The 2008 U.S. Auto Market Collapse*

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

August 14, 2018

Abstract

New vehicle sales in the U.S. fell nearly 40 percent during the last recession, causing significant job losses and unprecedented government interventions in the auto industry. This paper explores two potential explanations for this decline: falling home values and falling households' income expectations. First, we establish that declining home values explain only a small portion of the observed reduction in vehicle sales. Using a county-level panel from the episode, we find: (1) A one-dollar fall in home values reduced new vehicle spending by about 0.9 cents; and (2) Falling home values explain approximately 19 percent of the aggregate vehicle spending decline. Next, examining state-level data from 1997-2016, we find: (3) The short-run responses of 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 from miles traveled, responds very little to house price shocks. We also detail the sources of the differences between our findings (1) and (2) from existing research. Second, we establish that declining 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 liquid wealth and distribution of auto vintages and exhibits a large vehicle sales decline in response to a moderate decline in expected permanent income. In response to the decline in permanent income, households delay replacing existing vehicles, allowing them smooth the effects of the income shock without significantly adjusting the service flow from their vehicles. Combining our negative results regarding housing wealth with our positive model-based findings, we interpret the auto market collapse as consistent with existing permanent income based approaches to durable goods consumption (e.g., Leahy and Zeira (2005)).

Preliminary and incomplete draft: Please do not circulate without permission of the authors.

*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. First draft: December 2017.

† 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.

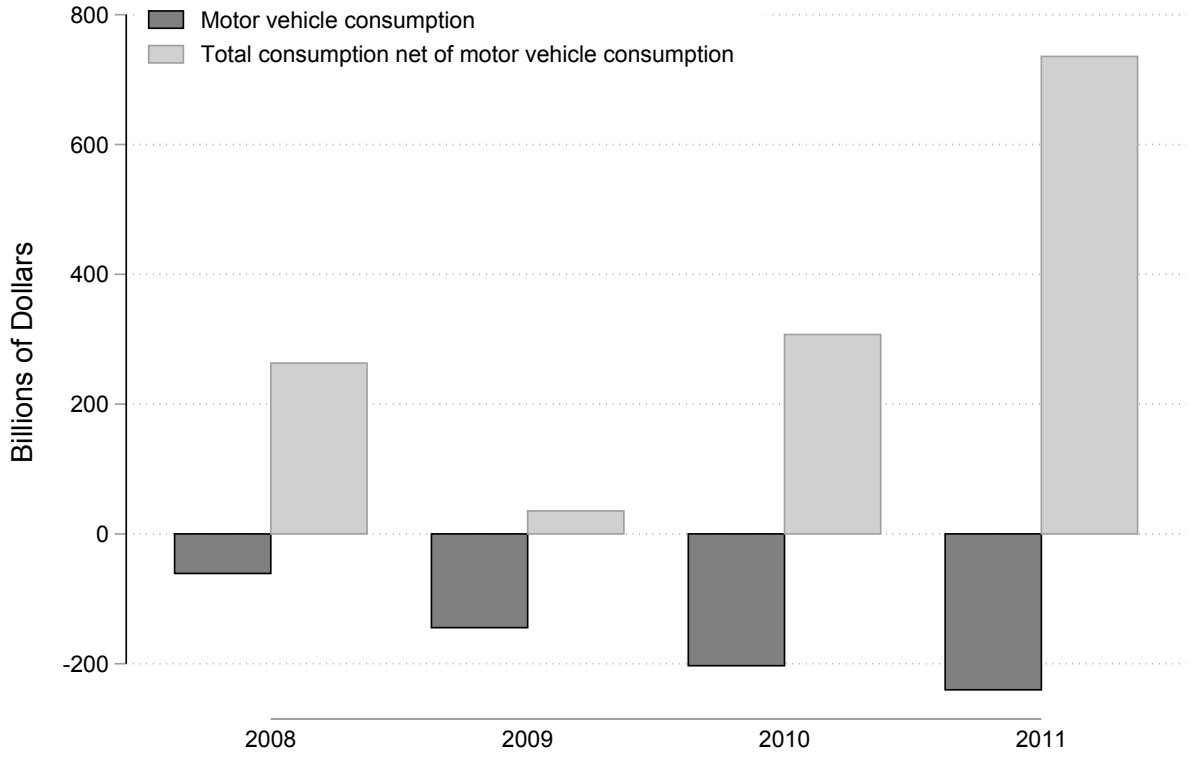
§ University of Richmond, smehkari@richmond.edu.

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

1 Introdution

The decline in autos purchased played a large role in the consumption decline during the last recession. Figure 1 plots the accumulated change in 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 vehicle sales decline was intense and violent. In one 12 month period alone, sales fell by $107 billion.[2] By Spring of 2009, Chrysler and General Motors faced bankruptcy. This led the U.S. government to use TARP funds to bailout both. At one point, the federal government owned 61

[1] Our usage of the phrase motor vehicle consumption here follows 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.

percent of General Motors.[3]

Despite the bailout, the decline in vehicle sales had a devastating impact. Over a two-year period, employment in the motor vehicle industry fell over 45 percent, excluding 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."

This paper explores two potential explanations for the auto market collapse: falling home values and changes in income expectations. We begin with the housing market. One view holds that, as homeowners see house prices fall, they internalize this as a reduction in wealth and respond by cutting auto purchases. This effect might be stronger if homeowners use home equity to purchase cars. With falling house prices, homeowners become more borrowing constrained which only intensifies the fall in auto sales.

In the first part of this paper, we exploit variation in home value changes to assess the role of home prices in explaining the auto sales collapse. We regress auto sales on home values across U.S. counties and show that a one dollar decline in home values reduced auto spending by 0.9 cents. Estimated as an elasticity, a 1 percent decline in home prices caused a 0.5 percent decline in auto sales. This relatively weak response helps explain our second finding: falling home values explain only about 19 percent of the auto sales reduction during the period.[4] In the historic auto market collapse, declining home values played a small part.[5]

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. Aggregate household debt-to-income 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 their falling housing wealth.

If leverage effects were quantitatively important in the aggregate during the last 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

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

[4] Later in the paper, we compare our findings to those of MRS, who run similar regressions and report a much larger response of auto sales to home values.

[5] Our paper, like many others studying macroeconomic phenomenon, using cross-sectional regression methods 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.

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.

Finally, we examine the effect of home values on auto usage. We replace auto sales with vehicle miles travelled in our state-level regressions and show that miles travelled was 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, home price shocks did not disrupt auto usage.

Having established that house prices played a minor role in explaining the auto sales collapse, the natural question is: what caused the auto sales decline? According to the PIH, households will reduce current consumption when expected future income falls, even in 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 astructural cross-sectional regressions.

We then develop an alternative explanation for the auto market's collapse: falling future income expectations.[6]

The durability of autos together with the discrete nature with which individuals adjust their auto stocks may be important. During the last 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. As noted earlier, aggregate vehicle miles travelled changed very little despite the large and persistent drop in auto sales starting in 2008.

With this in mind, we build a model with nondurable consumption, savings and infrequent auto purchases. In the model, individuals are subject to transitory idiosyncratic level income shocks and persistent aggregate income growth rate shocks. The latter shocks are calibrated to drive moderate swings in expected permanent income. Individuals optimally respond to negative shocks of this kind by delaying auto replacement. The model matches both the distribution of liquid wealth in the U.S. as well as the cross sectional distribution of autos by vintages in the period prior to the 2007-2009 recession.

