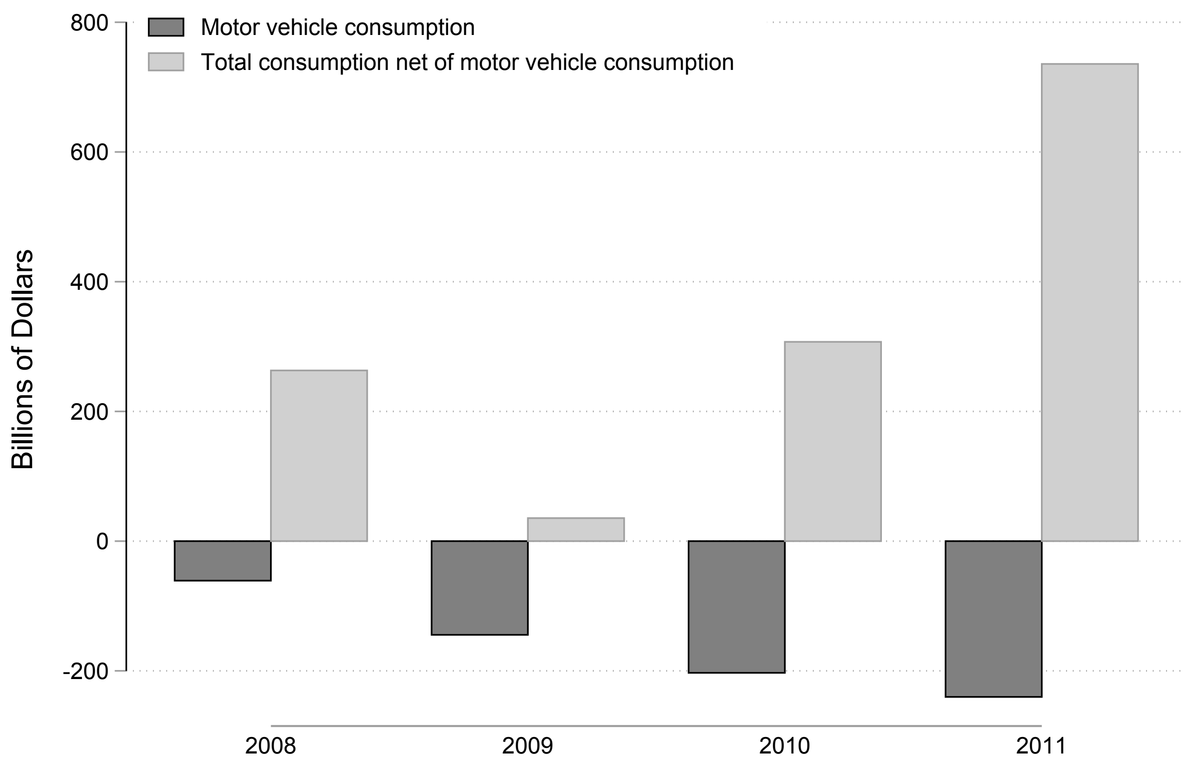

Notes: Data are annual and from the Bureau of Economic Analysis. Each bar corresponds to the accumulated change in X_t measured as $\sum_{t=2008}^{\delta} (X_t - X_{2007})$.

The new vehicle sales decline was intense and violent. In one 12-month period alone, personal new vehicle sales fell by $107 billion.[2] By spring 2009, Chrysler and General Motors faced bankruptcy. This led the U.S. government to use Troubled Asset Relief Program (TARP) funds to

[1] Our usage of the phrase motor vehicle consumption here follows U.S. Bureau of Economic Analysis (BEA) terminology. Later in the paper, we associate investment in the stock of durables with consumption and distinguish it from the consumption of the service flow from the stock of vehicles in the economy.

[2] This is a nominal seasonally adjusted rate between 2007Q4 and 2008Q4.

bailout both. At one point, the federal government owned 61 percent of General Motors.[3]

Despite the bailout, the decline in new vehicle sales had a devastating impact. Over a 2-year period, employment in the motor vehicle industry fell over 45 percent, which excludes additional knock-on effects reverberating through upstream and downstream industries.

The story is not a new one. As Martin Zimmerman (1998), then-chief economist at Ford Motor Company, wrote "I cannot think of an industry more cyclical or more dependent on the business cycle than the auto industry."

We consider three candidate explanations: rising oil prices, falling house prices and falling household income expectations. First, the role of oil prices in explaining the decline in motor vehicle sales and consumption in 2009 more generally in 2008-9 has been offered by both economists and the popular press. In a 2012 article for *The Atlantic*, Thompson (2012) writes: "For folks shacking up in the exurbs, higher gas bills ate into mortgage money. For companies, higher energy bills shocked productivity. Classic oil-shock + housing development arrested + financial crisis = Great Recession." Hamilton (2009) presents a related view from an academic perspective.

We assess the role of oil prices by examining 40 years of data on the aggregate relationship between oil prices and auto sales. Treating oil price changes as exogenous, we use regression analysis to show that oil price hikes that preceded the auto sales collapse explain about 15 percent of the auto sales decline.

Next, the housing-explanation proponents contend that as homeowners saw house prices fall, they internalized this as a wealth reduction and responded by cutting auto purchases. This effect might be stronger if homeowners used home equity to purchase cars. With falling house prices, homeowners became more borrowing constrained, which only intensified the fall in auto sales.

We exploit variation in home value and price changes to assess the role of house prices in explaining the auto sales collapse. We regress new auto sales on home values across U.S. counties and show that a one dollar decline in home values reduced household new auto spending by between 0.5 and 0.7 cents. Overall new auto spending, i.e., from consumers, businesses and government, fell by between 0.9 and 1.2 cents in response to the same change. This relatively weak response helps explain our second finding: falling home values explain between 16 and 19 percent of the overall new auto sales reduction during the period. In the historic auto market collapse, declining home values played a small part.[4]

The relatively mild responses of auto sales to home value changes might seem surprising given the attention researchers have placed on household leverage during the period. The aggregate household debt-to-income ratio rose from roughly 0.75 in 1997 to its peak of 1.2 in 2009. According to one view, over-levered households should have dramatically cut back auto purchases because of

[3] "GM and Chrysler, owned by the government, lobby the government," *The Washington Post*, January 13, 2011.

[4] Our paper, like many others studying macroeconomic phenomenon using cross-sectional regressions, suffers from the potential complication associated with estimating relative rather than aggregate effects of shocks or policy changes. Nakamura and Steinsson (2014), Dupor and Guerrero (2017) and Dupor and Guerrero (2018) present discussions and suggest strategies for comparing aggregate and relative effects in the context of fiscal multipliers.

their falling housing wealth.

If leverage effects were quantitatively important in the aggregate during the 2007-2009 recession, then one might expect to see even smaller responses of auto sales to home values outside of that period. We test this possibility by estimating similar responses using a panel of annual state-level data from 1997-2017. The state-level data are based on the same underlying house prices, but we replace vehicle sales counts with BEA motor vehicle consumption data.

Our state-level estimates of the response elasticities of motor vehicle consumption to house price changes are broadly in line with our results described above. There is a positive and statistically significant, but quantitatively mild, effect of house prices. The short-run responses (i.e. 1 to 3-years) are somewhat larger; however, at longer horizons (e.g., 5 years) leverage has little effect on the causal impact of home values on vehicle sales.

Next, we examine the effect of home values on auto usage. We replace auto sales with vehicle miles traveled in our state-level regressions and show that miles traveled were nearly unaffected by changes in home values. As such, households were able to smooth the flow of services from the stock of vehicles, as measured by miles traveled, in response to house price shocks. From the households' perspective, house price shocks did not disrupt auto usage.

Having established that oil price and house prices played a relatively minor role in explaining the auto sales collapse, the natural question is: What caused the auto sales decline? According to the Permanent Income Hypothesis (PIH), households will reduce current consumption when expected future income falls, even in the absence of borrowing constraints or reductions in tangible wealth. Moreover, if the expected future income declines were broad-based, it may be difficult to identify this effect using a structural cross-sectional regressions.

We provide microeconomic survey evidence showing that many individuals decided it was a bad time to purchase a car; moreover, the surveys establish that poor current and expected future economic conditions were primary drivers of this increased aversion to auto buying.

These survey-based findings motivate an alternative explanation for the auto market's collapse: falling future income expectations.[5]

The durability of autos together with the discrete nature with which individuals adjust their auto stocks may be important. During the 2007-2009 recession, households may have cut back on new auto purchases and simultaneously maintained their driving patterns by continuing to use their existing autos for a period of time.

With this in mind, we build a model with non-durable consumption, savings, and infrequent repeated auto purchases. In the model, individuals are subject to transitory idiosyncratic level income shocks and an unanticipated persistent aggregate income growth rate shock. The latter

[5]This brings to mind De Nardi, French and Benson (2012), who study how large a decline in future income would be required to explain the observed fall in total real personal expenditures based on a permanent income model. That paper finds that a large persistent decline in expected future income is capable of causing the decline in aggregate consumption. De Nardi, French and Benson (2012) use a model with only non-durable consumption to perform their calculations.

shock is calibrated to drive a mild decline in expected permanent income. Individuals optimally respond to negative shocks of this kind by delaying auto replacement. The model matches the cross-sectional distribution of autos by vintages and the liquid-wealth-to-income ratio in the period prior to the 2007-2009 recession.

The shock amplification mechanism is very strong: a 3.1 percent decline in expected permanent income drives a 70 percent decline in aggregate new vehicle sales. Furthermore, augmenting the model with a richer set of household expectations allows the model to match 65 percent of the overall new vehicle spending decline in the data. (i.e. the portion of the decline in the data that is not explained by oil prices and falling home values).

Our paper relates to several lines of research. McCully, Pence and Vine (2019) report that very few households in the U.S. purchase cars with home equity lines of credit or proceeds from cash-out refinancing. Auto buyers that do use these sources are affluent and have ample access to credit. They also explore whether household use home equity extraction to indirectly support car purchase. In this case too, they find while home equity extraction leads to a statistically increase in auto loan originations, its overall impacts on car purchase is very small.

Other papers link the house price decline during the 2007-2009 recession to the drop in consumer spending in the U.S. These include Mian, Rao and Sufi (2013) who find a strong positive relationship between house prices and both durable and non-durable consumption during the period. Based on county-level data, they report that consumption increases by 5.4 cents from a one dollar increase in housing wealth with 43 percent of this increase (2.3 cents per dollar) coming from auto spending (including new auto, used auto, and auto parts/accessories). Our results contrast from theirs in that we focus on new auto spending, while their paper's results look at a broader basket of auto sales that also include used cars and auto parts/accessories. In Mian and Sufi (2018), a follow-up paper, the authors updated their previous results by mainly removing sales by used car dealers and auto parts/accessories stores. This resulted in an updated figure of 1.8 cents per dollar being attributed to auto spending, however, still includes service and used car sales by new car dealers still contrasting their result from ours.

Kaplan, Mitman and Violante (2020a) find a positive relationship between house prices and non-durable consumption in the cross section during the 2007-2009 recession. General equilibrium analyses regarding consumption and the housing market include Garriga and Hedlund (2020) and Kaplan, Mitman and Violante (2020b). Most of the papers on the consumption response to home value changes focus on non-durable consumption. They do not model non-housing durable goods; whereas we do in our paper.[6] Lehnert (2004) is a seminal early contribution on the consumption effects of home value changes.

Our economic model's mechanism has been described in existing theoretical work. Leahy and Zeira (2005) present a model with infrequent durable goods purchases in which the timing decision

[6]These include, for example, Berger, et.al. (2018), Corbae and Quintin (2015), and Favilukis, Ludvigson and Van Nieuwerburg (2016).

of auto purchases amplifies and propagates shocks. In response to negative shocks, individuals who were going to purchase durable goods postpone their purchases.[7]

There are other —at least partial —explanations for the auto sales collapse. Benmelech, Meisenzahl and Ramcharan (2017) argue that the disruption in the asset-backed commercial paper market reduced the availability of auto loans and caused up to 31 percent of the auto sales fall during the episode. Gavazza and Lanteri (2021) investigate the effects of secondhand auto price on the dynamics of new car purchase. Using a model-approach, they find that during financial crisis, debt-constrained households postpone the replacement of their older cars, leading to a fall in the demand for mid-quality cars and depressing their prices, which in turn affects the new car replacement cost for the wealthy households and thus decreases overall new-auto sales. Another explanation focuses on the mismatch between the increased demand for higher efficiency cars, in light of positive oil price shocks, and the lack of supply of efficient vehicles by some major auto manufacturers.

Finally, using the Consumer Expenditure Survey, Attanasio et. al. (2022) decompose car expenditure into an extensive margin and an intensive margin, i.e., whether to buy a car and the size of car conditional upon the decision to purchase. They find that while most recessions are associated with a decline in the adjustment along the extensive margin, the Great Recession in addition witnesses a large decline in the intensive margin adjustment. Similar to our paper they also construct a partial equilibiurm model with Ss-type decisions and study the effects of income shocks in such a model.

Section 2 examines the impact of oil prices on the 2008 auto sales collapse using historical aggregate data. Section 3 presents our county-level house price findings from the 2007-2009 recession. Section 4 presents our state-level house price findings using data from the past two decades. It finds a weak response of auto sales, and also miles traveled, to home value changes. Section 5 presents a dynamic permanent income model augmented with auto purchases in which declines in expected permanent income generate a large decline in aggregate autos purchased. The final section recaps.

2 National Vehicle Sales and the Price of Oil

Our first candidate explanation for the auto collapse is the oil price hike that occurred before and around the same time as the decline in vehicle sales.[8] Between 2007Q1 and 2008Q1, the price of oil increased roughly 50 percent. Over the following year, vehicle sales per capita fell by roughly the same amount.

To evaluate the oil price vehicle sales channel, we explore the historical relationship between vehicle sales and oil prices. We begin with oil price and national vehicle sales.[9] Let p_t and q_t denote

[7] Empirical work on autos and the permanent income hypothesis include Adda and Cooper (2000), Bernanke (1984) and Eberly (1994).

[8] See, for example, Hamilton (2008) for the case of the 2007-2009 recession and Edelstein and Kilian (2009) who consider a longer time period in the U.S.

[9] The original series, from the U.S. Energy Information Agency and the U.S. Bureau of Economic Analysis, are

the quarter-t log real price of oil and log quantity of vehicles (per capita) sold. We regress the four quarter change in p_t on q_t for the sample 1977Q1 - 2019Q4. We plot these series on Figure 2.

Figure 2: Oil price and new vehicle sales growth

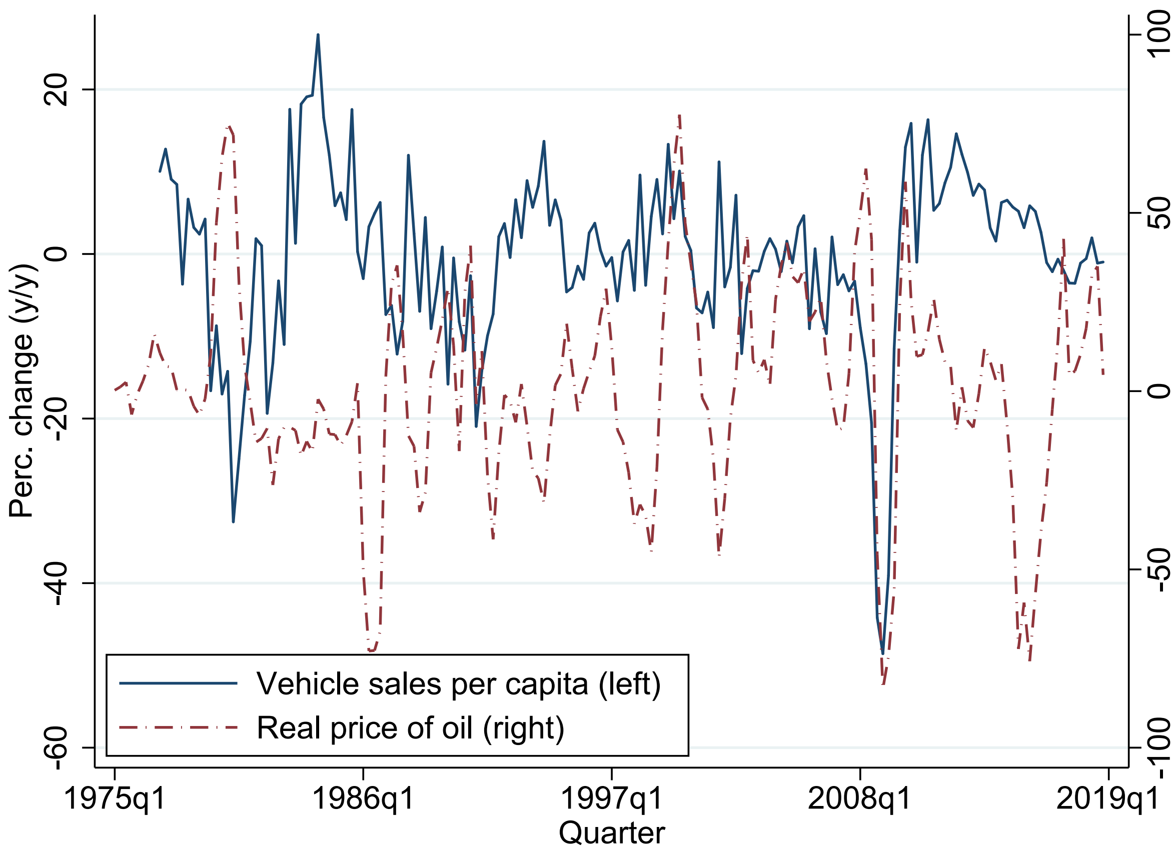

Notes: Data are from the BEA and EIA.