De Nardi, French and Benson (2012) estimate the decline in permanent income expectations during the early part of the last recession. They estimate an 11 percent decline towards the low end of their bounds. Using this 11 percent decline, our model generates a persistent 50 percent decline in new vehicle sales.

Our paper relates to several lines of research. McCully, Pence and Vine (2015) report that very few households in the U.S. purchase cars with home equity lines or credit or proceeds from cash-out refinancing. Auto buyers that do use these sources are affluent and have ample access to

[6]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 nondurable consumption to perform their calculations.

credit. Along the same lines, it is somewhat incongruous to think about borrowing constraints as important for an acquiring autos, since a purchased auto is itself collateral for its corresponding auto loan.

Other papers link the home price decline during the last recession to the drop in consumer spending in the U.S. These include Mian, Rao and Sufi (2013), which finds a strong positive relationship between house prices and both durable and nondurable consumption during the period. Based on county-level data, MRS report that consumption increases by 5.4 cents from a one dollar increase in housing wealth. They find that 43 percent of this increase (2.3 cents per dollar) comes from new auto spending. We explore potential reasons for differences between our findings and MRS later in the paper. Kaplan, Mitman and Violante (2016) find a positive relationship between house prices and nondurable consumption in the cross section during the last recession. General equilibrium analyses regarding consumption and the housing market include Garriga and Hedlund (2017) and Kaplan, Mitman and Violante (2017).

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 of auto purchases amplify and propagate shocks. In response to negative shocks, individuals who were going to purchase the durable good postpone their purchases. Empirical work on autos and the permanent income hypothesis include Adda and Cooper (2000), Bernanke (1984) and Eberly (1994).

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. Another explanation focuses on the mismatch between the increased demand for higher efficiency cars, in light of positive oil price shocks, with the lack of supply of efficient vehicles by some major auto manufacturers.

Section 2 presents our county-level findings from the 2007-2009 Recession. Section 3 presents our state-level 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 4 presents a dynamic permanent income model with augmented with auto purchases in which declines in expected permanent income generate large decline in aggregate autos purchased. The final section recaps.

2 County-Level Analysis

2.1 Data and Econometric Model

Let $A_{i,t}$ denote the dollar value of new vehicles sold in county i in quarter t. We calculate auto counts from county-level auto registrations. 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 $26,400 in 2007 according to the Bureau of Transportation Statistics.[7] Values are expressed at an annual rate in thousands of dollars per household.

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 of 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. Again, values are expressed in thousands of dollars per household.

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. Let $\bar{p}_{i,t,\delta}$ and $\bar{p}^c_{i,t,\delta}$ denote the nation-wide 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 shocks. Cumulative responses give the change in auto sales accumulated over a specific horizon with respect to the accumulated change in home prices over the same horizon.[8]

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 Jorda (2005) the local projections approach.

Here, $X_{i,t}$ consist of a linear trend, seasonal dummies and a "cash for clunkers" dummy, which equals one 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

[7] Average new car auto price changed very little during the period considered.

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

response to a 1 percent increase in house prices (cumulative through horizon δ). We shall 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 1 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 equal 0.60 (SE=0.03).

Table 1: Cumulative sales elasticities to home values changes, (county-level, 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	-0.004*	-0.018***	-0.038***
growth rate)	(0.002)	(0.004)	(0.007)
House Price (lag	0.020***	0.042***	0.064***
growth rate)	(0.006)	(0.011)	(0.018)
Cash for Clunker FE	0.343***	0.600***	0.576***
	(0.037)	(0.073)	(0.112)
Quarter	0.025***	0.262***	0.558***
	(0.003)	(0.007)	(0.011)
R2	0.39	0.58	0.66
N	14916	14916	14916

Notes: The dependent variable is the cumulative percentage change in 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 = home price.

Interestingly, the cumulative response of auto sales falls rather than increases in response to an accumulated change in home prices. In a standard adjustment cost model, if changes the growth rate of purchases of a good lead to additional convex costs, this would lead to a gradual increasing cumulative response to a positive wealth shock. On the other hand, the decreasing cumulative

response seen in Table 1 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 prices 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 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 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 value of one-year's auto purchases.

This implies that, at the 3-year horizon, each dollar of additional housing wealth increase 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 a *marginal propensity to consume*, directly.

2.2 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 one in the 2009Q3-2010Q1. We also include one lag of the change in auto sales and house vales at $t-1$ (i.e., $\Delta V_{i,t-1}$ and $\Delta A_{i,t-1}$). 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 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 housing values over the same δ quarter horizon.

This term more precisely describes what we actually can measure given the data available than some other language used in existing research, such as a marginal propensity to consume (MPC). The language MPC is not suitable in the current context. First, vehicles are durable goods and individuals consume the service flow from their stock of durables rather than consume their investment in durables. Second, while our data tells 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.

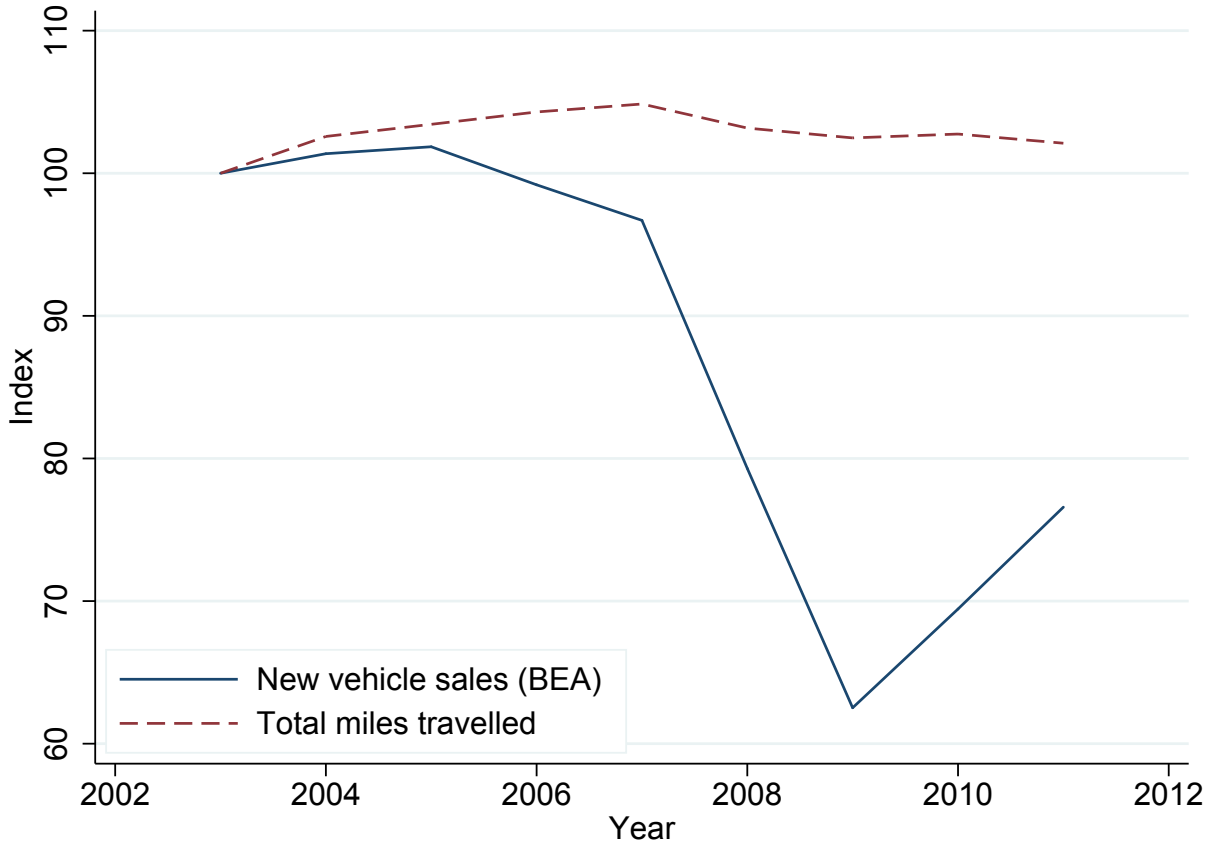
To this point, Figure 2 plots indices calculated from the number of new autos sold along with the total vehicle miles travelled during the period. While new auto sales falls dramatically, total vehicle miles travelled changes very little. The flow of services associated with autos was been nearly unchanged, suggesting that households were largely able to smooth consumption of auto usage.

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 occurred when a vehicle is acquired by a county's resident. 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 2 presents the 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 declining in the horizon. We focus particular attention on the 3-year horizon, since a related paper (MRS 2013) examines three year changes throughout. The 3-year AQR equals 0.009 (SE=0.003). This means that a one dollar increase in home values is associated with a 0.9 cent increase in auto sales. Reassuredly, the estimate of the AQR (i.e., 0.9 cents) is similar to the value approximated using the elasticity estimate of the last section (i.e., 1.2 cents).

Again, the response of new vehicle acquisitions is declining in the horizon over which the model is estimated, which suggests a dynamic aspect to the demand for vehicle in response to changes in home values.

Figure 2: Auto sales and vehicle miles travelled



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

Table 2: Cumulative vehicle acquisition rates, county-level panel

	(1)	(2)	(3)
	Coef./SE	Coef./SE	Coef./SE
1-yr Cum Home Val	0.015**	-	-
Change	(0.006)		
2-yr Cum Home Val	-	0.011***	-
Change		(0.003)	
3-yr Cum Home Val	-	-	0.009***
Change			(0.003)
Vehicles sold (lag	0.011	0.045	0.114**
change)	(0.018)	(0.034)	(0.056)
House value (lag	0.012	0.046	0.116**
change)	(0.018)	(0.034)	(0.056)
Cash for Clunker FE	0.380***	0.474***	0.341
	(0.064)	(0.161)	(0.285)
Quarter	0.028**	0.225***	0.470***
	(0.012)	(0.026)	(0.044)
R2	0.13	0.19	0.21
N	14916	14916	14916

Notes: The dependent variable 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 includes seasonal fixed effects (not reported). Standard errors are robust with respect to heteroskedacity and autocorrelation.

2.3 A Cross-Sectional Specification

In a well-known paper, MRS (2013) also estimate the response of new vehicle sales to home value changes. They use a cross-section rather than panel analysis and use changes in home values rather cumulative changes. In their baseline specification, they estimate a coefficient–analogous to our AQR–equal to 2.3 cents. Thus, their estimate is over 200 percent larger than ours.

To compare our findings with MRS (2013), we modify our specification to: (a) use a cross-section, (b) study changes in auto sales and home values, and (c) strip out some of the control variables used above.

Our dependent variable is the change in the dollar value of auto acquisitions between the first half of 2007 and the first half of 2009 in county j.[9] We choose 2009 as the end year because it follows the collapse of vehicle sales that began in September of 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 also it 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. 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.[10]

Table 3 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 0.010 (SE = 0.002). That is, an increase in housing value of one dollar in a county is associated with a one cent increase in new vehicles acquired 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 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 -1.33 (SE = 0.07). The weighted average of the dependent variable is -1.65. This implies that 81 percent (= -1.33 / -1.65) of the typical 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.

[9] The average new auto price fluctuated very little between 2006 and 2009.

[10] MRS (2013) do not include lagged variables in their specifications.

Table 3: Vehicle acquisition responses (AQR) of new vehicle sales to change in home values

	(1) Coef./SE	(2) Coef./SE	(3) Coef./SE	(4) Coef./SE
Home value change 07H1-09H1	0.010*** (0.002)	0.011*** (0.002)	0.010*** (0.002)	0.009*** (0.002)
Home value change 06H1-07H1	-	-0.007 (0.005)	-0.008 (0.005)	-0.008* (0.005)
Income pc (2006)	-	-	-0.005* (0.003)	-
Nonbank finance loan share	-	-	-	-1.914*** (0.354)
Intercept	-1.331*** (0.070)	-1.299*** (0.075)	-1.009*** (0.158)	-0.553*** (0.135)
Frac. explained by home value declines	0.191	0.217	0.202	0.169
R2	0.117	0.121	0.127	0.161
N	1243	1243	1242	1243

Notes: The dependent variable is the change in auto sales (annualized, 000s 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.

To a great extent, auto sales fell because the average household in most counties was cutting back on auto purchases, and not because of decline sales of the average household solely in counties that experienced dramatic house price declines.

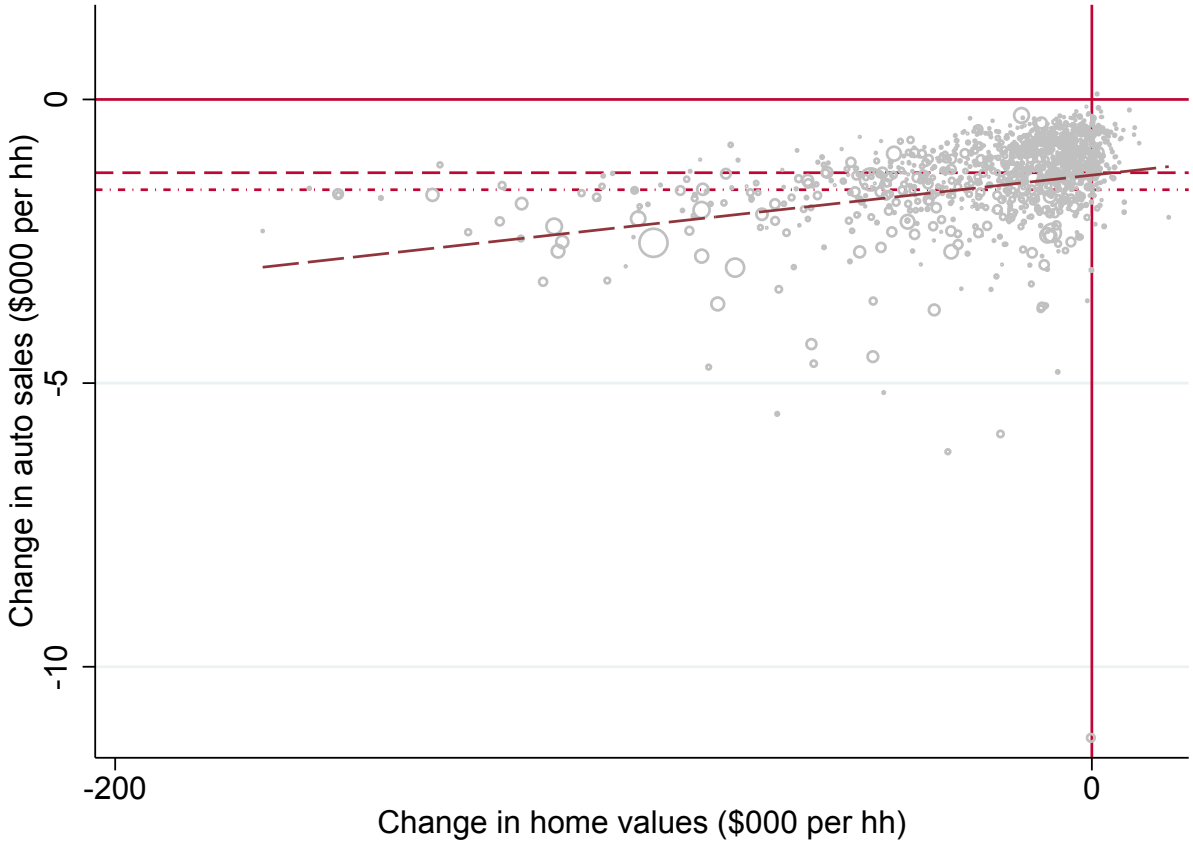
One can also see the limited role of housing in explaining the auto market collapse by applying the following counterfactual to our regression results. Take the vector of observations of house value changes in the sample and change every negative values to equal zero. Next, compute the fitted values 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 3 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 3 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 3: Response of new vehicle sales to change in home values