We exclude the 2007-2010 period from the regression. Because oil prices are endogenous, we instrument oil price growth using the Baumeister-Hamilton oil supply shock series. Moreover, we include three lags of the one quarter growth rate of oil prices and of vehicle sales as controls.

Column (1) of Table 1 presents the estimate for this specification. The coefficient on oil price growth equals -0.15 (SE=0.07). The coefficient is significantly different from zero and of the expected sign. The estimate implies that a 50 percent increase in the price of oil due to an exogenous negative oil supply shock causes a 7 percent decline in vehicle sales. Quantitatively, the shock is incapable of explaining most of the 40 to 50 percent year-over-year decline in auto sales that occurred during this period.

monthly. They are time averaged to get quarterly data.

Table 1: Response of national one-year vehicle sales growth to oil prices, two-stage least squares

	(1) Coef./SE	(2) Coef./SE	(3) Coef./SE	(4) Coef./SE
Real oil price growth (y/y)	-0.15**	-0.12**	-0.13	-0.10
	(0.07)	(0.06)	(0.09)	(0.07)
R^2	0.24	0.35	0.13	0.46
N	153	153	165	153

Notes: Dependent variable = Vehicle sales per capita growth rate (y/y). Column (1) includes four lags of the one-quarter growth rate of oil prices and of vehicle sales and excludes 2007-2010. Column (2) adds four lags of the one-quarter growth rate of house prices. Column (3) adds the 2007-2010 period. Column (4) reproduces column (1) with the addition of four lags of the federal funds rate. * $p < .1$, ** $p < .05$, *** $p < .01$. Standard errors are robust with respect to heteroskedacity and autocorrelation.

In the next section, we explore the role of house prices on auto sales. As such, in this section, we wish to isolate the effect of oil prices on auto sales by controlling for house prices. Column (2) adds to the regression four lags of the house price growth rate as additional control variables. The coefficient on oil prices remains statistically significant but the effect is somewhat dampened, with a coefficient equal to -0.12 (SE=0.06).

Column (3) adds the 2007-2010 period to the sample. In this case, the coefficient is not statistically different from zero and quantitatively similar to the previous specification. Column (4) replaces the house price growth rates with four lags of the change in the federal funds rate in order to control for the effects of monetary policy. The coefficient on oil price growth equals -0.10 (SE=0.07). Based on this final estimate, one would expect auto sales would decline 5 percent ($= -0.10 \times .5$) solely due to the oil price increase. This implies only a small contribution of oil prices to overall collapse of the auto market (between one-tenth and one-fifth of the decline in auto sales that occurred at the time).

Furthermore, the price of oil spiked downward shortly after the upward movement. In 2009Q1, the oil price was 83 percent lower than in the fourth quarter that preceded it. There was, however, no rapid offsetting bounce back in auto sales.

Our findings may at first seem to stand in start contrast with Hamilton (2009), who finds that nearly one-half of the decline in PCE motor vehicle and parts spending can be explained by oil shocks between 2006 and 2008.[10] One reason for the difference in results is that Hamilton (2009) studies the behavior of motor vehicles and parts spending. We plot an index of this series as the dash-dotted line in Figure 3. Note that his auto spending measure falls by only about 22 percent. The dashed line (our vehicle sales measure) falls by over 40 percent. Thus, Hamilton faces a much smaller hurdle with the respect to the amount of the decline that he needs to explain.

[10]See, for example, his Figure 17.

Figure 3: Three quantity measures of auto spending

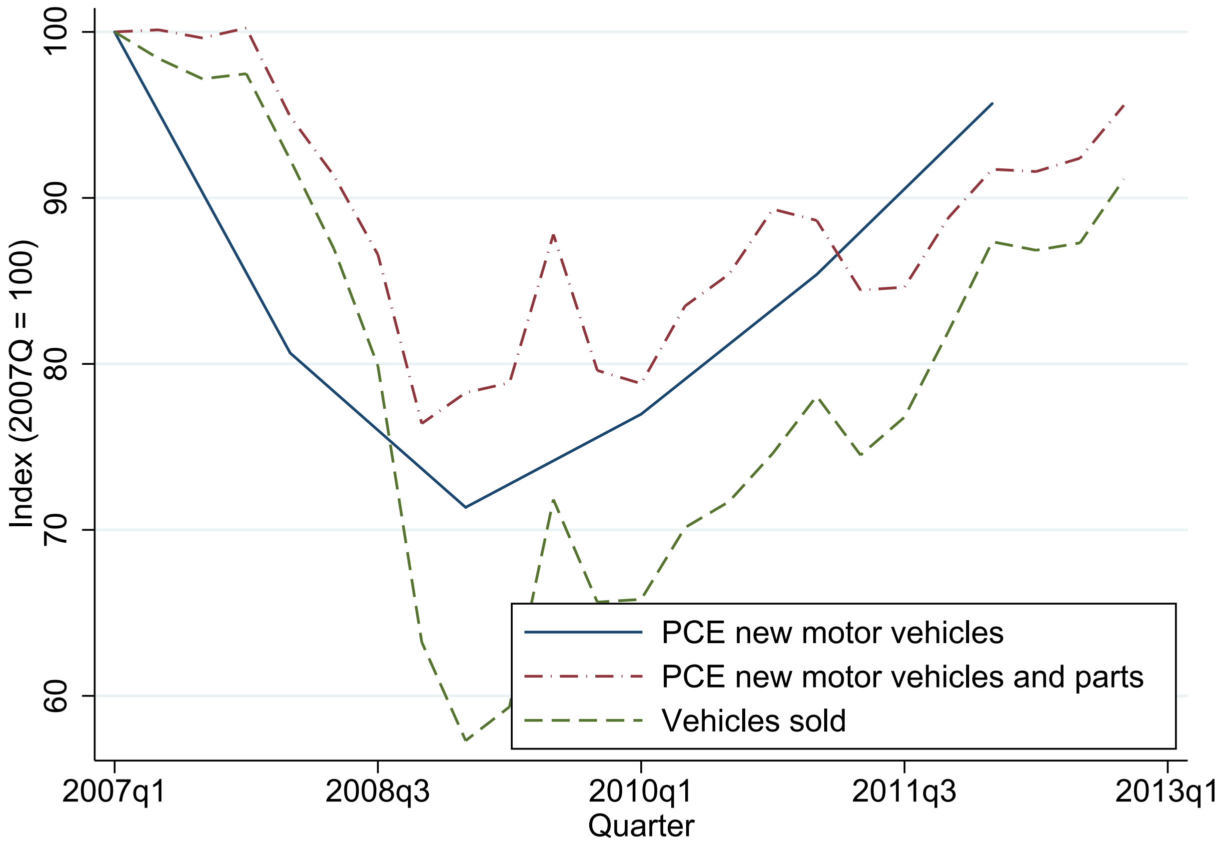

Notes: Data are from the BEA.

Auto sales did not bounce back quickly following its market's 2008 collapse. If oil prices had been a primary driver of the 2008 episode, then one would expect similar persistent changes in auto sales following other oil price shocks. To examine whether this was the case, we look beyond one year auto sales changes and analyze the three-year cumulative decline in auto sales, which we define as

$$q_t^c = \frac{1}{4}(100)\sum_{j=1}^{12}(q_{t+j-1} - q_{t-1})$$

In the year of the collapse, this decline was between 74 and 100 percent (depending on the particular base quarter used). A 100 percent decline means that one-year's worth of auto sales were lost over a three year period relative to the base quarter.

Table 2: Response of national three-year cumulative vehicle sales growth to oil prices, two-stage least squares

	(1) Coef./SE	(2) Coef./SE	(3) Coef./SE	(4) Coef./SE
Real oil price	-0.32	-0.21	-0.19	-0.07
growth (y/y)	(0.30)	(0.22)	(0.27)	(0.28)
R^2	0.19	0.49	0.35	0.39
N	147	147	159	147

Notes: Dependent variable = 3-year cumulative vehicle sales per capita growth rate (y/y). Column (1) includes four lags of the one-quarter growth rate of oil prices and of vehicle sales and excludes 2007-2010. Column (2) adds four lags of the one-quarter growth rate of house prices. Column (3) adds the 2007-2010 period. Column (4) reproduces column (1) with the addition of four lags of the federal funds rate. * $p < .1$, ** $p < .05$, *** $p < .01$. Standard errors are robust with respect to heteroskedacity and autocorrelation.

We replicate our analysis from the previous table using this medium-run auto sales response as the dependent variable. These results are present in Table 2. Besides the change in dependent variables, all other aspects of each column are identical.

As expected, each of the coefficients are negative. Each of the coefficients is somewhat larger (in absolute value) then its corresponding one-year growth specification. However, none of the coefficients is statistically different from zero. The point estimates indicate that oil prices explain only a small fraction of the auto sales decline measured this way. Based on our final specification, column (4), the coefficient 0.07 implies that a 50 percent increase in oil prices is associated with a 3.5 percent ($= -0.07 \times .5$) decline in auto sales. As such, we conclude that oil prices were not a major driver of the collapse.

3 Vehicle Sales and House Prices - County-Level Analysis

3.1 Data and Econometric Model

Let $A_{i,t}$ denote the dollar value of new vehicles sold overall in county i in quarter t. We calculate auto counts from county-level auto registrations, which include vehicles sold to households, businesses and government.[11] The vehicles acquired include those gotten via straight cash purchases, trade-in purchases, leases, etc. To go from quantities to dollar values, we multiply the quantity of autos by the nationwide average new auto price, which was in the range of $26,200 to $26,950 during the period according to the Bureau of Transportation Statistics (BTS).[12]

At times we distinguish the household response of vehicle spending from that of overall vehicle spending. We map from overall spending to household spending by using the aggregate share of new vehicle purchases by consumers as a fraction of total new vehicle purchases (from consumers, businesses and government), which is available annually from the BTS. Over 40 percent of new cars were sold to businesses and government during this period.[13]

Let $V_{i,t}$ denote the dollar value of the owner-occupied housing stock in county i in quarter t. CoreLogic constructs monthly house price data at the county level; however, these are reported as indices rather than dollar amounts. To go from indices to dollar prices, we begin with the county-level median house price available from the 2000 U.S. Census. Then we multiply this Census house price by the gross growth rate of the Corelogic index between the month of interest and January 2000. Let $P_{i,t}$ denote the current dollar price of an owner-occupied house, calculated according to the procedure.

To calculate the value of the county-level housing stock, we multiply $P_{i,t}$ by the number of households in owner-occupied housing from the 2006 Census.

Let $a_{i,t,\delta} = \log(A_{i,t+\delta-1}) - \log(A_{i,t-1})$. Next, let $a^c_{i,t,\delta}$ be the cumulative percentage increase in auto sales over a δ quarter horizon relative to a quarter $t-1$ baseline in county i:

$$a^c_{i,t,\delta} = \frac{1}{4}\sum_{j=1}^{\delta} a_{i,t,j}$$

The variables $p_{i,t,\delta}$ and $p^c_{i,t,\delta}$ are defined similarly.[14]

Let $\bar{p}_{i,t,\delta}$ and $\bar{p}^c_{i,t,\delta}$ denote the nationwide averages of their county-level counterparts, where the averages are weighted by the number of households in a county. Defining these variables as such permits us to estimate the dynamic, cumulative responses of auto sales to home value shocks. Cumulative responses give the change in auto sales accumulated over a specific horizon with respect

[11] Vehicles includes autos, light trucks and SUVs.

[12] Average new car auto prices changed very little during the period considered. The federal data on auto prices stop in 2010, and we use the 2010 price for later years where required. See Table 14 in the appendix for the time series of the average new vehicle price over this period.

[13] See Table 15 in the appendix for the time series on the share of new autos purchased by households.

[14] The use of accumulated growth rates means that our resulting regression coefficients can be interpreted as areas under impulse response functions. Ramey and Zubairy (2017) argue compellingly that this is an useful way to summarize dynamic responses to shocks.

to the accumulated change in house prices over the same horizon.[15]

First, we estimate the elasticity of vehicle sales to house price changes, using:

$$a^c_{i,t,\delta} = \phi_\delta p^c_{i,t,\delta} + \beta_\delta X_{i,t} + v_{i,t,\delta} \tag{1}$$

for $\delta = 1, ..., D$. By the form it takes, equation (1) implements the Jorda (2005) local projections approach.

Here, $X_{i,t}$ consist of a linear trend, seasonal dummies and a "Cash for Clunkers" dummy, which equals 1 in 2009Q3 through 2010Q1. We also include one lag of the growth rate in auto sales and house prices at $t-1$ (i.e., $a_{i,t-1,1}$ and $p_{i,t-1,1}$). The sample covers 2007Q2 through 2010Q2.

The coefficient ϕ_δ is then the cumulative percentage increase in auto sales through horizon δ in response to a 1 percent increase in house prices (cumulative through horizon δ). We call this the dynamic sales elasticity or simply the sales elasticity. The estimation uses least squares and is weighted by the number of households in a county. We report heteroskedasticity and autocorrelation corrected (HAC) standard errors throughout the paper.

Table 3 reports the sales elasticities at various horizons. Note that the largest potential sample size falls as we move to longer horizons because we lose observations as we extend the horizon of the cumulative responses. To make estimates more comparable, every estimate is based on the observations for the 3-year-horizon sample.

Column (1) reports a one-year elasticity equals 1.08 (SE=0.059). Columns (2) and (3) report the 2- and 3-year-horizon responses. The responses are all positive and statistically different from zero. Moreover the responses fall with the horizon. The 3-year sales elasticity equals 0.60 (SE=0.03).

Interestingly, the cumulative response of auto sales decreases rather than increases in response to an accumulated change in house prices. In a standard adjustment cost model, if changes to the growth rate of purchases of a good lead to additional convex costs, this would lead to a gradually increasing cumulative response to a positive wealth shock. On the other hand, the decreasing cumulative response seen in Table 3 may be due to the durable nature of autos.

A short-run increase in vehicle sales in response to a positive house price shock is not simply an immediate increase in sales with no related dynamic effects. Rather, an increase in house prices could in part generate greater sales immediately because the now-richer households pull consumption from the future to the present.

The control variable coefficients are all statistically different from zero and of the expected signs. The Cash for Clunkers fixed-effect coefficient is positive, indicating (very sensibly) that sales growth was stronger over horizons that included the government incentive program. The coefficient on lagged house price growth is positive, suggesting a somewhat delayed reaction of vehicle sales to auto prices. Finally, the coefficient on vehicle sales is negative. This is likely due to the durable

[15] Later in the paper, we estimate the regressions in growth rates rather that cumulative changes for a specific cross-section. The main findings using either approach are similar.

Table 3: Cumulative overall new sales elasticities to house price changes, county level panel, least squares

	(1) Coef./SE	(2) Coef./SE	(3) Coef./SE
1-yr cum HP growth	1.079***	-	-
	(0.059)		
2-yr cum HP growth	-	0.696***	-
		(0.041)	
3-yr cum HP growth	-	-	0.599***
			(0.033)
Vehicles sold (lag growth rate)	-0.001*	-0.005***	-0.009***
	(0.001)	(0.001)	(0.002)
HP (lag growth rate)	0.005***	0.011***	0.016***
	(0.001)	(0.003)	(0.004)
Cash for Clunker fixed effect	0.086***	0.150***	0.144***
	(0.009)	(0.018)	(0.028)
Quarter	0.006***	0.066***	0.140***
	(0.001)	(0.002)	(0.003)
R^2	0.39	0.58	0.66
N	14916	14916	14916

Notes: The dependent variable is the cumulative percentage change in new auto sales at the appropriate horizon. * $p < .1$, ** $p < .05$, *** $p < .01$. Regressions weight each observation by the number of households in the county and include seasonal fixed effects (not reported). Standard errors are robust with respect to heteroskedacity and autocorrelation. HP = house price.

nature of autos. Intuitively, a recent past period of intense accumulation of the stock of autos likely reduces the need to invest in autos in the near future. The coefficients on the second- and third-quarter seasonal dummies, not reported here, are positive and statistically different from zero.