Notes: The long-dash line is the best fit from a weighted regression of changes in auto sales on changes in home values. Circle sizes are proportional to the number of households in each county. The short-dash line corresponds to the regression intercept, i.e. the best linear predictor of auto sales in a county that saw no change in home values. The dash-dotted line is the unconditional weighted average of 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 3 adds the lagged change in home prices as a control, which brings us closer

to the panel specification used earlier.[11] 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.[12] 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 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.

2.4 Reconciling Our Findings with MRS (2013)

As stated in the introduction, MRS (2013) find a strong positive relationship between house prices and both durable and nondurable consumption during the period. Based on county-level data, MRS report that consumption increases by 5.4 cents from a one dollar increase in housing wealth. They find that 43 percent of this increase (2.3 cents per dollar) comes from new auto spending. The differences between their and our findings may be the scaling they used to map quantities of autos sold to the values of those autos sold.

For our sample, the weighted average of the dependent variable is -1.79, or $-\$1,790$.[13] This is substantially smaller than a similar variable reported in MRS, which equals $-\$3,300$, between 2006 and 2009. We contend that the MRS figure is likely too high.

To show this, we offer the following calculation. First, multiply MRS's weighted average by the number of households in the U.S. This implies a fall in sales of roughly $383 billion between 2006 and 2009.[14] The BEA reports that total auto sales fell by 6.5 million units over this period. Based on the MRS total value number and the BEA sales count, one would infer an average vehicle price equal to $\$58,900$. This is more than double the average car price in the U.S. during this period.

Column (5) estimates the identical specification as (4) except we use the higher value for the average new car price to compute the dependent variable. In this case, the AQR increases to X.X, which is very similar to the value reported in MRS (2013).

In a potentially related data issue, MRS use aggregate Census Annual Retail Sales Data to

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

[12] 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.

[13] Our sample is limited primarily by house price data availability as explained below.

[14] We compute $383 billion as $\$3300 \times 116$ million.

assign a dollar value to auto sales based in each county according to that county's share of new car sales. The particular category they use, as best we can tell, is "New Car Dealers" retail sales.

However, this amount includes revenue from retailing new vehicles "in combination with activities, such as repair services, retailing used cars and selling replacement parts and accessories." Unless one could strip out the value of these other activities, one would overstate the value of new cars sold in a particular county using the MRS approach. If, for example, a new car dealership sold a new car for $\$30,000$ and took a trade in that it was able to resell for $\$25,000$, then dealership would record $\$55,000$ in sales to the Census.[15]

An alternative approach would be to consider both new and used vehicles in constructing the dependent variables. This would generate additional problems. First, the vehicle counts are based on registration data. An auto handed down from a mother to a son in which the registration changed would be counted as a sale. Similarly, someone moving a car's registration from one state to another would be counted as a sale, without having any offsetting reduction from the place where the person relocated. More generally, there would remain an implicit double-or-more counting as once-new cars were sold as used cars and those used cars were sold again as used cars. Also, spending on used cars is not reflected in GDP. While potentially important for some questions, the reshuffling of used vehicles amongst households does not directly impact the quantity of newly produced goods and services in the economy.

How one translates shares of new autos sold into dollar values of new autos sold matters crucially for the AQR, but not for computing the percentage contribution of declining home values towards declining auto sales. Scaling up or down the left-hand side variable by a fixed proportion changes each coefficient in the regression as well as average vehicle sales by the same factor. Since it is the intercept coefficient relative to the mean of the dependent variable that determines the aggregate importance of house prices towards auto sales, the scaling factor cancels out in the numerator and denominator. Thus the second main finding our paper—the general inability of house prices to explain the auto market collapse—is unrelated to the auto count scaling issue.

On another matter, one variable that we do not consider is a measure of total net worth. It would be interesting to examine the influence of total net worth (housing and non-housing) on auto sales. However, estimating net worth at the county-level requires financial wealth information, which requires some imputation. MRS impute financial wealth by calculating the share of each counties' dividend plus interest earnings relative to national dividend plus earnings. They then divide the aggregate financial wealth from the Flow of Funds according to those shares. Following this procedure, we found that the weighted average of financial wealth equals over $300,000 per household in 2007. By contrast, the median financial wealth from the Survey of Consumer Finances in that year equaled $\$120,000$.[16] As such, we interpret this imputation method as providing

[15]The questionnaire sent by the Census to dealers states that, in filling out their surveys, the value of trade-ins should be included as partial payment.

[16]Bucks et. al. (2009).

unreliable results and do not study net worth.

3 State-Level Panel Analysis

3.1 Data and Econometric Model

In this section, we compare the relationship between home prices and vehicle sales during the last recession period relative to the remainder of the last 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.[17] This variable is available at the state level at an annual rate beginning in 1997.[18]

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 this into a 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 and 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 home 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 home 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 horizons regressions).

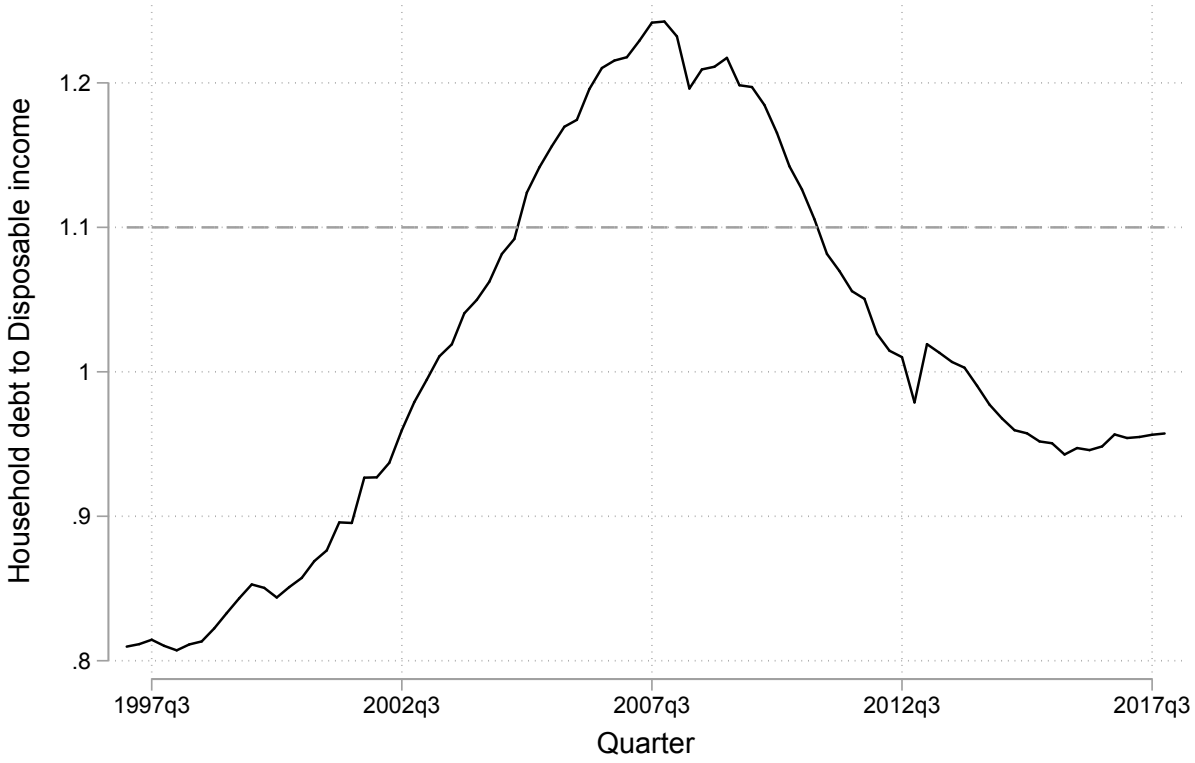
3.2 Results

Columns (1) through (3) of Table 4 contain the elasticity estimates for the full sample. All three cumulative elasticities are positive and statistically different from zero. At the 1-horizon, the

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

[18] Personal consumption expenditures of motor vehicles includes 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 4: Ratio of household debt to disposable personal income



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

coefficient equal 0.92 (SE=0.10). A one percent increase in home prices over one year leads to a 0.92 percent increase in motor vehicle consumption over the corresponding 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 2. For example, at the 1-year horizon, the benchmark county-level estimate equals 1.08 (SE = 0.06).

Recall that in Section 2, 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 the county-level data, this implies that over the entire 1997-2017 period, the response of motor vehicle consumption to home 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 MRS (2012), which finds 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 for how individuals adjust auto sales. Highly leveraged individuals may choose or be forced to react quickly in adjusting 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-leverage households.

3.3 Miles Traveled and Home 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 travelled provides a more direct measure of vehicle services provided. Therefore, we next estimate the relationship between the growth in miles travelled and home 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

Table 4: 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 one year lag of the one year growth rate of auto sales and house prices as controls. Lev. = leverage.

level using the weighted average the number of households in each county.

We then run a regression where the dependent variable is the analogous change in the log of miles travelled 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 5. Columns (2), (4), and (6) add additional controls. These are 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 travelled accumulated over the corresponding 3 years.

Thus, the effect of house prices on vehicle miles travelled is very small. By comparison, the cumulative elasticity of vehicle sales to house price changes over a 3-year horizon, from Table 1, equaled 0.60. The elasticity of house prices on vehicle sales, as we already explained was a modest effect, is more than fifteen times as large as the effect of home values on miles travelled.

At least aggregated to the state level, there is little evidence that house prices changes disrupted the service flow of provided by vehicles to a significant extent. Thus, households were likely effective at smoothing the effects of home price shocks on their vehicle usage during the period. The economic model developed and calibrated in the next section is motivated by this observation that one can see a large change in investment in durable goods purchases is consistent with only a small change in the flow of services delivered by the stock of a durable good.

4 Consumption and Auto Purchases

4.1 The Idea and the Mechanism

This section develops a permanent income model augmented with an auto-purchase choice. In the model, at multiple points over his life, an individual pays a fixed cost 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 economy-wide 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.[19] In the aggregate, there is

[19]See Leahy and Zeira (2005) for a theoretical exploration of this mechanism.

Table 5: Cumulative response elasticities of miles traveled to home 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 = home price.

a large downward response of vehicle sales to a negative shock to expected permanent income.

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 income growth rate to zero, similar to that experienced during the last recession, delivers a 40 percent auto sales decline on impact. The 40 percent decline in auto sales is roughly equal to that experienced in the second half of 2008. In contrast, a model that simply treated auto purchases as part of nondurable consumption would have a elasticity that would be much too low to match the observed 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 home 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 home price changes have no direct effect on own wealth and therefore auto purchases. The effect 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 result—the quite limited role for house prices in explaining the new auto sales decline—was established using cross sectional 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 survey data on future income expectations in calibrating our model.

4.2 The Model

Our model consists of a unit mass of individuals indexed $i \in [0, 1]$. Each individual i earns an exogenous stochastic income and maximizes lifetime utility by choosing a stream of savings, nondurable consumption, and vehicle purchases. We calibrate the model so that a period lasts one year. The individual buys only newly produced, i.e. not pre-owned, cars.[20]

Let income be given by $\tilde{Y}_{i,t} = \exp(y_{i,t}) Z_t$, where Z_t indexes aggregate income and evolves

[20]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.

according to $Z_t/Z_{t-1} = 1 + g_t$. Also, g_t evolves according to a two-state Markov chain 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 standard deviation σ_ε and is i.i.d over time and individuals. Let $y_{i,-1}$ and Z_{-1} be positive and given as initial conditions.

The expected utility function is

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

where $\tilde{C}_{i,t}$ is consumption and $v_{i,t}$ is the vintage of the auto currently owned by the individual, and U_N and U_D give the utility of nondurable and durable goods respectively. We assume,

$$U_N (C) = \frac{C^{1-\sigma}}{1-\sigma} \tag{3}$$

Utility, therefore, is increasing and concave in nondurable consumption.

We further assume that each individual owns exactly one car, and the utility of owning a car depends on how old the car is:

$$U_D (v) = \begin{cases} -\alpha \left[e^{\chi v} - e^{\chi \bar{v}} \right] & \text{if } v \leq v_c \\ -\alpha \left[e^{\chi v_c + \chi \xi (v - v_c)} - e^{\chi \bar{v}} \right] & \text{if } v > v_c \end{cases} \tag{4}$$

Thus, the utility an individual gets from owning a car is decreasing and convex in the vintage of the vehicle. Vintages take interger values from 0 to J, with 0 indicating the individual has a brand new car, and $\bar{v}$ gives the average vintage of the car. If individual i buys a new car at t, the vintage resets to $v_{i,t+1} = 0$. Otherwise, $v_{i,t+1} = v_{i,t} + 1$. An auto purchased J periods ago becomes inoperable and must be replaced. Furthermore, the elasticity of the disutility is χ if the vintage is younger than v_c and decreases to $\chi\xi$ if the vintage is older. We use a step function for the disutility function above to reflect the fact that vehicles depreciate in value at a faster rate in their early age.

Our analysis studies responses to shocks, added up across a large number of individuals, after the economy has reached a steady-state distribution. As such, the starting values $(y_{i,-1}, Z_{-1}, v_{0,i})$ will not play a role in our quantitative results. In the model's stochastic steady state, the vintages held by individuals are heterogeneous.

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

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

where borrowing is not allowed (i.e., $\tilde{W}_{i,t} \geq 0$) and the price of a new auto is given as $\tilde{F}_t$ in each

period.

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: $\tilde{F}_t = Z_t F$. Without rising auto prices, as income trended 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 car prices (which rise in expected price over time).