Under a set of simplifying assumptions, one can map an elasticity reported here into a derivative: specifically, the per dollar change in overall vehicle spending in response to a one dollar increase in home values. Suppose new auto prices and the home ownership rate are roughly unchanged over the period. Then this derivative is approximately equal to the corresponding estimated elasticity times the ratio of the value of new vehicles sold relative to the value of the housing stock, averaged over the same period. For the 2006-2009 period, this ratio is approximately 0.02. In other words, the value of the housing stock is about 50 times greater than the value of one year's overall auto purchases.

This implies that, at the 3-year horizon, each dollar of additional housing wealth increases overall auto sales by 1.2 cents ($= 0.02 \times 0.599$). This is an approximation. In the next subsection, we estimate this derivative, sometimes called the *marginal propensity to consume (MPC)*, directly.

3.2 Vehicle Acquisition Responses

Next, we estimate the model using cumulative changes in levels rather than cumulative changes in growth rates. All of the variables in this subsection are reported in per household terms. Define

$$V^c_{i,t,\delta} = \frac{1}{4}\sum_{j=1}^{\delta} (V_{i,t+\delta-1} - V_{i,t-1})$$

and let $A^c_{i,t,\delta}$ be defined analogously. The regression specification is

$$A^c_{i,t,\delta} = \beta_\delta V^c_{i,t,\delta} + \Gamma_\delta S_{i,t} + \varepsilon_{i,t}$$

Here, $S_{i,t}$ consist of a linear trend, seasonal dummies and a "Cash for Clunkers" dummy, which equals 1 in the 2009Q3-2010Q1. We also include one lag of the change in auto sales and home values at $t-1$ (i.e., $\Delta A_{i,t-1}$ and $\Delta V_{i,t-1}$ and). As before, the regressions are weighted by the number of households in the county.

This second model has a straightforward interpretation. We call the coefficient β_δ the overall vehicle acquisition response, or acquisition response (AQR). It is the cumulative dollar change in vehicle acquisitions in a county over a δ-quarter horizon in response to a one dollar cumulative increase in home values over the same horizon.

AQR more precisely describes what we actually can measure given the data available than some other language used in existing research, such as the MPC. MPC, however, is not suitable in the current context. First, vehicles are durable goods; households and businesses consume the service flow from their stock of durables. Second, while our data tell us something about durable goods investment, through new car registrations, we do not know the extent to which individuals disinvested in vehicles by scrapping or selling their existing stock. An individual who purchases a new car but "trades in" a similar but slightly used car may experience a very small increase in the flow of services despite the new car purchase.

To this point, Figure 4 plots indices calculated from the number of new autos sold along with the total vehicle miles traveled during the period. While new auto sales fall dramatically, total vehicle miles traveled change very little. This suggests that the flow of services associated with autos was been nearly unchanged, meaning that households were largely able to smooth consumption of auto usage.

Figure 4: Total auto sales and vehicle miles traveled

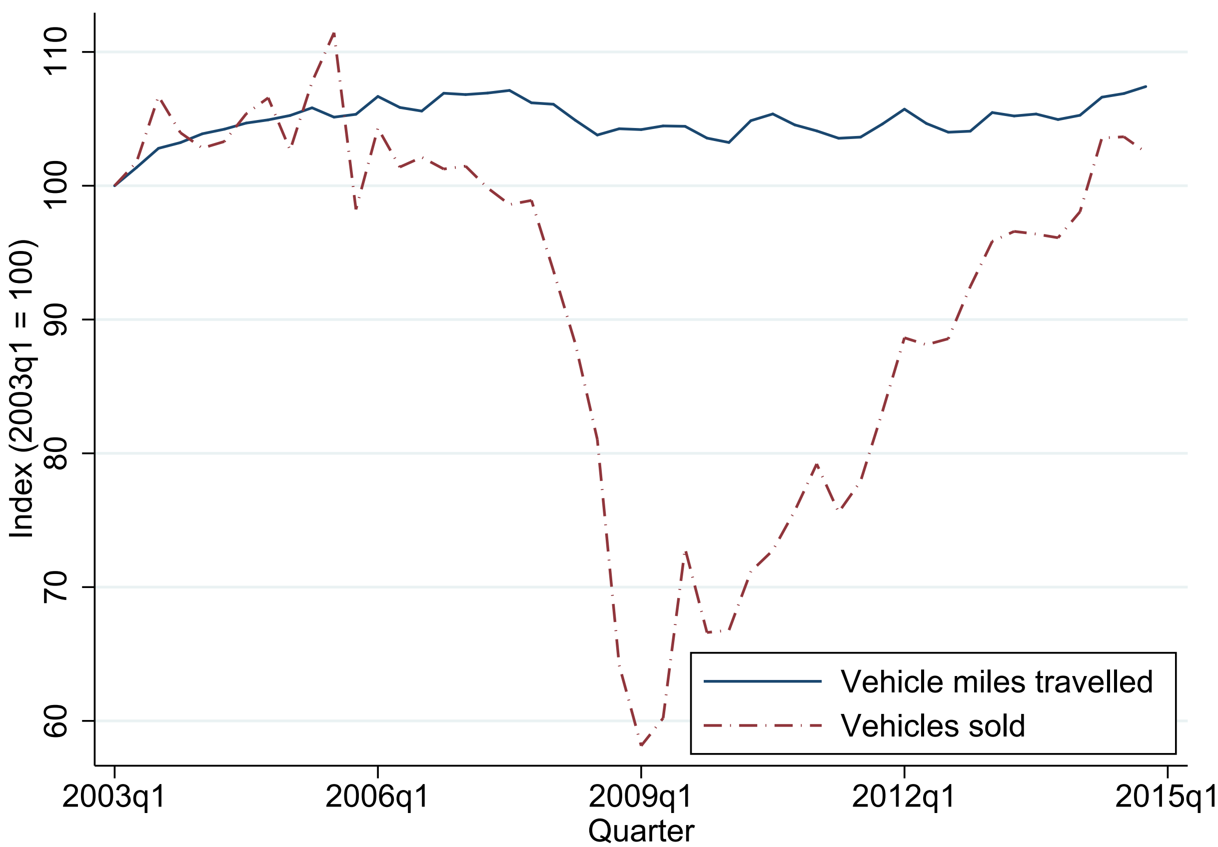

Notes: Auto sales are from the Bureau of Economic Analysis measure of the quantity of new vehicles sold. Miles traveled is from the Federal Highway Administration.

A slightly better, but also deficient, term might be *marginal propensity to spend (MPS)*. Since we do not know the frequency or value of trade-ins for new vehicle purchases, we cannot infer how much out-of-pocket spending occurs when a vehicle is acquired. Even apart from trade-ins, many vehicles are rented or leased. In this case, a person who acquires an auto would spend only a fraction of the auto's full purchase price.

Table 4 presents the overall (i.e., inclusive of consumers, businesses and government) AQR at three different horizons. Examining columns (1) through (3), note that the AQR is positive and statistically different from zero at each horizon. The response is nearly unchanged across horizons. We focus particular attention on the 3-year horizon, similar to other papers in the literature for example Mian, Rao and Sufi (2013). The 3-year AQR equals 0.012 (SE=0.001). This means that a one dollar increase in home values is associated with a 1.2 cent increase in auto sales. Reassuringly, the estimate of the AQR is equal to the approximated value using the elasticity estimate of the previous section.

Table 4: Cumulative overall vehicle acquisition responses, county-level panel

	(1) Coef./SE	(2) Coef./SE	(3) Coef./SE
1-yr Cum HP Change	0.012***	-	-
	(0.002)		
2-yr Cum HP Change	-	0.011***	-
		(0.001)	
3-yr Cum HP Change	-	-	0.012***
			(0.001)
Vehicles sold (lag	-0.131	-0.391	-0.628
change)	(0.244)	(0.521)	(0.867)
HP (lag change)	0.008	0.014	0.016
	(0.006)	(0.013)	(0.021)
Cash for Clunkers	0.343***	0.447*	0.569
fixed effect	(0.110)	(0.243)	(0.416)
Quarter	0.055***	0.282***	0.504***
	(0.008)	(0.018)	(0.033)
R^2	0.15	0.22	0.22
N	13673	13673	13673

Notes: The dependent variable is the cumulative change in vehicle sales at the appropriate horizon. * $p < .1$, ** $p < .05$, *** $p < .01$. Regressions weight each observation by the number of households in the county and include seasonal fixed effects (not reported). Standard errors are robust with respect to heteroskedacity and autocorrelation.

Vehicle registrations include those of businesses, governments and consumers. Thus, our overall AQR reflects the contribution of both types of buyers.

To distinguish the overall AQR from the household AQR, we map from overall spending to household spending by using the aggregate share of new autos purchased by consumers as a frac-

Table 5: Cumulative household vehicle acquisition responses, county-level panel

	(1) Coef./SE	(2) Coef./SE	(3) Coef./SE
1-yr Cum HP Change	0.009*** (0.001)	-	-
2-yr Cum HP Change	-	0.007*** (0.001)	-
3-yr Cum HP Change	-	-	0.007*** (0.001)
Vehicles sold (lag change)	-0.209 (0.218)	-0.471 (0.454)	-0.730 (0.763)
HP (lag change)	-0.002 (0.003)	0.004 (0.006)	0.005 (0.010)
Cash for Clunker fixed effect	0.136** (0.057)	0.215* (0.125)	0.240 (0.215)
Quarter	0.010** (0.004)	0.095*** (0.009)	0.202*** (0.016)
R^2	0.12	0.17	0.18
N	13673	13673	13673

Notes: The dependent variable is the cumulative change in vehicle sales to households at the appropriate horizon. * $p < .1$, ** $p < .05$, *** $p < .01$. Regressions weight each observation by the number of households in the county and include seasonal fixed effects (not reported). Standard errors are robust with respect to heteroskedacity and autocorrelation.

tion of total new auto purchases (from consumers, businesses and government), which is available annually from the BTS. Over 40 percent of new cars are sold to businesses and government during this period.[16]

Table 5 in the appendix presents the household AQR regressions based on this measure of vehicle sales. The coefficient should be interpreted as the response of vehicles acquired for personal use to a change in home values. The coefficients at the three horizons are between 0.007 and 0.009.

3.3 A Cross-Sectional Specification

As further robustness of our results and for consistency with previous research, such as Mian, Rao and Sufi (2013), we next use cross-section rather than panel analysis and use changes in home values rather cumulative changes.

Our dependent variable is the change in the dollar value of overall auto acquisitions between the first half of 2007 (2007H1) and the first half of 2009 (2009H1) in county j. We choose 2009 as the end year because it follows the collapse of vehicle sales that began in September 2008. It excludes the second half of 2009 because this period contains a transitory spike in sales due to the Cash for Clunkers program. We choose the starting year as 2007 because it precedes the auto market's collapse and is the first year of data available to us. Our independent variable is the change in the value of the housing stock in each county between 2007H1 and 2009H1.

We estimate the cross-sectional model in a way that necessitates fewer control variables than in our panel regressions. First, we take differences over the same half years, therefore we do not require seasonal dummies. Second, the estimation sample ends before implementation of Cash for Clunkers, which eliminates the need for the corresponding fixed effect. We estimate the model with and without lagged changes in vehicle sales and home values.

Table 6 contains the first set of regressions. It reports HAC standard errors and uses observation weights given by the number of households in each county. Column (1) contains the simplest specification. The coefficient on the change in home values equals 0.01 (SE = 0.002). That is, an increase in home value of one dollar in a county is associated with a 1.0 cent increase in new vehicles acquired (by households, businesses and government) in that county. In this specification, the coefficient equals the AQR. As with the cumulative response, there is a muted, but statistically significant and precisely estimated, increase in overall auto acquisitions in response to increases in home values.

Next, the intercept coefficient plays an important role in the study. The intercept coefficient can be interpreted as the best linear predictor of the change in auto sales in a county with no change in home values. Its value equals -1.42 (SE = 0.07). The weighted average of the dependent variable

[16]The time series for this share appears in Table 15 in the appendix. This partition implicitly assumes that the share of vehicles sold to individuals in a county approximates that counties' share of vehicles sold to both individuals and businesses.

is -1.65. This implies that 80 percent (= -1.42 / -1.65) of the typical new auto sales change in a county is captured by the intercept rather than being associated with the change in home values.

The reduction in sales by vehicle manufacturers was nationwide, occurring largely in regions with and without depressed house prices. There is a small effect of declining home values on vehicle sales no doubt, but most of the decline in vehicle acquisitions is captured in the regression intercept.

Table 6: Overall AQR of new vehicle sales to changes in home values

	(1) Coef./SE	(2) Coef./SE	(3) Coef./SE	(4) Coef./SE
HP change	0.010***	0.011***	0.011***	0.009***
2007H1-2009H1	(0.002)	(0.002)	(0.002)	(0.002)
HP change	-	-0.007	-0.008	-0.008*
2006H1-2007H1		(0.005)	(0.005)	(0.005)
Income pc (2006)	-	-	-0.006**	-
			(0.003)	
Non-bank finance	-	-	-	-1.968***
loan share				(0.368)
Intercept	-1.422***	-1.389***	-1.042***	-0.622***
	(0.073)	(0.078)	(0.164)	(0.140)
Frac. explained by home value declines	0.185	0.209	0.193	0.163
R^2	0.114	0.118	0.126	0.156
N	1243	1243	1242	1243

Notes: The dependent variable is the change in auto sales (annualized, thousands of dollars per household). * $p <$ .1, ** $p <$ .05, *** $p <$ .01. "Fraction explained by home value declines" is the proportion of the average change in auto sales due to falling home values. Changes in variables are computed from 2007H1 to 2009H1. Regressions weight each observation by the number of households in the county.

To a great extent, new auto sales fell because the average household and business in most counties cut back on auto purchases, and not because of declining sales of the average car owner mainly in counties that experienced dramatic house price declines.

One can also see the limited role of housing in explaining the auto market collapse by applying the following counterfactual to our regression results. Take the vector of observations of home value changes in the sample and change every negative value to instead equal zero. Next, compute the fitted values from the regression using the non-negative modified vector. These fitted values are the econometric model's best predictor of the auto sales changes for the counties had there been no observed house price declines.

Next, divide the weighted average of this auto sales change predictor by the weighted average of the actual sample auto sales changes. This ratio is the fraction of the change in auto sales that can be explained without allowing for house price declines. The row labelled "Fraction explained by home value declines" in Table 6 reports this ratio subtracted from one. In Column (1), only 19 percent of the auto sales decline is explained by reductions in home values.

Figure 5 contains a scatter plot corresponding to this specification. The long-dashed line indicates the best-fit line from the weighted regression. Its slope is the AQR. The best-fit line intersects the vertical axis at the regression intercept. This is the best estimate of the county-average change in auto sales in a county that saw no house price change between 2007H1 and 2009H1. We plot the unconditional weighted average of the change in auto sales as the horizontal dash-dotted line.

Figure 5: Response of overall new vehicle sales to changes in home values

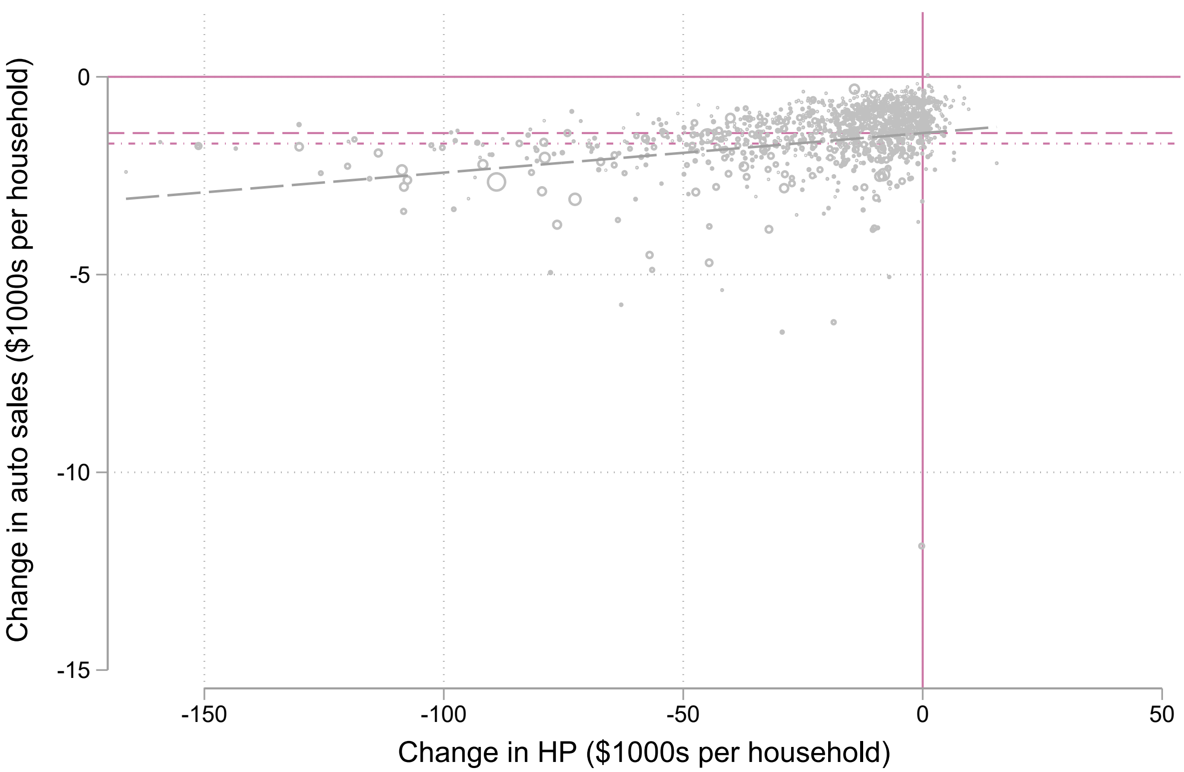

Notes: The long-dashed line is the best fit from a weighted regression of changes in overall new auto sales on changes in home values. Circle sizes are proportional to the number of households in each county. The short-dashed line corresponds to the regression intercept, i.e., the best linear predictor of overall new auto sales in a county that saw no change in home values. The dash-dotted line is the unconditional weighted average of the change in auto sales. hh=household. Changes in variables are computed from 2007H1 to 2009H1. The auto sales change in annualized.

The close proximity of the two horizontal lines indicates that changes in home values are explaining only a small fraction of the observed decline in auto sales. This is *despite the fact* that there is a statistically significant relationship between auto sales and home values. If the aggregate decline in auto sales had been entirely accounted for by home value changes, then the intercept would be zero or, equivalently, the short-dashed line would lie on the horizontal axis.

Column (2) in Table 6 adds the lagged change in home prices as a control, which brings us closer to the panel specification used earlier.[17] Both of our two main results—a low response of autos to home value changes and a low fraction of vehicle sales explained by declining home values—are maintained in this specification.

Column (3) adds income per household to the regression.[18] The coefficient on income is negative: lower-average-income counties had a smaller increase in auto purchases ceteris paribus. The AQR is nearly unchanged.

Column (4) adds the pre-recession share of auto loans provided by non-bank finance companies as an additional control. Benmelech, Meisenzahl and Ramcharan (2017) find that this was an important driver of auto sales. They argue that a negative shock to the asset-backed commercial paper market during the financial crisis reduced credit availability in regions that had relied on non-bank finance companies. The coefficient on the non-bank finance loan share is of the expected sign; however, the inclusion of the variable has only a small effect on the AQR response to home value changes.

Table 7: Household AQR of new vehicle sales to changes in home values

	(1) Coef./SE	(2) Coef./SE	(3) Coef./SE	(4) Coef./SE
HP change	0.005***	0.006***	0.006***	0.005***
2007H1-2009H1	(0.001)	(0.001)	(0.001)	(0.001)
HP change	-	-0.004	-0.004	-0.004*
2006H1-2007H1		(0.003)	(0.003)	(0.003)
Income pc (2006)	-	-	-0.002	-
			(0.001)	
Non-bank finance	-	-	-	-1.036***
loan share				(0.188)
Intercept	-0.645***	-0.628***	-0.539***	-0.224***
	(0.036)	(0.039)	(0.083)	(0.072)
Frac. explained by HP declines	0.210	0.238	0.228	0.186
R^2	0.126	0.130	0.132	0.172
N	1243	1243	1242	1243

Notes: The dependent variable is the change in household auto sales (annualized, thousands of dollars per household). * $p < .1$, ** $p < .05$, *** $p < .01$. "Fraction explained by home value declines" is the proportion of the average change in auto sales due to falling home values. Changes in variables are computed from 2007H1 to 2009H1. Regressions weight each observation by the number of households in the county. hh=household.

[17] We do not add the lagged change in vehicle sales because of lack of available data.

[18] We also use the average 2006 income per household, which is calculated from IRS data as the adjusted gross income in a county divided by the number of filers in that county.

Table 7 presents results analogous to Table 6 except that we use household rather than overall new auto sales in calculating the dependent variable. Each household AQR is in the range of 0.005 to 0.006 across the specifications. A one dollar decline in home values leads to a 0.5 to 0.6 cent decline in new auto purchases by consumers.

4 National Vehicle Sales and House Prices: State-Level Panel Analysis

4.1 Data and Econometric Model

In this section, we compare the relationship between house prices and personal new vehicle sales during the 2007-2009 recession period relative to the remainder of the past two decades. We cannot repeat the exact analysis because of data limitations.

We lack vehicle count data before 2007, and instead use personal consumption expenditures on motor vehicles in this section of the paper.[19] This variable is available at the state level at an annual rate beginning in 1997.[20]

This will move us from a county-level quarterly analysis to a state-level annual one. Let $G_{i,t}$ denote the per capita motor vehicle consumption in state i in year t. The raw data are nominal and we translate them into real series using the Consumer Price Index.

Our independent variable is based on county-level house price indices constructed by CoreLogic. Our annual variable is averaged across monthly observations. State-level house price indices are constructed by using the county-level averages weighted by the number of households in the county in 2007. We use house prices rather than home values on the right-hand side because the number of homes is not available for the entire sample.

Our estimation equation is

$$g^c_{i,t,\delta} = \phi_\delta p^c_{i,t,\delta} + \beta_\delta D_{i,t} + v_{i,t,\delta}$$

for $\delta = 1, ..., H$.

Census-region fixed effects, the lagged one-year growth rate of house prices and motor vehicle consumption are included as controls. We estimate the model using least squares at the 1-year, 3-year and 5-year horizons. To make estimates comparable, for each horizon we use the 5-year horizon sample (which implies that we drop some observations for the shorter-horizon regressions).

[19] Motor vehicle consumption also includes motor vehicle parts.

[20] Personal consumption expenditures of motor vehicles include net purchases of used vehicles, measured as dealer margins and net transactions, and the value of new vehicles purchased, as described in NIPA documentation. Dealer margins, for the most part, include the difference between the selling price and the dealer's acquisition cost. They also include wholesale margins for vehicles sold by wholesalers to dealers. According to NIPA documentation, net transactions consist primarily of the "wholesale value of purchase by persons from dealers less sales by persons to dealers."

Figure 6: Ratio of household debt to disposable personal income

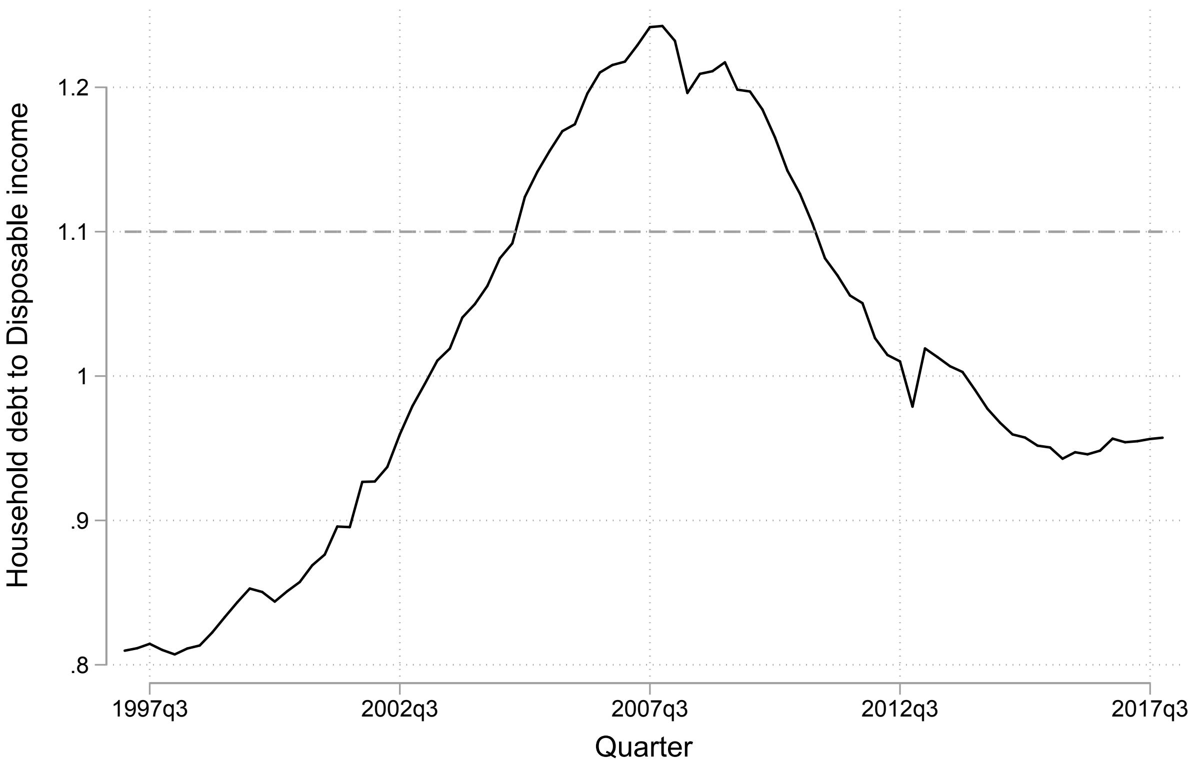

Notes: Data sources are the Federal Reserve Board and the Bureau of Economic Analysis.

4.2 Results

Columns (1) through (3) of Table 8 contain the elasticity estimates for the full sample. All three cumulative elasticities are positive and statistically different from zero. At the 1-year-horizon, the coefficient equals 0.92 (SE=0.10). A 1 percent increase in house prices over 1 year leads to a 0.92 percent increase in motor vehicle consumption over that year.

The cumulative elasticity is declining with the length of the horizon. At the 3-year horizon, the coefficient equals 0.66 (SE = 0.05). At the 5-year horizon, the coefficient equals 0.48 (SE = 0.04). Our estimates from the state-level panel are similar to those from the county-level recession-period results in Section 3. For example, at the 1-year horizon, the benchmark county-level estimate equals 1.08 (SE = 0.06).

Recall that in Section 3, the county-level elasticity implies an AQR that is statistically significant and positive, but quantitatively small. Since the state-level analysis finds a similar elasticity to that of the county-level data, that over the entire 1997-2017 period the response of motor vehicle consumption to house prices was also quantitatively small.

Next, we estimate the model for two different sub-periods: the high-leverage period (2005-2011) and the low-leverage period (1997-2004 and 2012-2017). Here we see some evidence that the auto sales response to house price changes was stronger in the high-leverage period relative to the low-leverage period. At the 3-year horizon, the cumulative elasticity equals 0.31 (SE=0.05) for the low-leverage sample. The corresponding value for the high-leverage sample equals 0.74 (SE=0.07). This is consistent with the evidence from Mian, Rao and Sufi (2013), which find that in the cross-section, counties with higher average leverage tend to have larger consumption responses to changes in house prices.

Examining Columns (6) and Columns (9) adds some nuance to this finding. At the 5-year horizon, the differences in elasticities has almost disappeared. At this horizon, the cumulative elasticities equal 0.43 (SE=0.05) and 0.55 (SE=0.07) for the low- and high-leverage samples, respectively.

Important dynamic considerations may influence how leverage interacts with housing wealth to effect how individuals adjust auto sales. Highly leveraged households may choose or be forced to react quickly to adjust auto purchases when housing wealth falls; however, following the shock their purchasing patterns begin to look more and more like the otherwise similarly affected low-leveraged households.

4.3 Miles Traveled and House Prices

Because of its durability, investment in vehicles provides a poor basis to measure the marginal utility of consumption of vehicle services. This marginal utility is better reflected by the service flow from the stock of vehicles. We contend that vehicle miles traveled provide a more direct measure of vehicle services provided. Therefore, we next estimate the relationship between the

Table 8: Cumulative response elasticities of auto sales to house price changes at various horizons (state-level panel, least squares)

	Full Sample			Low Lev. Sample			High Lev. Sample		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE	Coef./SE
1-yr HP growth	0.918***	-	-	0.417***	-	-	0.932***	-	-
	(0.098)			(0.085)			(0.110)		
3-yr cum HP growth	-	0.659***	-	-	0.305***	-	-	0.739***	-
		(0.052)			(0.050)			(0.073)	
5-yr cum HP growth	-	-	0.480***	-	-	0.425***	-	-	0.548***
			(0.039)			(0.052)			(0.069)
R2	0.42	0.58	0.62	0.30	0.50	0.61	0.35	0.56	0.63
N	714	714	714	357	357	357	357	357	357

Notes: The dependent variable is the accumulated percentage change in auto sales relative to a year $t-1$ baseline. * $p < .1$, ** $p < .05$, *** $p < .01$. Regressions weight each observation by the number of households in the state. Each estimate includes census-region fixed effects, the 1-year lag of the 1-year growth rate of auto sales and house prices as controls. Lev. = leverage.

growth in miles traveled and House prices.

Monthly miles traveled are available at the state level from the Federal Highway Administration beginning in 2006. We time-average monthly miles traveled up to the quarterly frequency. We similarly take quarterly averages of the monthly house price data and aggregate these to the state level using the weighted average of the number of households in each county.

We then run regressions where the dependent variables are the analogous change in the log of miles traveled at alternative horizons (1, 2 and 3 years). Our independent variable is the analogous house price variables at the corresponding horizons.

We estimate the regression via least squares with weights given by the number of households in each state and include state fixed effects in our baseline specification. These are presented in Columns (1), (3) and (5) of Table 9. Columns (2), (4), and (6) add additional controls: real income per household, the quarter-to-quarter growth rate in oil prices and the first lag thereof. Including these additional variables has very little impact on house price elasticities at every horizon.

Across each horizon and for alternative specifications, the elasticity is estimated within the range 0.019 and 0.038. Each is statistically different from zero at conventional confidence levels. Take, for example, Column (5), with an estimate of 0.038 (SE=0.016). If house prices on average increase by 10 percent accumulated over a 3-year period, then one would expect a 0.38 percent increase in vehicle miles traveled accumulated over the corresponding 3 years.

Thus, the effect of house prices on vehicle miles traveled is very small. By comparison, the cumulative elasticity of vehicle sales to house price changes over a 3-year horizon, from Table 3, equaled 0.60. The elasticity of house prices on vehicle sales, already explained as modest, is more than fifteen times as large as the effect of home values on miles traveled.

There is little evidence that house price changes significantly disrupted the service flow provided by vehicles. Thus, households were able to smooth the effects of house price shocks on their vehicle usage during the period. The economic model developed and calibrated in the next section is motivated by this observation: one can see a large change in investment in durable goods alongside only a small change in the flow of services delivered by the stock of durable goods.

5 National Vehicle Sales and Expected Income

5.1 Survey Evidence on Economic Conditions

We look to individual-level survey data to find other potential explanations for the auto market collapse besides house prices. The Michigan Survey of Consumers asks questions regarding consumers' likelihood of buying a car as well as their reason for their answer.

Figure 7 plots the fraction of respondents who state that, in the current quarter, it is an unfavorable time to purchase an auto over time. The figure shows an upward spike at the time of the auto market collapse. The survey also asked respondents to state why it is either a favorable

Table 9: Cumulative response elasticities of miles traveled to house price changes at various horizons (state-level panel, least squares)

	(1) Coef./SE	(2) Coef./SE	(3) Coef./SE	(4) Coef./SE	(5) Coef./SE	(6) Coef./SE
1-yr cum HP growth	0.035*	0.020	-	-	-	-
	(0.021)	(0.021)				
2-yr cum HP growth	-	-	0.036**	0.019	-	-
			(0.018)	(0.017)		
3-yr cum HP growth	-	-	-	-	0.038**	0.019
					(0.016)	(0.016)
Lagged income, oil prices	No	Yes	No	Yes	No	Yes
R2	0.01	0.04	0.01	0.05	0.02	0.06
N	2197	2147	2197	2147	2197	2147

Notes: The dependent variable is the cumulative percentage change in miles traveled relative to a year $t-1$ baseline. * $p < .1$, ** $p < .05$, *** $p < .01$. Regressions weight each observation by the number of households in the state. Each estimate includes state fixed effects. HP = house price.

or unfavorable time to purchase a car.