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

$$S\left(W, v, y, g\right) = \max\{S^R\left(W, v, y, g\right), S^N\left(W, v, y, g\right)\}$$

where S^R and S^N denote the values associated with car replacement and car retention, respectively.

$$S^R\left(W, v, y, g\right) = \max_{C, W'} \left\{\log\left(C\right) - \alpha\left[1 - e^{\chi \bar{v}}\right] + E\left[S\left(W', 0, y', g'\right)|y, g\right]\right\}$$

$$S^N\left(W, v, y, g\right) = \begin{cases} \max_{C,W',v'} \left\{\frac{C^{1-\sigma}}{1-\sigma} - \alpha\left[e^{\chi v'} - e^{\chi \bar{v}}\right] + E\left[S\left(W', v', y', g'\right)|y, g\right]\right\} & \text{if } v \leq v_c \\ \max_{C,W',v'} \left\{\frac{C^{1-\sigma}}{1-\sigma} - \alpha\left[e^{\chi v_c + \chi\xi(v' - v_c)} - e^{\chi \bar{v}}\right] + E\left[S\left(W', v', y', g'\right)|y, g\right]\right\} & \text{if } v > v_c \end{cases}$$

subject to:

$$C = \left(1 + r\right) W + \left(1 + g\right)\left(y - W'\right) - \left(1 + g\right) F \times \mathbf{1}\left(v' = 0\right)$$

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

with $v_i' \leq J$ and $W_i' \geq 0$. The evolution for each the exogenous stochastic variable is given by a two-state markov chain in the aggregate variable g and (2). The prime superscript advances time by one period.

Let $\hat{W} = \hat{W}(v; y, g)$ denote the level of W that leaves an individual with income level y indifferent between car replacement and car retainment. In particular, an individual with wealth at or below $\hat{W}$ will not replace his car. Thus, individuals described by their state variables (W, v, y) will begin the subsequent period with car vintage given by

$$v' = \begin{cases} 0 & \text{if } W \leq \hat{W} \\ v + 1 & \text{if } W \geq \hat{W} \end{cases} \tag{5}$$

Given (5), we can explicitly define the evolution of the distribution of individuals over car vintages. At each date, the fraction of individuals associated with each car vinatge is given by

the predetermined vector $\Theta_t = \{\theta_{v,y}\}$, where each $\theta_{v,y}$ describes the number of individuals with income level being y currently owning vintage v cars, and $\theta_{0,y}$ denotes the number of individuals with income y that replace their cars in the previous period. The evolution of the cross-sectional distribution over time is determined as follows. Let α_{vy} denote the vector of adjustment rates, then the fraction of individuals associated with each car vintage in the support at Θ is summarized by equations (6) and (7) below.

$$\theta'_{0,y} = \sum_{v=0}^{J} \alpha_{v,y}\theta_{v,y} \tag{6}$$

$$\theta'_{v,y} = \theta_{v-1,y}(1-\alpha_{v-1,y}), v = 1, 2, ..., J \tag{7}$$

Finally, economy wide aggregate variables are calculated as

$$\text{Aggregate Variable} = \int_0^1 \text{Individual Variable}_i di$$

4.3 Calibration

We start by assuming $\beta = 0.94$ and $\sigma = 1$ for log utility, which are standard assumptions.

Next, we assume $g \in (0, \bar{g})$ with transition matrix

$$\Pi_g = \begin{bmatrix} p & 1-p \\ 0 & 1 \end{bmatrix}$$

where $\bar{g} > 0$.

This form allows us to highlight a few features of the 2007-2009 recession. Suppose the individuals in an economy are currently in and are expected to stay in the high income growth state, which we assume is 2 percent per annum. We can solve the problem for a large number of individuals and compute the steady-state distribution of wealth and auto vintages. The expected present value of income (with current income normalized to 1) is then given by:

$$EPV_H = \frac{1+r}{r-\bar{g}}$$

Then, suppose that the economy unexpectedly enters into a temporary "growth slow down" putting the economy into the zero-growth state temporarily. This low growth aggregate state continues with probability p each period. In the period of the shock,

$$EPV_L = \frac{1+r}{1+r-p} + \frac{1-p}{1+r-p} EPV_H$$

De Nardi, French and Benson (2012) examine individual-level survey data on income in order to compute the decline in expected permanent income at the time of the last recession. The survey

data they look at only asks individuals about income expectations five years out. As such they need to make assumptions about further out income expectations. Using their most conservative estimate, expected permanent income fell by 11 percent. To hit that target we choose r and p such that $\log\left(EPV_L/EPV_H\right) = 0.88$.

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 nondurable consumption, the real interest rate consistent with steady-state growth would be: $1 + r^* = \left(1 + \bar{g}\right)/\beta$.

$$1 + r^* = \left(1 + \bar{g}\right)/\beta$$

In our model, this would lead to oversavings 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.

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

Experimenting with κ, we set $\kappa = 0.005$ Given the values of $\bar{g}$ and β, this implies $r = 0.08$. With r and $\bar{g}$ determined, we match an 11 percent decline in expected permanent income by assuming $p = 0.8$.

Next, we calibrate the vehicle cost F. The Kelly Blue Book price of a new vehicle in 2008 was approximately \$27,000. Although outside our model, it makes sense to account for a trade-in value of the car being scrapped. If this equals 1/3 of a new car price, then the net cost of a new vehicle equals \$18,000. Using a household income of \$60,000, we set $F = 0.3$.

We follow standard assumptions on the individual income process. Stated at an annual rate, $\rho = 0.95$ and $\sigma_\varepsilon^2 = 0.015$. We discretize the process for $y_{i,t}$ on a grid with N points, where $N = 8$, using the Rouwenhorst method. We discretize the wealth state space into H evenly spaced points between 0 and $\bar{W}$, in which $H = 500$ and $\bar{W} = 40$. We set $J = 30$.

The only remaining parameters are those associated with the disutility of having an older car. These are α, χ, ξ, v_c, and $\bar{v}$, These parameters largely influence the average age of a vehicle (originally purchased new) in operation and the shape of the corresponding vintage density. Data on the later can be gleaned from the Survey of Consumer Finance. We choose these parameters to match this distribution. Figure 6 plots a kernel density estimate of this distribution in 2007. The parameter values are summarized in Table 6.

4.4 The Solution Method and Policy Function

We solve the individual's problem by discrete discounted dynamic programming. The total state space has dimension 198,400 $(= H \times (2 \times N) \times (J + 1))$.

Figure 5 plots the optimal cutoff (log) wealth as a function of the individual's current (log) income and vintage of his current auto. The cutoff wealths appears as numbers on the chart,

Table 6: Parameter Values

Parameter	Value	Description	Motivation
β	0.94	Discount Factor	Standard value for annual model
σ	1	Elasticity of Substitution	Log utility
g	0.02	Steady state (high) income growth rate	Standard value
r	0.08	Interest rate	Consistent with g and β values
p	0.8	Prob. of remaining low income growth	Chosen to match income growth expectations
ρ	0.95	Persistence of income level	Standard value
σ_ε	0.015	Standard deviation of shocks to income	Standard value
F	0.3	Price of car relative to income	New car price (after trade in) = 30% of avg. income.
α	0.03	Relative disutility from car vintage	
χ	0.15	Elasticity of disutility for $v \leq v_c$	
$\chi\xi$	0.022	Elasticity of disutility for $v > v_c$	Jointly chosen match to the car vintage distribution
v_c	8	Elascitity change cutoff	
$\bar{v}$	10	Average car vintage	

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 his auto if current log wealth is above 2.5 and does not replace if the current log wealth is below 2.5%.[21] 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 Results from the Economic Model