Figure 7: Fraction of individuals reporting that the current quarter is a bad time to buy an auto, 1995Q1 to 2014Q4

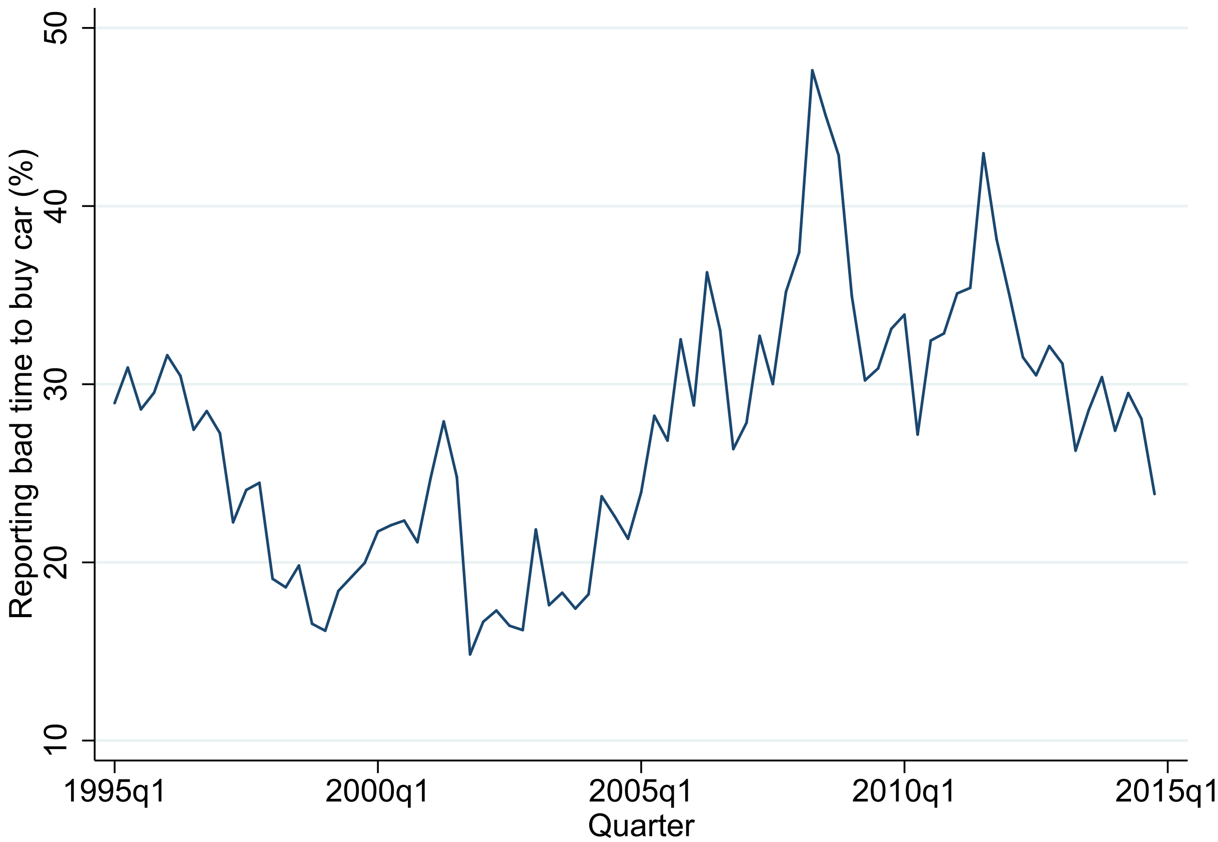

Notes: Source is the Michigan Survey of Consumers.

We take a subset of these responses and group them into one of two categories. The first category is credit and debt conditions, both at the individual level and nationwide.[21] The second is economic conditions.[22] Other categories, such as changes in the price of gasoline, are not included here. Figure 8 plots the fraction of respondents who answered that it was an unfavorable time for the reasons in to one of these categories.

The figure shows almost no change in the fraction motivated by credit and debt conditions and a dramatic increase in the fraction motivated by economic conditions at the time of the collapse.

[21]The specific answers are described in survey documentation as: debt or credit is bad; larger/higher down payment required; interest rates are high, will go up; and credit hard to get, tight money.

[22]The specific answers are: people cannot afford to buy now, times bad; people should save money, bad times ahead.

Figure 8: Fraction of individuals reporting credit/debt or economic conditions as reason it is an unfavorable time to buy an auto, 1995Q1 to 2014Q4

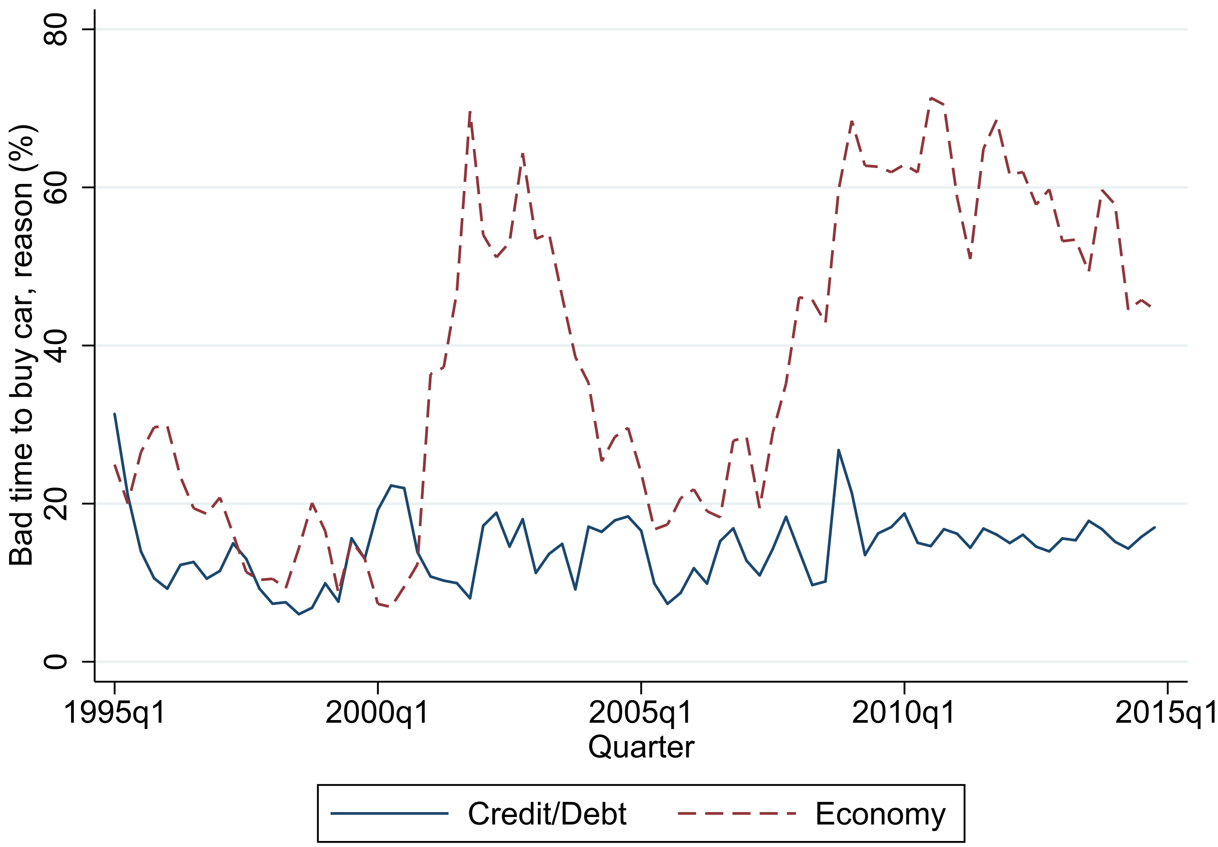

Notes: Source is the Michigan Survey of Consumers.

We take this as further evidence that debt and credit changes due to house price declines in that period did not significantly impact auto purchases. On the other hand, the "economic conditions" reason motivates our dynamic model of auto purchases, which uses shocks to current and expected future income as the driving force for the auto market collapse.

Respondents are also asked about the expected future growth of their own home values. This allows us to compare the relative importance of house price expectations versus perceived economic conditions on self-reported attitudes toward auto buying. We estimate a probit regression of the likelihood of buying a car dummy variable using a panel of respondents between 2004Q1 and 2018Q2. The left-hand-side variable equals 1 if the respondent answers that it is a favorable time to buy a car in the current quarter. The right-hand-side variables are the expected percentage increases in own house prices, the probability of an individual losing their job over the next year, log income, and a dummy for whether an individual predicts favorable economic conditions over the next five years. Alternative specification include or exclude time and region fixed effects.[23]

Table 10: Marginal effect on probability of reporting that it is a good time to buy a car, 2004 - 2018

	(1) Coef./SE	(2) Coef./SE	(3) Coef./SE	(4) Coef./SE
Percentage increase	0.006***	0.005***	0.006***	0.005***
inhouse prices	(0.000)	(0.000)	(0.000)	(0.000)
Positive 5-yr	0.175***	0.176***	0.177***	0.174***
economic outlook	(0.006)	(0.006)	(0.006)	(0.006)
Prob of losing job	-0.076***	-0.074***	-0.071***	-0.075***
	(0.011)	(0.011)	(0.010)	(0.011)
Log of income	0.068***	0.067***	0.068***	0.068***
	(0.003)	(0.003)	(0.003)	(0.003)
Increase in house	-	-	-	0.015**
prices				(0.007)
Decrease in home	-	-	-	-0.016**
price				(0.006)
Region fixed effects	Yes	No	Yes	Yes
Quarter fixed effects	Yes	Yes	No	Yes
N	34427	34427	34427	34427

Notes: The dependent variable equals 1 if an agent views the present as a good time to buy a car. Data are from the Michigan Survey of Consumers. * $p < .1$, ** $p < .05$, *** $p < .01$.

Table 10 displays the marginal effects for when explanatory variables are set equal to their means in the sample. It shows that an individual's personal view about the overall economy is a key determinant of their attitude toward purchasing autos. Consider Column (1). Keeping other

[23]Respondents are classified into one of four regions.

variables at their means, an individual that moves their opinion about the economic performance from positive to negative reduces the probability of having a positive car buying attitude by more than 17 percent.

The regression also shows that although house price growth expectations influence car buying attitudes, the marginal effect is small. According to column (1), an individual who expects their house price to further decline by 1 percent from the average expectation sees an 0.5 percent reduction in having a positive attitude toward car buying. Columns (2) and (3) alter fixed effects. The results are nearly unchanged. Column (4) adds two dummy variables, whether the own house price increased over the past year and decreased over the next year. The coefficients are of the expected sign, but have little impact on the remaining coefficients.

5.2 The Idea and the Mechanism

Next, we develop a permanent income model augmented with an auto-purchase choice to illustrate the effect of expected income changes on the auto purchase decision. In the model, at multiple points over its lifetime, a household pays a fixed price to buy a new car. The utility associated with owning a car is decreasing in the vehicle's age. There are idiosyncratic shocks to income and aggregate shocks to the growth rate of economywide average income. In the model, car owners experience relatively small changes in the marginal disutility of holding on to an old vehicle when expected income falls. Delaying auto replacement is an effective way to smooth the path of the marginal utility of consumption in response to the negative shock.

The calibrated model exhibits a large short-run decline in new vehicle purchases in response to weaker expected income growth going forward. A slowing of the real income growth rate to -2 percent, similar to that experienced during the 2007-2009 recession, delivers an over 70 percent auto sales decline on impact. This decline in auto sales is large, similar to the roughly 40 percent decline experienced in the second half of 2008. In contrast, a model that simply treats auto purchases as part of non-durable consumption would have an elasticity that would be much too low to generate a large decline in auto sales during the episode.

We abstract from several real-world features of the auto market, such as car loans and leasing. The power of our approach is to show that, even absent these frictions, a largely standard permanent income model can quantitatively replicate salient features of the 2008 auto market collapse.

We do not directly model the housing decision. This is because, earlier in the paper, we establish that house price fluctuations explain only a small fraction of the auto sales decline. Moreover, for many individuals, house prices are unlikely to influence the auto buying choice. For a homeowner planning to stay put, a house price decline largely nets out to a zero effect because it reduces tangible wealth but also increases the user cost of staying in the home. Also, survey data indicate that very few individuals use home equity to purchase vehicles. For a renter not close to the margin of buying a home, negative house price changes have no direct effect on their own wealth and therefore auto purchases. The effects on consumption for the two remaining groups, renters close to buying homes and homeowners close to selling homes, work in opposite directions and therefore are likely to be largely offset in the aggregate.

Finally, we note that our paper's first two results—the quite limited role for oil prices and house prices in explaining the new auto sales decline— are established using data without bringing a specific economic model to the table. It might seem natural that we investigate the role of income and future income expectations using cross-sectional data as well. Unfortunately, highly disaggregate (e.g., county-level) future income expectations data are not available. As such, we change approaches by shifting to a calibrated economic model. Note, however, that we will use the limited available survey data on future income expectations in calibrating our model.

5.3 The Model

Our model consists of a unit mass of households indexed $i \in [0,1]$. Each household i earns an exogenous stochastic income, and maximizes lifetime utility by choosing a stream of savings, non-durable consumption, and vehicle purchases. We calibrate the model so that a period lasts one year. The household buys only newly produced (not pre-owned) cars.[24]

Let income be given by $\tilde{Y}_{i,t} = \exp(y_{i,t}) Z_t$, where Z_t indexes aggregate income and evolves according to $Z_t/Z_{t-1} = 1 + g_t$. Also, g_t evolves according to a two-state Markov chain, $\{g_L, g_H\}$, and $y_{i,t}$ evolves according to a first-order autoregression:

$$y_{i,t} = \rho y_{i,t-1} + \varepsilon_{i,t} \tag{2}$$

where the innovation $\varepsilon_{i,t}$ has mean zero, standard deviation σ_ε and is i.i.d over time and households. Let $y_{i,-1}$ and Z_{-1} be positive and given as initial conditions.

The expected utility function is

$$U_{i,t} = \sum_{j=0}^{\infty} \beta^j E_t \left[U_N\left(\tilde{C}_{i,t+j}\right) + \alpha U_D\left(v_{i,t+j}\right)\right]$$

where $\tilde{C}_{i,t}$ is consumption and $v_{i,t}$ is the vintage of the auto currently owned by the household, and U_N and U_D give the period utility of non-durable and durable goods, respectively. We assume the utility is increasing and concave in non-durable consumption and given as

$$U_N\left(\tilde{C}\right) = \log(\tilde{C}) \tag{3}$$

We further assume that each household owns exactly one car, and the utility of owning a car depends on the depreciated value of the car,

[24]Our calibration will match data on autos originally purchased new. Thus, one should think about the new and pre-owned car markets as segmented, with our analysis solely focused on the former.

$$U_D(v) = \log(\tilde{P}_{v,Car}) \tag{4}$$

The depreciated value of a vintage v car is

$$\tilde{P}_{v,Car} = (1-\delta_v)\tilde{P}_{0,Car} \tag{5}$$

where δ_v is the accumulated depreciation rate of a car that is v years old and $\tilde{P}_{0,Car}$ is the price of a new car.

Next, the individual's wealth $\tilde{W}_{i,t}$ evolves according to

$$\tilde{W}_{i,t+1} = (1+r)\,\tilde{W}_{i,t} + \tilde{Y}_{i,t} - \tilde{C}_{i,t} - (\tilde{P}_{0,Car} - \tilde{P}_{v,Car}) \times \mathbf{1}\,(v_{i,t+1} = 0)$$

where borrowing is not allowed (i.e., $\tilde{W}_{i,t} \geq 0$) and the price of buying a new auto is given as $(\tilde{P}_{0,Car} - \tilde{P}_{v,Car})$, which is the post trade-in cost of buying a new car.

To allow us to write the individual's problem in recursive form, we assume that the new car price is a constant fraction of the current average-income index: $\tilde{P}_{0,Car} = Z_t P_{0,Car}$. Without rising auto prices, as income trends upwards, the vehicle vintage distribution would pile up at $v = 0$. Furthermore, this relationship between auto prices and income is reflected in the data. The real price of autos generally increases over time; however, the price was flat during the 2007-2009 recession—the same time that income growth was very low. Note also that this assumption biases us towards finding a smaller auto purchase response to the shock, because the price effect during the low-income-growth period pushes individuals to purchase cars (which rise in expected price over time) sooner rather than later.

If we define $C_{i,t} = \tilde{C}_{i,t}/Z_{t-1}$ and $W_{i,t} = \tilde{W}_{i,t}/Z_{t-1}$, consumption and wealth respectively, then we can express household i's optimization problem recursively in the transformed system:

$$V(W, v, y, g) = \max\{V_N(W, v, y, g)\,, V_B(W, v, y, g)\}$$

where V_N and V_B denote the values associated with car retention (not buying a new car) and car replacement (buying a new car), respectively, with

$$V_N(W, v, y, g) = \max_{C,W'}\left\{\log(C) + \alpha\log(P_{v,Car}) + E\left[V\left(W', v+1, y', g'\right)|y, g\right]\right\}$$

subject to

$$W' + \frac{1}{1+g}C = \frac{1+r}{1+g}W + y$$

and

$$V_B(W, v, y, g) = \max_{C,W'}\left\{\log(C) + \alpha\log(P_{v,Car}) + E\left[V\left(W', 0, y', g'\right)|y, g\right]\right\}$$

subject to

$$W' + (P_{0,Car} - P_{v,Car}) + \frac{1}{1+g}C = \frac{1+r}{1+g}W + y$$

both with $W_i' \geq 0$. The prime superscript advances time by one period.

Note that

$$v' = \begin{cases} 0 & \text{if auto is purchased} \\ v+1 & \text{otherwise} \end{cases}$$

To understand the household's car buying decision, let $\hat{W} = \hat{W}(v; y, g)$ denote the level of W that leaves a household with income level y and current car with vintage v indifferent between car replacement and car retainment. The car replacement decision can then be written as

$$v' = \begin{cases} 0 & \text{if } W \geq \hat{W} \\ v+1 & \text{if } W \leq \hat{W} \end{cases} \tag{6}$$

5.4 Calibration

We start by assuming $\beta = 0.94$, which is a standard assumption.

Next, we assume $g \in (g_L, g_H)$ with transition matrix

$$\Pi = \begin{bmatrix} \pi_{L,L} & 1 - \pi_{L,L} \\ 1 - \pi_{H,H} & \pi_{H,H} = 1 \end{bmatrix}$$

We use data from the Michigan Survey of Consumers to choose the values $g_L = -0.02$ and $g_H = 0$. Figure 9 plots the average one-year-ahead expected real growth of personal income across respondent's to the Michigan survey. Expected growth is for the upcoming year and subtracts each respondents expected nominal income growth minus their expected inflation over the following year. The figure shows a clear break around the time of the auto sales collapse. The red-dashed line is the average annual expected growth in 2006 and 2007. It equals approximately 0 percent. The average annual expected growth over the following four years equals approximately -2 percent.

We assume that $\pi_{L,L} = 2/3$, which is consistent with a decline in income growth that would last 3 years in expectation, after which it would revert to its original growth path. Furthermore, we set $\pi_{H,H} = 1$ to indicate the economy is initially in steady state with a constant growth rate of $g_{SS} = g_H$, which implies an expected decline in permanent income equal to 3.1 percent.[25]

To calibrate r, we rely on wealth data. In the standard permanent income model, if the economy were in the high growth steady state, then given log preferences over non-durable consumption, the

[25]See the appendix for a derivation of this relationship.

Figure 9: Average one-year ahead real income growth expectations

Notes: Source is the Michigan Survey of Consumers. For each quarter, we calculate the difference between reported expected 1-year-ahead nominal income growth net of reported expected 1-year-ahead inflation.

real interest rate consistent with steady-state growth would be

$$1 + r^{*} = (1 + g_{SS}) / \beta$$

In our model, this would lead to over saving with average wealth being too high. For this reason, we introduce a wedge κ that lowers the rate of return on savings and thus discourages wealth accumulation

$$r = \left(\frac{1 + g_{SS}}{\beta}\right) - 1 - \kappa$$

where $\kappa = 0.005$, the values of g_{SS} and β, implies $r = 0.059$ and implies a 3.1 percent decline in expected permanent income by assuming $\pi_{L,L} = 0.66$.

Next, we calibrate the vehicle cost $P_{0,Car}$. The Kelly Blue Book price of a new vehicle in 2008 was approximately $27,000. Using a household income of $60,000, the average price of a car is 45% of household income, and thus we set $P_{0,Car} = 0.45$.

To calibrate the accumulated depreciation rate of a car, we look at the annual used car prices for a 2006 Chervolet Malibu and a 2006 Honda Accord from when they were sold new in 2006 through being 12 years old in 2018. We find that these cars roughly depreciated in value 30% the first year and then 15% on an accumulated basis every subsequent year. Thus, we set $1 - \delta_1 = 0.7$ and for $v > 1$ we set $1 - \delta_v = 0.7 \times 0.85^v$. [26]

Finally, we set $\rho = 0.95$ and calibrate the standard deviation of the household's income process, σ_ε, and the relative disutility, α, to match both the empirical liquid assets/income ratio and the distribution of car vintages, for households that own a car that they bought new.

The empirical liquid assets/income ratio comes from the 2007 Survey of Consumer Finances (SCF). Liquid assets are defined as the sum of money market accounts, checking accounts, saving accounts, call accounts, and prepaid cards (variable: liq). Income includes wage income, business farm income, rent income, interest and dividend income, self employed income, etc. (variable: income). In Table 11 we compare various percentiles of the liquid asset/income ratio in both the SCF data and the calibrated model. The SCF data are for the subsample of households that own at least one car, van, minivan, SUV, or pickup that they bought new. Weights are accounted for when calculating percentiles in the SCF data.

The car vintage distribution gives us the probability that a household x years later still owns a car that they originally bought new. Figure 10 plots both the empirical distribution and the distribution from our calibrated model. The empirical car vintage distribution comes for the 2007 Survey of Consumer Expenditures (CEX). To determine the distribution, we consider all automobiles, trucks, minivans, vans, and SUVs. There is an unexplained dip in the distribution for new cars in year 0, possibly due to incomplete data. In calibrating our model, we impute this value as

[26] Many reasons have been proposed for this asymmetry in depreciation values. One explanation is the lemons problem, which may be most severe for new cars, as shown in House and Leahy (2004).

Table 11: Liquid assets/income ratio

Percentile	SCF	Model
25th	0.04	0.00
50th	0.11	0.15
75th	0.30	0.46
90th	0.75	0.95
99th	3.80	2.72

Notes: Liquid assets/income Ratio = variable:liq/variable:income. SCF data are from 2007 and for the subsample of households that own at least one car, van, minivan, SUV, or pickup that they bought new. Weights are accounted for when calculating percentiles in the SCF data.

Table 12: Parameter Values

Parameter	Value	Description	Motivation
β	0.94	Discount factor	Standard value for annual model
r	0.059	Interest rate	Set a little below $(1+g_{SS})/\beta$
$P_{0,Car}$	0.45	Price of new car relative to income	New car price = 45% of average income
$1-\delta_1$	0.7	Car depreciation after 1 year	Car depreciates 30% first year
$1-\delta_v$	0.7×0.85^v	Car depreciation after v years	and 15% every subsequent year
ρ	0.95	Persistence of income level	Standard value
$g_H = g_{SS}$	0.00	Steady state income growth rate	From Michigan Survey of Consumers
g_L	-0.02	Steady state income growth rate	From Michigan Survey of Consumers
$\pi_{H,H}$	1.0	High to high transition probability	High income growth is the steady state
$\pi_{L,L}$	0.66	Low to low transition probability	Set to match income slowdown expectations
σ_ε	0.002	Standard deviation income shocks	Set to match liquid asset/income distribution from SCF
α	0.06	Relative utility consumption vs. car	Set to match car vintage distribution from CEX

a linear trend based on data for 1- and 2-year-old cars.

All the parameter values are summarized in Table 12.

5.5 The Solution Method and Policy Function

We solve the household's problem by discrete discounted dynamic programming. The wealth state space is discritized into $H = 500$ evenly spaced points between 0 and $\bar{W} = 80$, the car vintage distribution is capped at $J_{max} = 30$, and the $y_{i,t}$ process is discretized on a grid with $N = 8$ points using the Rouwenhorst method. The total state space has dimension 240,000 ($= H \times (2 \times N) \times (J+1)$). The value of $\bar{W}$ and J_{max} are picked to ensure that the results are invariant to small changes in them.

A household with wealth at or below $\hat{W}$ as defined in equation (6) will not replace its car. Figure 11 plots this optimal cutoff wealth (in logs) as a function of the individual's current income

Figure 10: Distribution of auto vintages, data and economic model

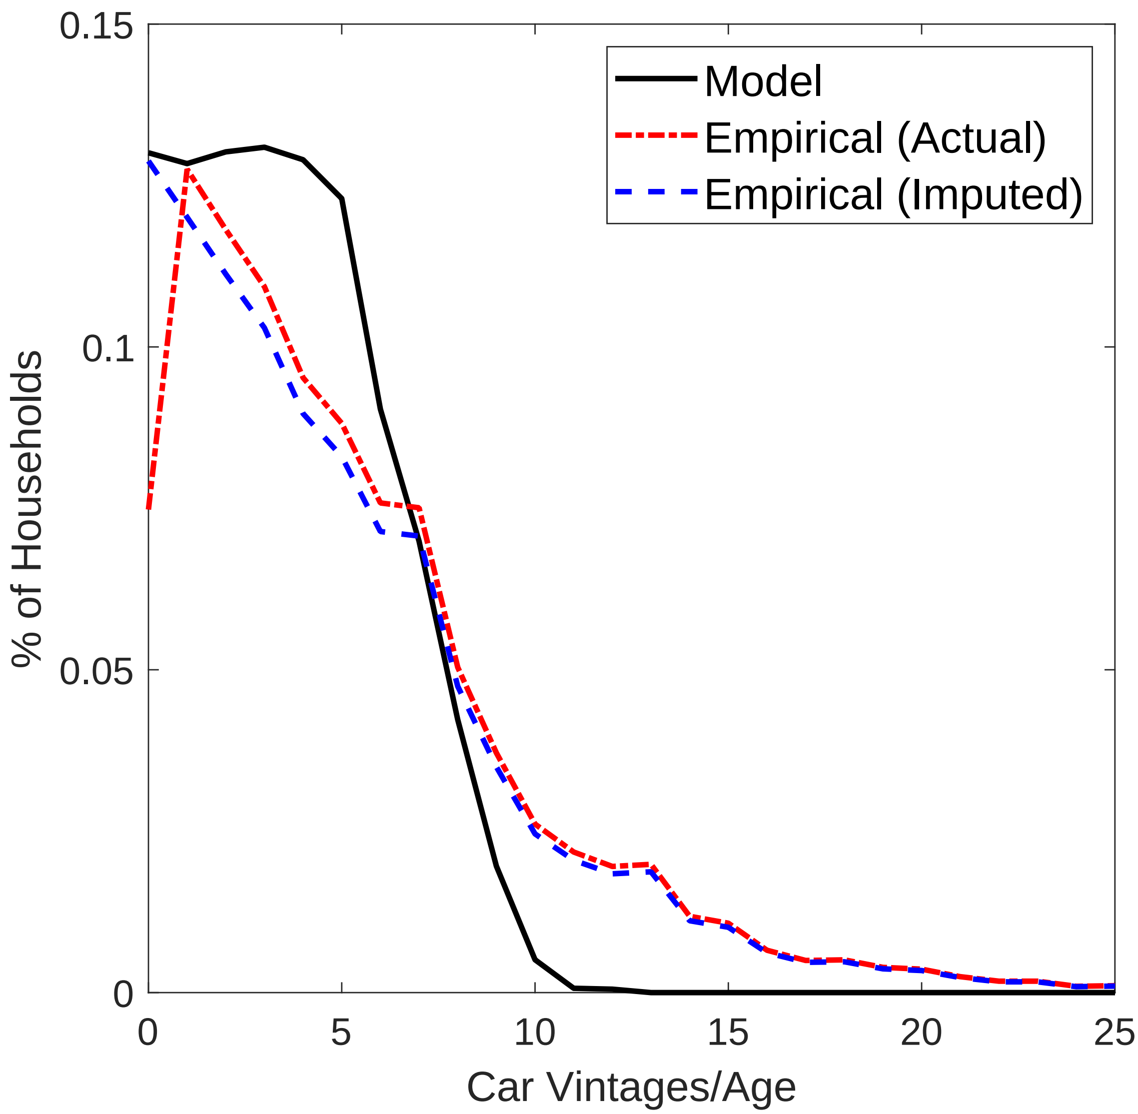

Notes: The non-solid lines plots the 2007 distribution of vehicles, originally purchased new, and the solid line plots the steady-state distribution from the economic model preceding the aggregate income growth rate shock. The data comes from the CEX.

(in logs) and vintage of their current auto. The cutoff wealths appear as numbers on the chart, where those wealths label contour lines. For example, the contour labeled "2.5" gives pairs of current incomes and vehicle vintages for which the cutoff log wealth equals 2.5. At these pairs, the individual chooses to replace their auto if current log wealth is above 2.5 and chooses not to if the current log wealth is below 2.5.[27] These policy functions assume the individual expects to remain in the high-average-income growth state forever.

Each contour line is downward sloping, which implies that as current income falls, the individual will replace older vehicles at the same cutoff level. Also, the cutoff wealth is falling as contour lines move rightward and upward. This implies that the set of wealth values for which replacement is optimal becomes larger as the individual has higher income or has an older-vintage current auto.

5.6 Results from the Economic Model

We simulate the outcome of a household's decision problem for a long history of $T + Q$ periods. For the first T periods, we assume Z grows at $g_{SS} = g_H = 0$. At period $T + 1$, average-income growth slows to $g_L =$ -2 percent and remains at -2 percent through $T + Q$. However, beginning at $T + 1$ the forecasted law of motion for Z_t evolves according to Π. That is, average income begins each period in the low-growth state and is expected to escape to the absorbing high-growth state with probability $1 - \pi_{H,H}$. In each period, the idiosyncratic determinant of income, $\epsilon_{i,t}$, is drawn according to the process described before.

Next, we set $T = 1000$, to ensure the initial conditions on a household's state variables do not affect are results, and $Q = 4$. We repeat this simulation for 100,000 households and present the average results across these households. This ensures that the idiosyncratic paths of $\epsilon_{i,t}$ largely cancel out in our final simulation.

Figure 12 plots the impulse responses for average income, auto sales, and non-durable consumption in response to an unanticipated slowdown in income growth that occurs at time zero. Each variable is plotted as an index with base year $t = -1$. In the years preceding the shock, all three variables grow at 0 percent annually. At period zero, average-income growth unexpectedly becomes -2 percent and remains as such through period 5 (although households predict a 33 percent chance each year that average-income growth will increase to its initial steady-state growth rate of 0 percent.)

Non-durable consumption falls approximately 5 percent in response to the growth slowdown. The decline in auto sales is much more dramatic. On impact, auto sales fall by over 70 percent. This is consistent with the large decline seen during the 2008 auto market collapse.

Our results show that a permanent income model augmented with an auto purchase choice and quantitatively calibrated can explain a large portion of the decline in car sales during 2008. As such, the model also illustrates a mechanism to explain the auto market crash. Because autos are a

[27] Vintages take on integer values, so the contour lines between integers on the horizontal axis reflect interpolations.

Figure 11: Policy function for auto replacement

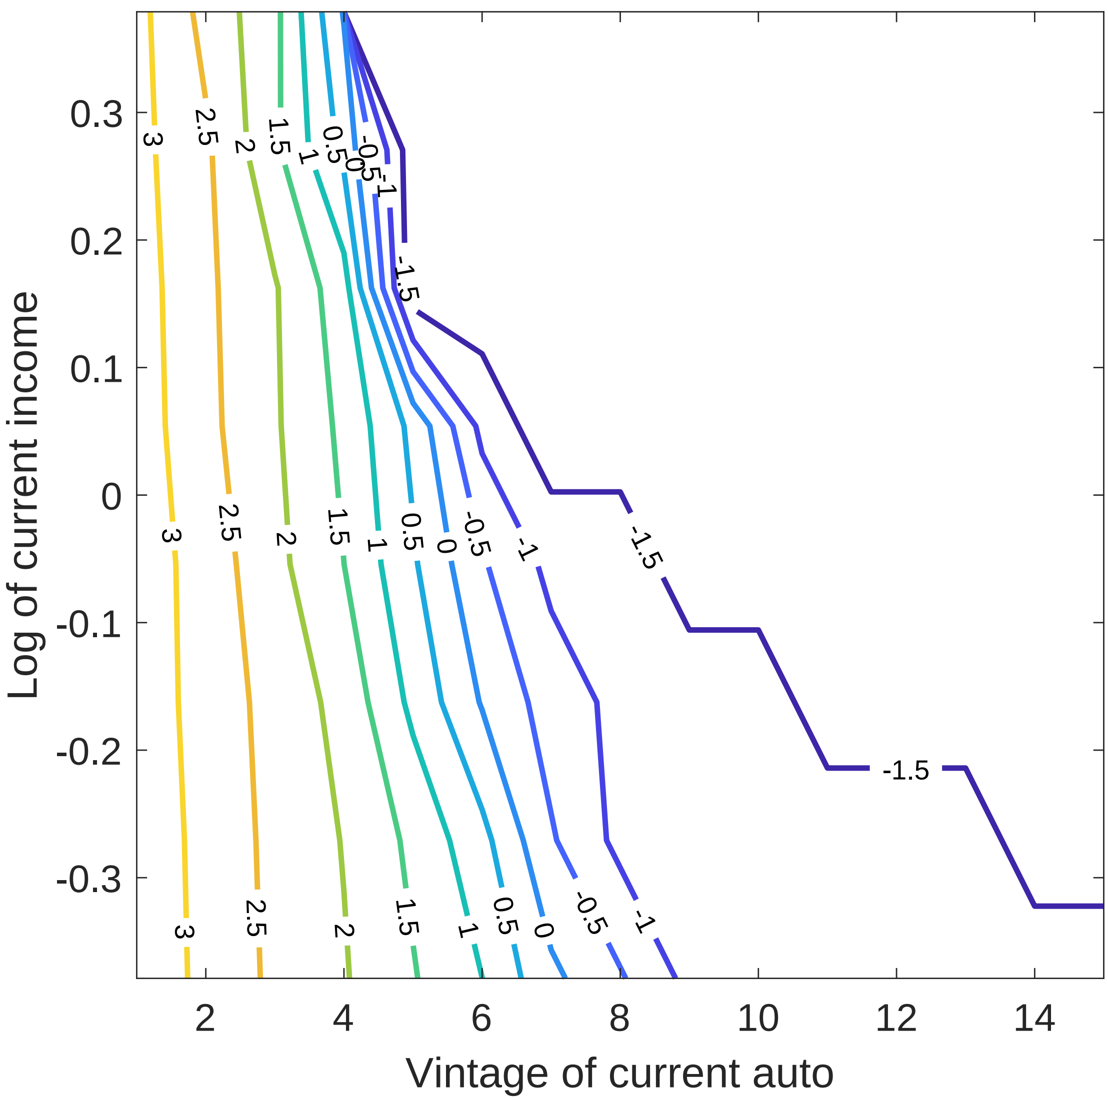

Notes: The optimal policy at each point in the state space is described by a cutoff log wealth. Contour lines reflect cutoffs, with the value of log wealth labeling the corresponding contour line.