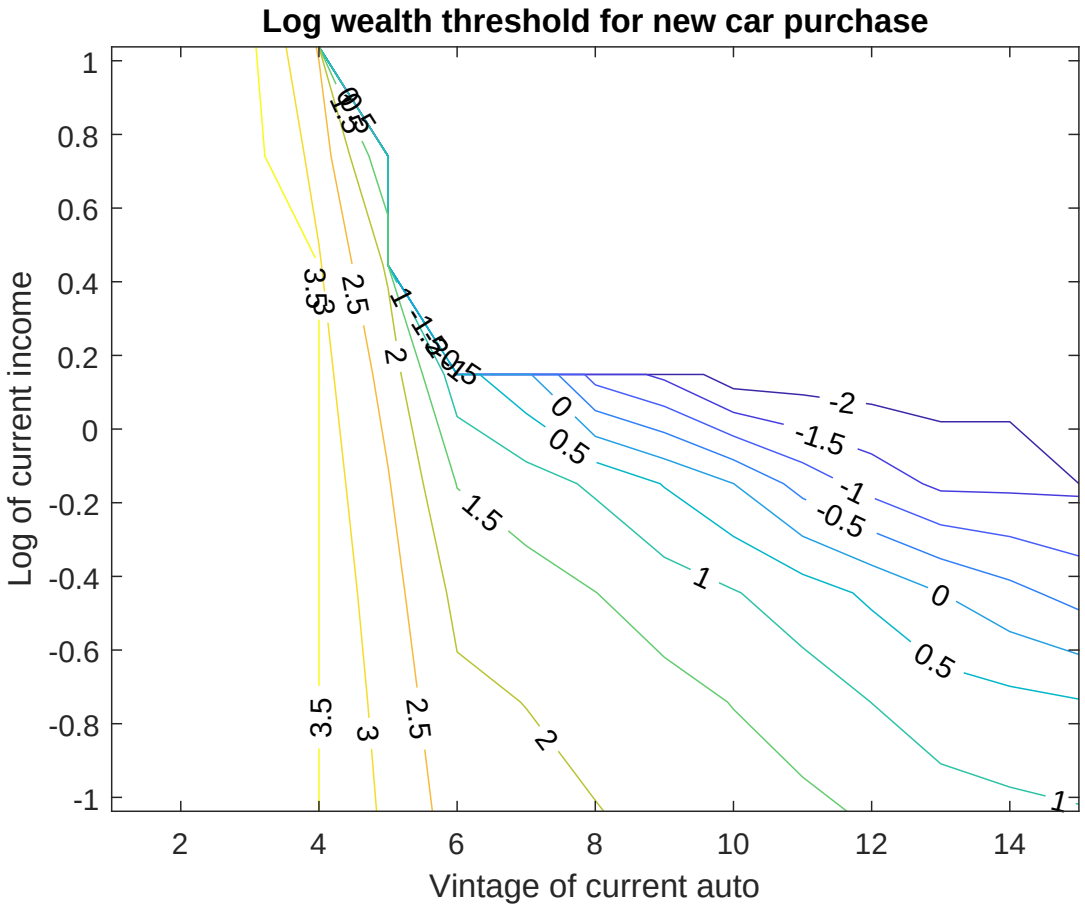
Using the policy function solved above, we simulate the outcome of an individual's decision problem for a long history of $T + Q$ periods. For the first T periods, assume Z_t grows at g percent and is expected to grow at 2 percent in all future periods. Figure 6 plots the steady state distributions of auto vintages from our economic model and data. This figure shows that our model generates an auto vintages distribution that can approximately match the distribution from data.

At period $T + 1$, average income growth slows to 0 percent. Beginning at $T + 1$, the forecasted law of motion for Z_t evolves according to Π_g. That is, average income begins each period in the low growth rate state after which it escapes to the absorbing high growth state with probability $1 - p$. In each period, the idiosyncratic determinant of income, $m_{i,t}$, is drawn according to the above process.

We set $T = 500$ and $Q = 5$. The realized path of net Z_t growth shall remain at zero from

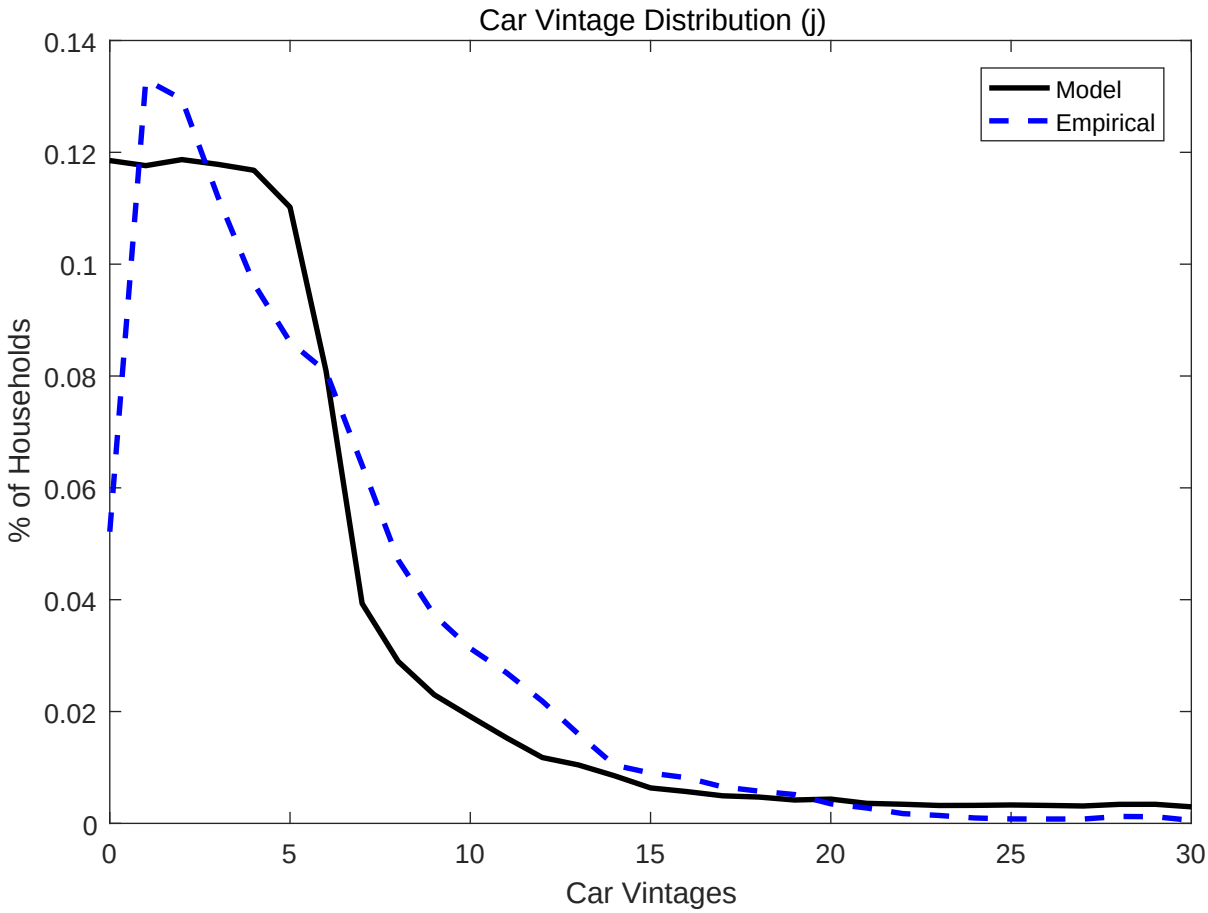
[21] Vintages take on integer values, so the contour lines between integers on the x axis reflect interpolations.

Figure 5: 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 labelling the corresponding contour line.