Figure 12: Response of variables to income growth slowdown shock

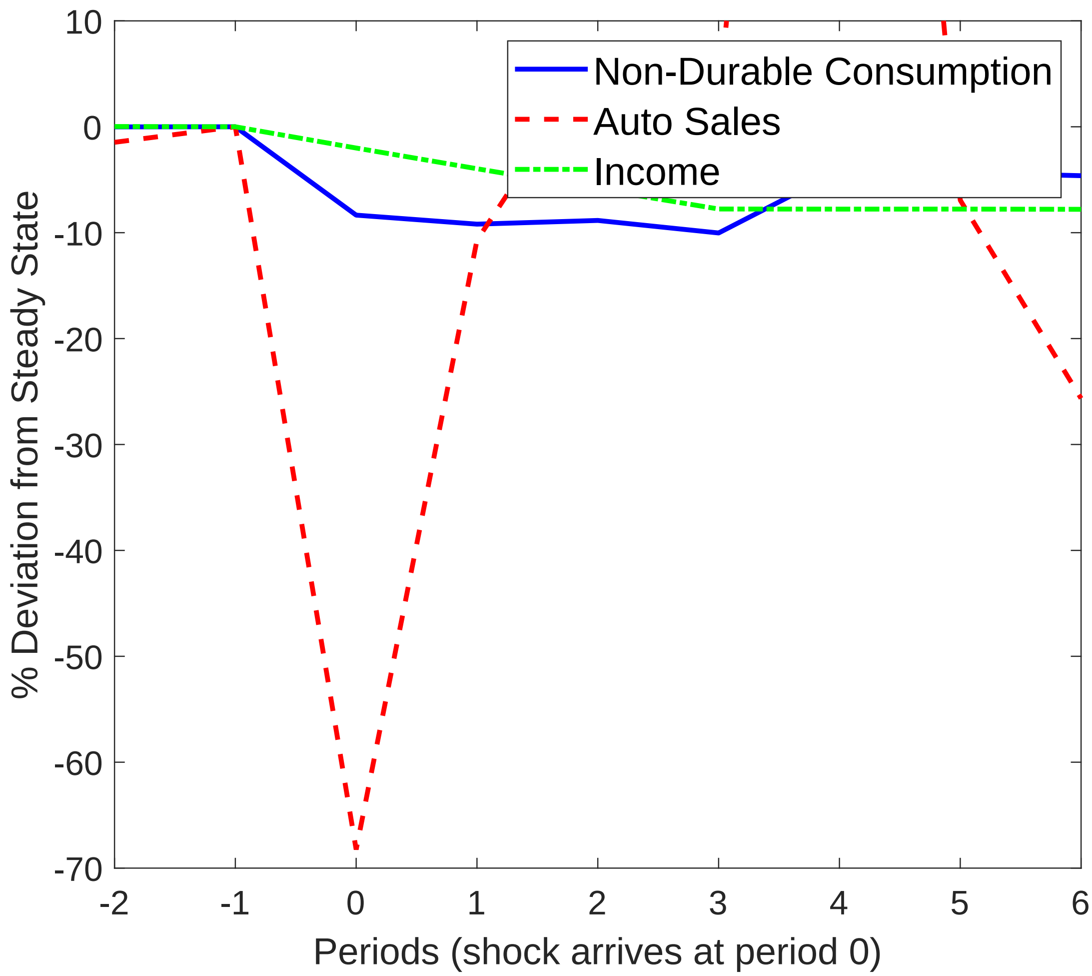

Notes: Paths are aggregated across 100,000 individual paths from the initial steady-state wealth, income, and auto vintage distributions.

durable good, many individuals respond to the decline in expected permanent income by delaying the replacement of their existing auto. The impact on the marginal disutility of having a slightly older car is smaller than the spike in marginal utility that would have occurred if an individual had instead dramatically reduced their non-durable consumption.

Leahy and Zeira (2005) previously developed a theoretical exploration of our paper's mechanism. They study the cyclical behavior of both durable and non-durable consumption. Their simplifying assumptions on preferences imply that consumers will not purchase small amounts of durable goods early in their lifetime. Therefore, the shocks can be entirely absorbed by the changes in the purchase timing of durable goods, and non-durable consumption can be fully insulated from the aggregate shocks.

Similar to Leahy and Zeira (2005), we study the role of the infrequent, discrete car replacement decisions in explaining consumption behaviors over the business cycle, but assume a more general setting of preferences that implies households consume both durable and non-durable goods in every period of their lifetime. Furthermore, we highlight how the standard permanent income model augmented with a discrete car replacement decision can explain the large fall in auto sales in the second half of 2008. We focus on the income shocks, which include both idiosyncratic shocks to individual income and aggregate shocks to the growth rate of economy wide average income. The delay in the timing of auto replacement allows us to explain the large decline in auto purchases in response to negative income shocks, which is essentially the mechanism of timing decisions as emphasized in Leahy and Zeira (2005). Unlike their model, where non-durable consumption is fully insulated from the aggregate shocks, the non-durable consumption in our model falls following a negative income shock, which is consistent with the 2008 experience.

5.7 Economic Model Augmented with Richer Belief Structure

We augment our model with a richer belief structure to allow for more persistence in beliefs, and thus a more persistent response to the income shock.

Households in this model believe there are four potential states $\{g_L^T, g_L, g_H^T, g_H\}$. g_L and g_H represent the high- and low-growth states as before. On the other hand, g_L^T and g_H^T represent low- and high-growth states, where the value of income growth $g_L^T = g_L$ and $g_H^T = g_H$, but the belief on how persistence these states are may be different than for g_L and g_H.

The transition matrix for switching between these states is as follows:

$$\Pi^B = \left\{ \begin{array}{cccc} \pi_{L^T L^T} & 0 & 0 & 1 - \pi_{L^T L^T} \\ 0 & \pi_{LL} & 1 - \pi_{LL} - \pi_{LH} & \pi_{LH} \\ 0 & 1 - \pi_{H^T H^T} & \pi_{H^T H^T} & 0 \\ 0 & 0 & 0 & 1 \end{array} \right\}$$

Furthermore, households in our economy have a hard time distinguishing between g_L^T and g_L

Table 13: Parameter Values

parameter	$\pi_{L^TL^T}$	π_{LH}	$\pi_{H^TH^T}$	ξ_{HL}	ξ_{LH}
estimate	0.062	0.018	0.431	0.289	0.011

and also g_H^T and g_H, so when a change in state occurs only a fraction ξ of households know the true state immediately. Household beliefs subsequently correct at a rate ξ too. In particular, ξ_{HL} represents the probability of knowing the true state when the economy changes from a high growth state to a low growth state and ξ_{LH} represents the probability in the opposite scenario.

Setting $\pi_{LL} = 0.66$ as before, we calibrate the remaining five parameters, $\{\pi_{L^TL^T}, \pi_{LH}, \pi_{H^TH^T}, \xi_{HL}, \xi_{LH}\}$, to minimize the distance between the model simulation and 65 percent of the actual auto sales decline that started in Sep 2008 through Aug 2015. We pick 65 percent, as from our empirical analysis that is the portion unexplained by oil and housing prices. Table 13 gives the calibrated values and Figure 13 plots the results.

As can be seen in the figure, our simple model augmented with a slightly richer set of beliefs is able to explain 65 percent portion of the decline in auto sales in this period. This positive result further provides evidence that beliefs about future income declines were driving the large drop in sales witnessed during the Great Recession.

6 Conclusion

Nationwide, new auto sales collapsed in 2008. Using a calibrated, dynamic stochastic consumption-savings model, we show that a widely experienced negative shock to permanent income is a strong candidate explanation for the collapse. The explanation is consistent with the permanent income hypothesis (PIH) adapted to include infrequent, discrete durable goods purchases. House price declines, on the other hand, explain only a small part of the auto sales decline.

A related explanation for the decline in auto sales is the increase in uncertainty that many researchers have associated with the 2007-2009 recession.[28] Bloom (2009) presents a model where irreversible investment in durable goods causes an increase in uncertainty to reduce purchases of durables. We note that a new vehicle purchase exhibits an aspect of irreversibility, because the resale value of a newly bought new auto falls dramatically immediately. Consistent with this story, Hassler (2001) finds auto expenditures in the U.K. declined dramatically with increases in uncertainty, proxied by stock market volatility.[29]

Our estimates speak to two important concepts in the economics of consumption: PIH and

[28] See for example Baker, Bloom and Davis (2016) and Jurado, Ludvigson and Ng (2015).
[29] See also Bertola, Guiso and Pistaferri (2005).

Figure 13: Response of variables to income growth slowdown shock

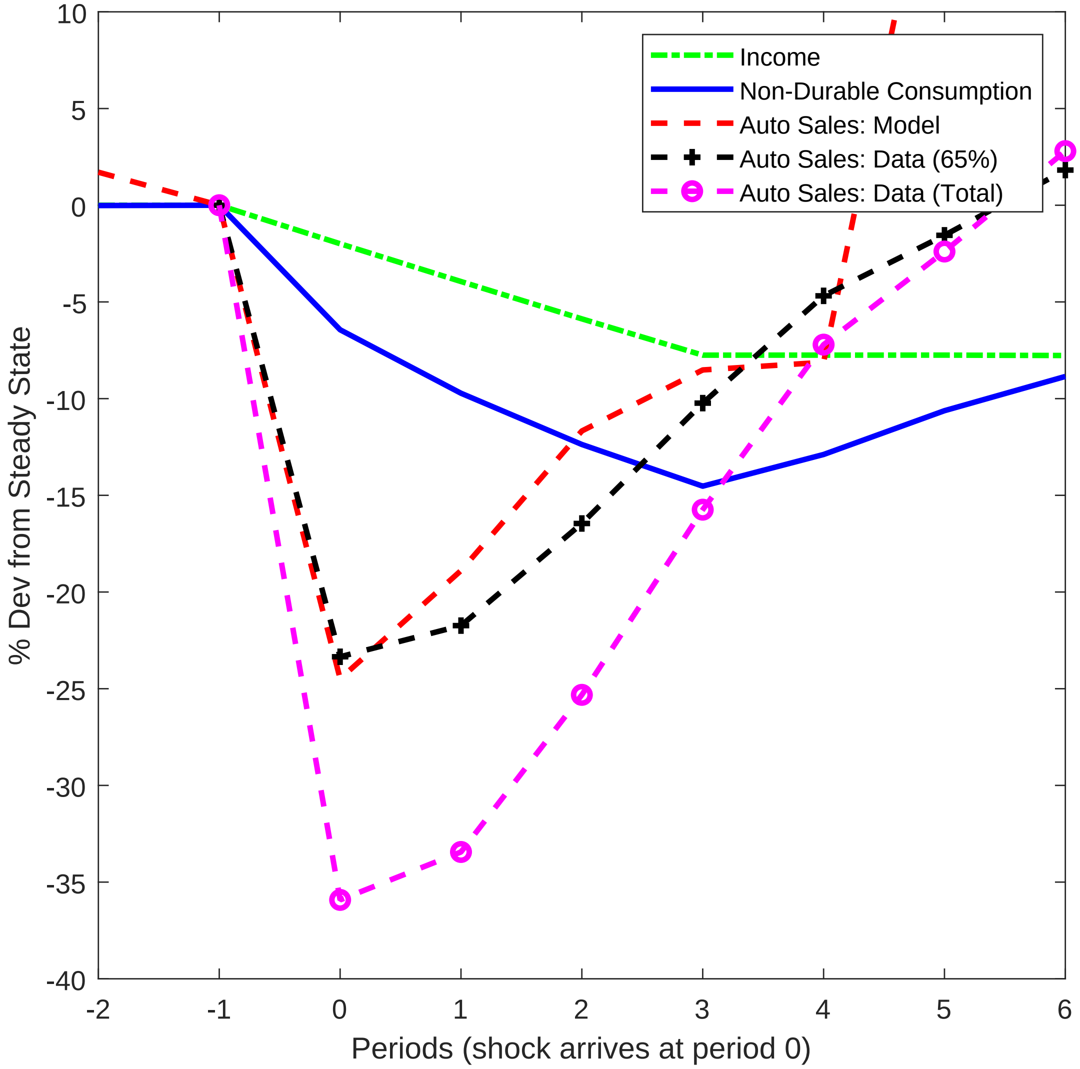

Notes: Paths are aggregated across 100,000 individual paths from the initial steady-state wealth, income, and auto vintage distributions.

consumption risk sharing. Our regressions do not directly test either theory; however, our findings do not violate either theory in a quantitatively important way.

In its modern form, the PIH states that households attempt to smooth the marginal utility of consumption in response to shocks. At a passing glance, it might seem that a 40 percent decline in auto sales would be an obvious violation of the PIH. That view, however, would confound investment in durable goods with the flow of services of the stock of durable goods. Based on the generally smooth series for aggregate vehicle miles traveled before, during and after the recession, one could conclude that the marginal utility from the services delivered by the stock of autos was little affected by the shock that drove down house prices.

One implication of consumption risk sharing is that the marginal utility of consumption is equated across regions even though shocks influence various regions with different intensities. The decline in house prices was very heterogeneous across U.S. counties. A strong positive correlation between house price changes and auto sales changes would have indicated a breakdown of cross-region consumption insurance. This strong positive correlation was not observed in the 2007-2009 recession, as evidenced by our low estimated AQR. From a broader perspective, investment in durables provides a poor measure of the marginal utility of consumption of durables as explained above. Therefore, without additional structure on preferences or else different data, examining auto sales regressions may constitute an inadequate approach for studying consumption risk sharing.

If, as we conjecture, households delayed replacing their existing autos with new ones in response to economic shocks, then one could see utility-reducing changes on households apart from miles traveled. For example, households may have spent additional dollars and time on maintaining used cars that they would have otherwise replaced. The aggregate evidence for this channel is weak: based on the Census Annual Retail Sales data, spending at stores supplying automotive parts, accessories and tires was nearly unchanged during the period.

Another possibility is that, although vehicle miles were smooth during the period, the typical quality of the driving experience could have been diminished because households did not replace their existing cars during the recession. For instance, those putting off buying a new car in cold weather climates may not have been able to enjoy heated seats, which were becoming more common in new vehicles during this period.

Even though the service flow from auto usage (measured by miles traveled) was almost imperceptibly disrupted during the 2007-2009 recession, the fall in auto demand did have dire consequences for those working in the auto and related industries. The associated fall in labor demand from these and other durable goods industries helped drive the national unemployment rate above 10 percent. This, paired with imperfect labor income risk sharing, meant the bulk of the welfare costs from the downturn was borne mainly by those who lost their jobs.

Understanding the strong sensitivity of demand for durable goods to economic shocks and how this interacts with these sectors' labor demand under imperfect labor income risk sharing merits further research.

References

Adda, J. and R. Cooper (2000), "Balladurette and Juppette: A Discrete Analysis of Scrapping Subsides," *Journal of Political Economy*, 108, 4, 778-806.

Attanasio, O., K. Larkin, M. Padula, and M. Ravn(2022), "(S)Cars and the Great Recession," *Econometrica*, 90, 5, 2319-2356.

Baker, S., N. Bloom and S. Davis (2016), "Measuring Economic Policy Uncertainty," *Quarterly Journal of Economics*, 131, 4, 1593-1636.

Benmelech, E., R. Meisenzahl and R. Ramcharan (2017), "The Real Effects of Liquidity During the Financial Crisis: Evidence from Automobiles," *Quarterly Journal of Economics*, 132, 1, 317-365.

Berger, D., V. Guerrieri, G. Lorenzoni and J. Varva (2018), "House Prices and Consumer Spending," *Review of Economic Studies*, 85, 3, 1502-1542.

Bernanke, B. (1984), "Permanent Income, Liquidity, and Expenditure on Automobiles: Evidence From Panel Data," *Quarterly Journal of Economics*, 99, 3, 587-614.