Figure 6: Distribution of auto vintages, data and economic model

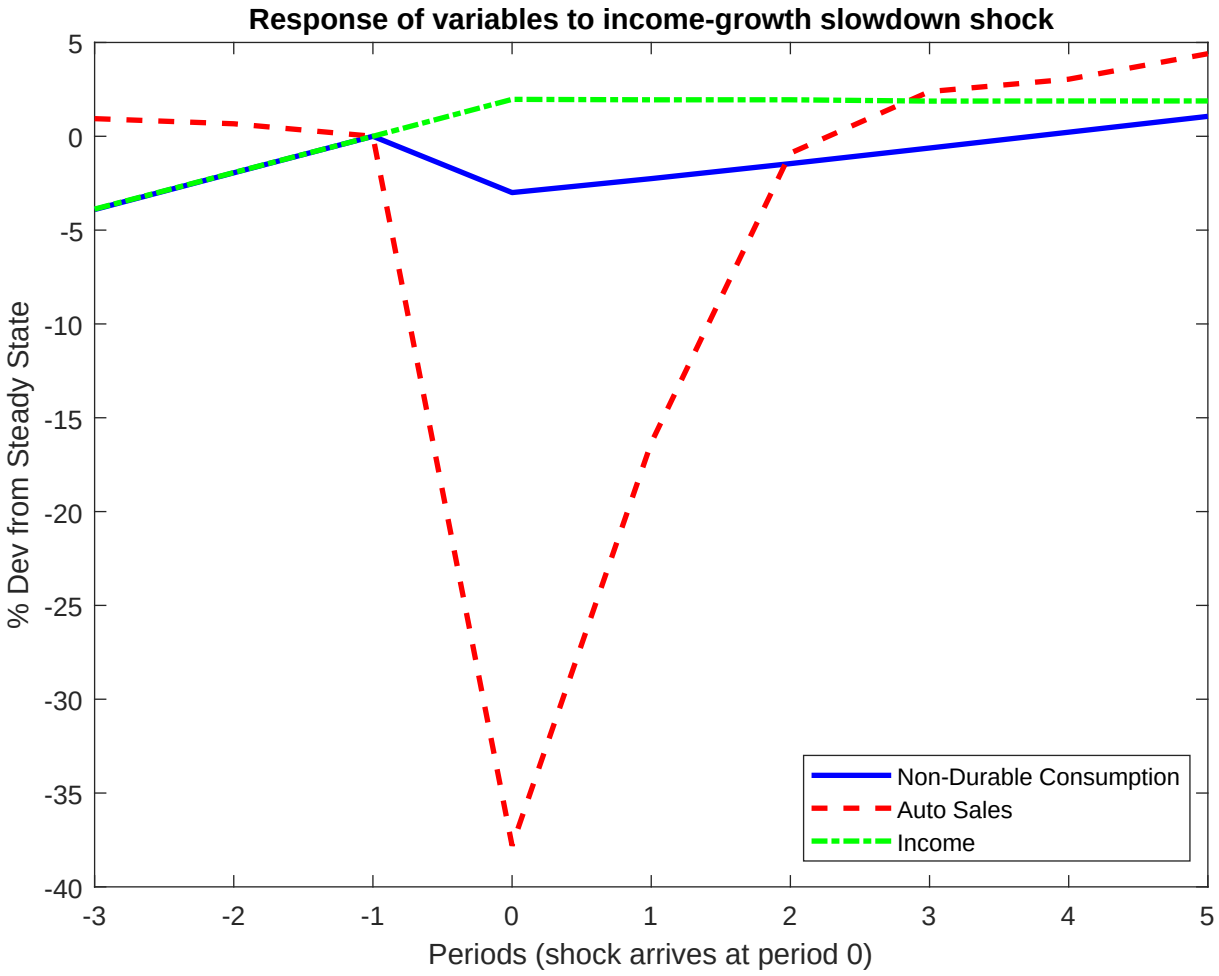


Notes: The dashed line plots the 2006 distribution of vehicles, originally purchased new, and the solid line plots the steady-state distribution from the economic model preceding the aggregate income growth rate shock.

$T+1$ through $T+Q$. Finally, we repeat this simulation for a large number of individuals H, where $H = 5000$. We experimented with T and H in order to make T sufficiently large that such that initial conditions on individuals' state variables do not affect the long-run, pre-shock distribution. We set H sufficiently large so that the idiosyncratic paths $m_{I,t}$ largely cancel out across individuals.

Figure 7 plots the impulse response for average income, auto sales and nondurable 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 2 percent annually. Income is exogenously growing. Consumption is a constant fraction of income along the steady state. The quantity of vehicles sold is constant, however the price of vehicles rises at the same rate as income. At period zero, average income growth unexpectedly becomes zero. It remains at zero through period 4 (although individuals predict a 20 percent

Figure 7: Response of variables to income growth slowdown shock



Notes: Paths are aggregated across 10,000 individual paths from the initial steady-state wealth, income and auto vintage distributions.

chance each year that average income growth will increase to its initial steady-state growth rate.)

Nondurable consumption falls approximately 5 percent in response to the growth slowdown. The decline in auto sales is much more dramatic. On impact, auto sales falls by nearly 40 percent. Also observe that the auto sales decline is persistent. Auto sales do not recover to the steady state value until the third year following the expected income shock. Because autos are a durable good, many individuals respond to the decline in expected permanent income by delaying the replacement of their existing automobile. 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 nondurable consumption.

6 Conclusion

A related explanation for the decline in auto sales is the increase in uncertainty that many researchers have associated with the last recession.[22] 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 after being acquired. Consistent with this story, Hassler (2001) finds auto expenditures in the U.K. declined dramatically with increases in uncertainty, proxied by stock market volatility.[23]

Our estimates speak to two important concepts in the economics of consumption: the permanent income hypothesis (PIH) and consumption risk sharing. Our regressions do not directly test either theory, quantitatively; however, our findings do not violate either theory in a quantitatively important way.

The PIH states that individuals consume out of permanent income (current wealth plus the expected discounted future income). If the market interest rate equals the rate of time preference, then in many models, households consume only the interest earned on their stock of permanent income. We find that the AQR is equal to 0.01, which is less than the long-run real interest rate in the post-WW U.S.[24]

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 the durable good with the flow of services of the stock of durable goods. Based on the generally smooth series for aggregate vehicle miles travelled 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 home 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 last 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 the durables as explained above. Therefore, without additional structure on preferences or else different data, 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

[22] See for example Baker, Bloom and Davis (2016) and Jurado, Ludvigson and Ng (2015).

[23] See also Bertola, Guiso and Pistaferri (2005).

[24] Of course, autos are only one component of consumption and our paper does not estimate the response of other consumption goods to housing wealth changes.

to economic shocks, then one could see utility-reducing changes on households apart from miles travelled. 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 travelled) was only mildly disrupted during the last 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.

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.

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.

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.

Garriga, C. and A. Hedlund (2017), "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession," FRB St. Louis Working Paper No. 2017-30.

Glover, A., J. Heathcote, D. Krueger and J. Ríos-Rull (2017), "Intergenerational Redistribution in the Great Recession," Working Paper.

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

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," NBER Working Paper No. w22232.

Kaplan, G., K. Mitman and G. Violante (2017), "The Housing Boom and Bust: Model Meets Evidence," NBER Working Paper No. 23694.

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," FEDS Working Paper No. 2004-63.

McCully, B., K. Pence and D. Vine (2015), "How Much are Car Purchases Driven by Home Equity Withdrawal? Evidence from Household Surveys," FEDS Working Paper No. 2015-106.

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.

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.

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