Bertola, G., L. Guiso and L. Pistaferri (2005), "Uncertainty and Consumer Durables Adjustment," *The Review of Economic Studies*, 72, 4, 973-1007.

Bloom, N. (2009), "The Impact of Uncertainty Shocks," *Econometrica*, 77, 3, 623-685.

Bucks, B., A. Kennickell, T. Mach and K. Moore (2009), "Changes in U.S. Family Finances from 2004 to 2007: Evidence from the Survey of Consumer Finances," *Federal Reserve Bulletin*, 95, A1-A56.

Corbae, D. and E. Quintin (2015), "Leverage and the Foreclosure Crisis," *Journal of Political Economy*, 123, 1, 1-65.

De Nardi, M., E. French and D. Benson (2012), "Consumption and the Great Recession," *Economic Perspectives*, Federal Reserve Bank of Chicago, QI, 1-16.

Disney, R., J. Gathergood and A. Henley (2010), "House Price Shocks, Negative Equity, and Household Consumption in the United Kingdom," *Journal of the European Economic Association*, 8, 6, 1179-1297.

Dupor, B. and R. Guerrero (2017), "Local and Aggregate Fiscal Policy Multipliers," *Journal of Monetary Economics*, 92, 16-30.

Dupor, B. and R. Guerrero (2018), "The Aggregate and Relative Economic Effects of Medicaid and Medicare Expansions," Working Paper.

Eberly, J. (1994), "Adjustment of Consumers' Durable Stocks: Evidence from Automobile Purchases," *Journal of Political Economy*, 102, 3, 403-436.

Edelstein, P. and L. Kilian (2009), "How sensitive are consumer expenditures to retail energy prices?," *Journal of Monetary Economics*, 56, 6, 766-779.

Favilukus, J., S. Ludvigson and S. Van Nieuwerburg (2016), "The Macroeconomic Effects of Housing Wealth, Housing Finance and Limited Risk Sharing in General Equilibrium," *Journal of Political Economy*, 125, 1, 140-223.

Garriga, C. and A. Hedlund (2020), "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession," *American Economic Review*, 110, 6, 1603-34.

Gavazza, A. and A. Lanteri (2018), "Credit Shocks and Equilibrium Dynamics in Consumer Durable Goods Markets," *Review of Economic Studies*, 88, 6, 2935-2969.

Glover, A., J. Heathcote, D. Krueger and J. Ríos-Rull (2020), "Intergenerational Redistribution in the Great Recession," *Journal of Political Economy*, 128, 10, 3730-3778.

Hamilton, J. (2008), "The Oil Shock and the Recession: Part 1," *Econbrowser*, Dec. 31.

Hamilton, J. (2009), "Causes and Consequences of the Oil Shock of 2007-08, *Brookings Paper on Economic Activity*, Economic Studies Program, The Brookings Institution, 40, 1 , 215-283.

Hassler, J. (2001), "Uncertainty and the Timing of Automobile Purchases," *The Scandinavian Journal of Economics*, 103, 2, 351-366.

House, C. and J. Leahy (2004), "An Ss Model with Adverse Selection," *Journal of Political Economy*, 112, 581-614.

Jorda, O. (2005), "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 1, 161-182.

Jurado, K., S. Ludvigson and S. Ng (2015), "Measuring Uncertainty," *American Economic Review*, 105, 3, 1177-1216.

Kaplan, G., K. Mitman and G. Violante (2016), "Non-Durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data," *Journal of Public Economics*, 189.

Kaplan, G., K. Mitman and G. Violante (2017), "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 128, 9, 3285-3345.

Leahy, J. and J. Zeira (2005), "The Timing of Purchases and Aggregate Fluctuations," *Review of Economic Studies*, 72, 1127-1151.

Lehnert, A. (2004), “Housing, Consumption, and Credit Constraints,” Finance and Economics Discussion Series Working Paper No. 2004-63.

McCully, B., K. Pence and D. Vine (2019), “How Much are Car Purchases Driven by Home Equity Withdrawal? Evidence from Household Surveys,” *Journal of Money Credit and Banking,* 51(5)

Mian, A., K. Rao and A. Sufi (2013), “Household Balance Sheets, Consumption, and the Economic Slump,” *The Quarterly Journal of Economics*, 128(4), 1687-1726.

Mian, A., A. Sufi (2018), “Household Balance Sheets, Consumption, and the Economic Slump: A Clarification,” working paper.

Nakamura, E. and J. Steinsson (2014), “Fiscal Stimulus in a Monetary Union: Evidence from U.S. Regions,” *American Economic Review,* 104(3), 753-92.

Poterba, J. (1991), “House Price Dynamics: The Role of Tax Policy and Demography,” *Brookings Papers on Economic Activity*, 2, 143-203.

Ramey, V. and S. Zubairy (2017), “Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data,” *Journal of Political Economy,* forthcoming.

Thompson, D. (2012), “Rick Santorum Is Right: Gas Prices Caused the Great Recession,” *The Atlantic*, February.

Zimmerman, M. (1998), “A View of Recessions from the Automotive Industry, Beyond Shocks: What Causes Business Cycles,” Federal Reserve Bank of Boston Conference Proceedings.

A Appendix

A.1 Additional aggregate car price and volume data

Figure 14 plots the time series for the average new vehicle price from the BTS. We use these prices to assign dollar values to the quantity of vehicles sold. Our data do not include a 2011 price; therefore, we use the 2010 price in this year.

Figure 14: Average new car auto price

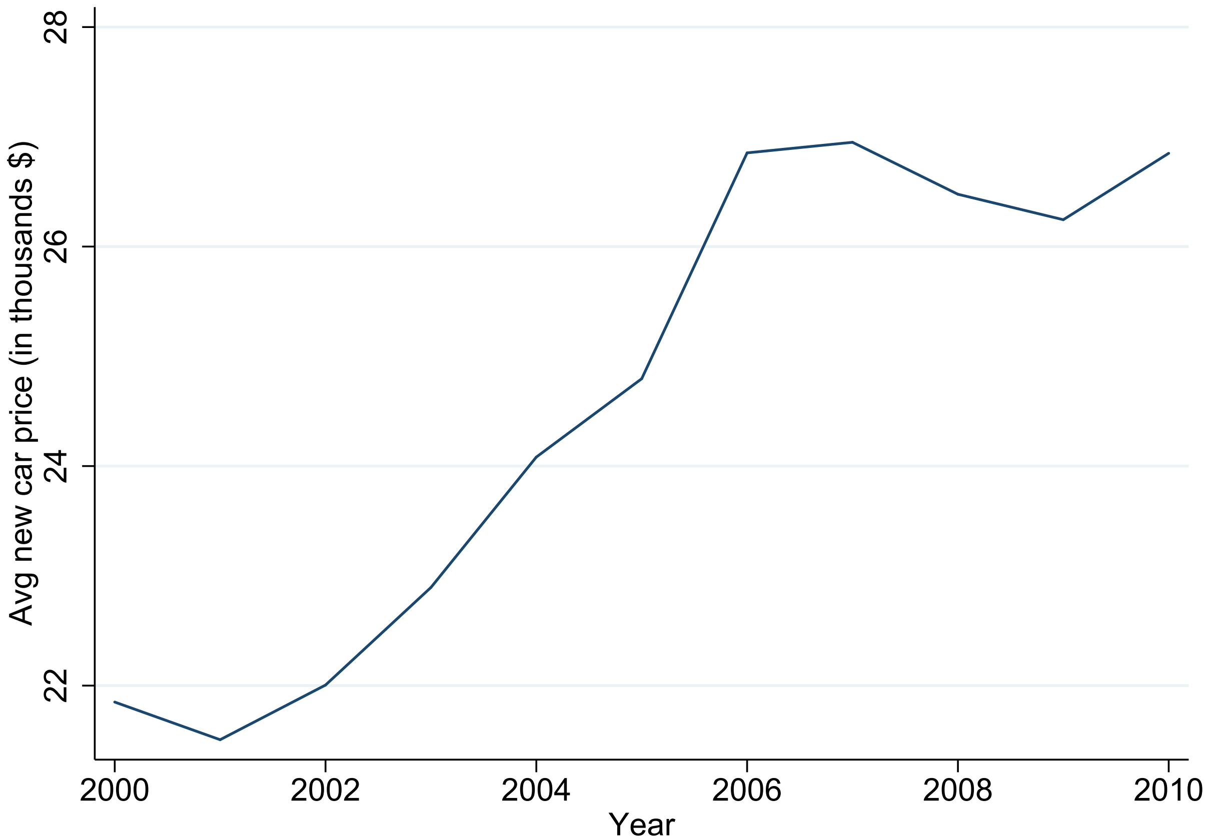

Notes: National Transportation Statistics, Values measured in thousands of current dollars.

Figure 15 plots the percent of new cars sold to consumers divided by the total new cars sold to consumers, businesses and governments. We use this share to allocate each county's total new vehicle sales to the quantity of sales to consumers.

Figure 15: Share retail sales of new cars to consumers

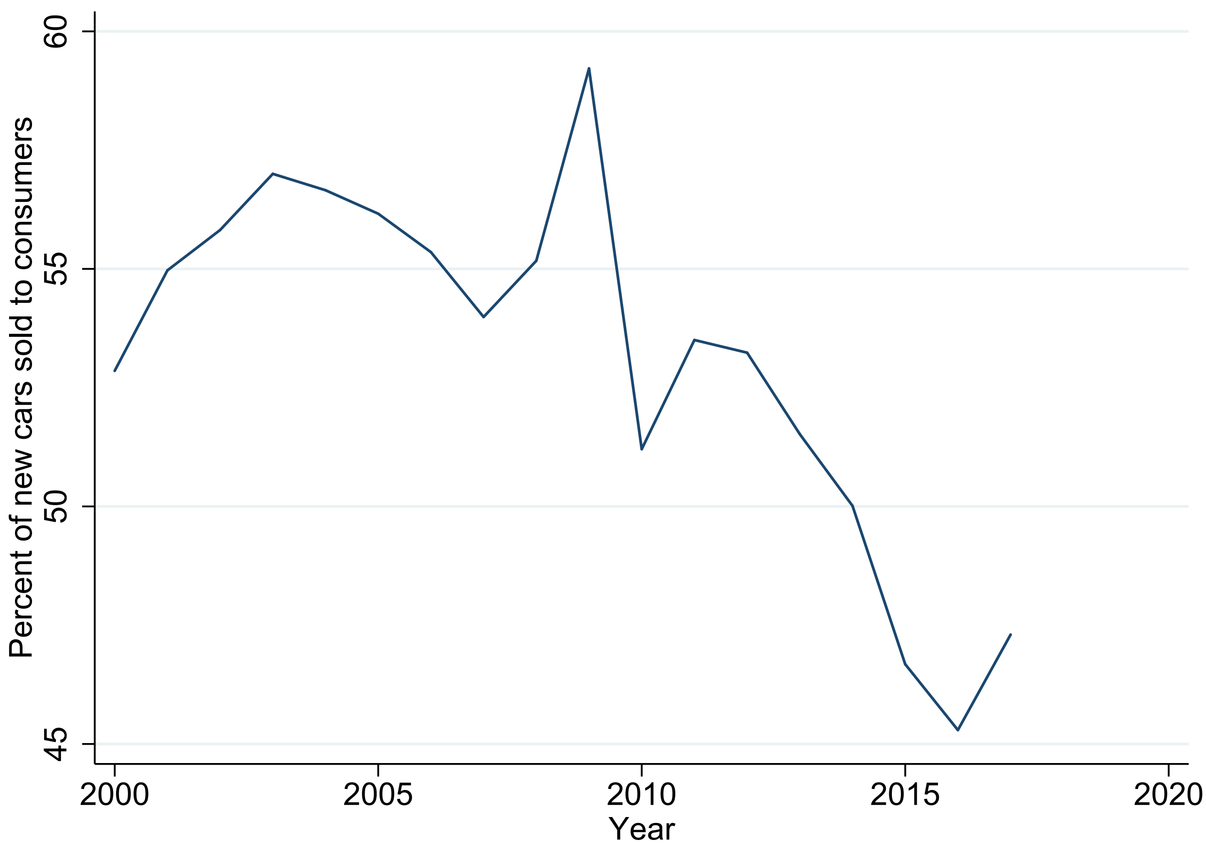

Notes: National Transportation Statistics, Tables 1-18. The numerator is the quantity of cars sold to consumers as a percentage of cars sold to consumers, businesses and governments.

A.2 Calibrating the average-income process

The values for r, g_H and g_L are given in the body of the paper. Suppose that individuals hold that income will remain in the low state for an expected length of 3 years. Given the Markov process for income growth, this implies $\pi_{L,L} = 0.66$.

To calibrate the expected present value of income in the low- and high-income process.

$$
\begin{aligned}
EPV_L &= 1 + \frac{\pi_{L,L}\,(1+g_L)}{1+r} EPV_L + \frac{(1-\pi_{L,L})\,(1+g_H)}{1+r} EPV_H \\
EPV_H &= \frac{1+r}{r-g_H}
\end{aligned}
$$

Solving for the ratio of the two income processes:

$$
\frac{EPV_L}{EPV_H} = \frac{r-g_H}{1+r-\pi_{L,L}\,(1+g_L)} + \frac{(1-\pi_{L,L})\,(1+g_H)\,(1+r)}{(r-g_H)\,[1+r-p\,(1+g_L)]}
$$

This equation implies that the expected decline in permanent income at the time of the shock equals 3.1 percent.

Table 14: Household new vehicle acquisition response (AQR) of sales to changes in home values, unweighted specification, county-level panel

	(1) Coef./SE	(2) Coef./SE	(3) Coef./SE
1-yr Cum HP Change	0.009***	-	-
	(0.001)		
2-yr Cum HP Change	-	0.007***	-
		(0.000)	
3-yr Cum HP Change	-	-	0.007***
			(0.000)
Vehicles sold (lag	-0.310**	-0.645**	-0.996**
change)	(0.140)	(0.287)	(0.479)
HP (lag change)	-0.003**	-0.002	-0.007
	(0.001)	(0.003)	(0.004)
Cash for Clunker	0.109***	0.216***	0.243**
fixed effect	(0.031)	(0.066)	(0.112)
Quarter	0.007***	0.078***	0.176***
	(0.002)	(0.005)	(0.008)
R^2	0.23	0.29	0.30
N	13673	13673	13673

Notes: The dependent variable is the cumulative change in vehicle sales to households at the appropriate horizon. * $p < .1$, ** $p < .05$, *** $p < .01$. Regressions include seasonal fixed effects (not reported). Standard errors are robust with respect to heteroskedacity and autocorrelation.

A.3 Additional Regression Specifications

This section contains additional regression results. The details of each specification are reflected in the table titles and notes.

Table 15: Household new vehicle acquisition response (AQR) of sales to changes in home values, time fixed effects included, county-level panel

	(1)	(2)	(3)
	Coef./SE	Coef./SE	Coef./SE
1-yr Cum HP Change	0.004***	-	-
	(0.001)		
2-yr Cum HP Change	-	0.005***	-
		(0.001)	
3-yr Cum HP Change	-	-	0.006***
			(0.001)
Vehicles sold (lag change)	-0.195	-0.441	-0.673
	(0.222)	(0.462)	(0.775)
HP (lag change)	0.002	0.009*	0.014*
	(0.002)	(0.005)	(0.008)
R^2	0.19	0.20	0.20
N			
N	13673	13673	13673

Notes: The dependent variable is the cumulative change in vehicle sales to households at the appropriate horizon. * $p < .1$, ** $p < .05$, *** $p < .01$. Regressions weight each observation by the number of households in the county and include time fixed effects. Standard errors are robust with respect to heteroskedacity and autocorrelation.

Table 16: Household new vehicle acquisition responses (AQR) of sales to changes in home values, Cash for Clunkers dummy in Fall 2009 only, county-level panel

	(1)	(2)	(3)
	Coef./SE	Coef./SE	Coef./SE
1-yr Cum HP Change	0.008***	-	-
	(0.001)		
2-yr Cum HP Change	-	0.006***	-
		(0.001)	
3-yr Cum HP Change	-	-	0.007***
			(0.001)
Vehicles sold (lag change)	-	-	-
HP (lag change)	0.000	0.007	0.009
	(0.003)	(0.007)	(0.011)
Cash for Clunker	0.287***	0.477***	0.639***
(Fall 09) fixed effect	(0.059)	(0.125)	(0.207)
Quarter	0.010***	0.094***	0.197***
	(0.002)	(0.005)	(0.010)
R^2	0.13	0.18	0.19
N	13673	13673	13673

Notes: The dependent variable is the cumulative change in vehicle sales to households at the appropriate horizon. * $p < .1$, ** $p < .05$, *** $p < .01$. Regressions weight each observation by the number of households in the county and includes seasonal fixed effects (not reported). Standard errors are robust with respect to heteroskedacity and autocorrelation.