How Big is the Wealth Effect? Decomposing the Response of Consumption to House Prices*

S. Borağan Aruoba

University of Maryland

Ronel Elul

FRB Philadelphia

Şebnem Kalemli-Özcan

University of Maryland

March 2018

Abstract

We investigate the effect of declining house prices on household consumption behavior during 2006–2009. We use an individual-level data set that has detailed information on borrower characteristics, mortgages and credit risk. Proxying consumption by individual-level auto loan originations, we decompose the effect of declining house prices on consumption into three main channels: wealth effect, household financial constraints, and bank health. We find a negligible wealth effect. Tightening household-level financial constraints can explain 40-45 percent of the response of consumption to declining house prices. Deteriorating bank health leads to reduced credit supply both to households and firms. Our dataset allows us to estimate the effect of this on households as 20-25 percent of the consumption response. The remaining 35 percent is a general equilibrium effect that works via a decline in employment as a result of either lower credit supply to firms or the feedback from lower consumer demand. Our estimate of a negligible wealth effect is robust to accounting for the endogeneity of house prices and unemployment. The contribution of tightening household financial constraints goes down to 35 percent, whereas declining bank credit supply to households captures about half of the overall consumption response, once we account for endogeneity.

JEL CLASSIFICATION: E32, O16.

KEY WORDS: financial crisis, mortgage, individual-level data, general equilibrium, bank health, credit supply
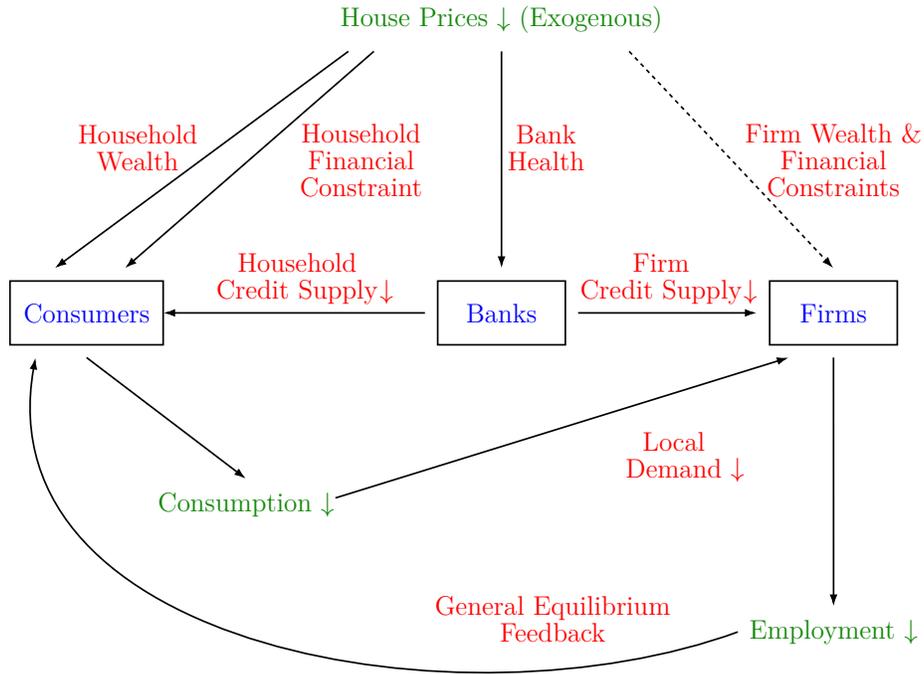
1 Introduction

The U.S. economy experienced a large financial crisis together with a housing bust in 2007–2008. A deep recession with significant declines of consumption, investment, and employment has followed. Although there is an extensive theoretical and empirical literature on the causes and consequences of the crisis, so far there is still no consensus on the role of the channels linking the housing bust to the recession.

There have been three main narratives of the crisis put forth in the literature. The first narrative is a wealth shock to consumers via a decline in their housing wealth, which lead them to cut their consumption, and this in turn lead to a decline in output. Mian and Sufi (2009), Mian and Sufi (2011), and Mian, Rao, and Sufi (2013) have been the main proponents of this view, where an increase in household leverage predict the subsequent crisis, de-leveraging and consumption decline. They show empirically a strong relationship between these variables and argue that the recession is due to this demand channel via declining consumption.[1] The second narrative is about households being financially constrained as a result of a shock to their housing wealth. When house values go down, the value of housing collateral falls and households' borrowing constraints get tighter, which in turn might prevent them from borrowing. Berger, Guerrieri, Lorenzoni, and Vavra (2015) and Kaplan, Mitman, and Violante (2016) have proposed models where this channel is important for the decline in consumption and the associated recession. Aladangady (2017) provides empirical support for this channel, where a large part of the response of consumption to changing house values are driven by credit-constrained households. The final narrative is about the shocks to the financial sector which tighten their financial constraints, and they in turn reduce credit supply to both households and firms. Households decrease consumption as a result, and firms cut down employment and investment. There is an extensive empirical debate on the effect of reduced credit supply on firm employment. While Duygan-Bump, Levkov, and Montoriol-Garriga (2015) and Greenstone, Mas, and Nguyen (2015) find that reduced credit supply can only account for less than one-tenth of the decline in employment, Chodorow-Reich (2014), Chen, Hanson, and Stein (2017) and Gilchrist, Siemer, and Zakrajek (2017) find that up to one-third of the employment decline may be driven by bank shocks.

Our goal in this paper is to quantify each of these narratives using detailed individual-level data, which include mortgage and credit risk information. We know from the existing literature that shocks to house values create a large consumption response at an *aggregate* (ZIP

---

[1]Philippon and Midrigan (2016), using aggregate data and a model, argue that household de-leveraging by itself cannot explain a large part of decline in employment and output.

Figure 1: House Prices and Consumption: Channels

House Prices ↓ (Exogenous)

Household Wealth

Household Financial Constraint

Bank Health

Firm Wealth & Financial Constraints

Household Credit Supply↓

Firm Credit Supply↓

Consumers

Banks

Firms

Consumption ↓

Local Demand ↓

General Equilibrium Feedback

Employment ↓

code or county) level, but we have only scant evidence about the channel(s) such a response operates through. The key shortcoming in most of the literature so far is the unavailability of individual-level consumption data and the inability to combine various individual-level controls in conjunction with more aggregate controls to identify these channels. It is not possible, for example, to know who is credit-constrained and where households are in their life cycle, which directly affects their housing demand, without individual-level data.

Figure 1 shows all of the possible channels that will lead to lower consumption as a result of an exogenous decline in house prices, where we show three players in the same locality: consumers, banks and firms. First, on the household side, we argue that as a result of declining house prices, there will be both a wealth effect, denoted with the arrow "household wealth" and a collateral shock effect, denoted with the arrow "household financial constraints." Although there are models that combine these effects under a single wealth effect,[2] we argue that their effects have to be quantified separately. Why this is important? In the standard permanent income model, a shock to housing wealth will have no effect on consumption since positive endowment effects will be canceled out by negative cost of living effects, as shown by Buiter (2008). In the context of the life-cycle model, if homeowners are

---

[2] See for example Kaplan, Mitman, and Violante (2016).

likely to sell their house in the future, there can be positive wealth effects via rising house prices as modeled by Sinai and Souleles (2005). In terms of the current debate, many theory papers argue that, to be able to match the large responses of consumption to house prices changes found in the data by Mian, Rao, and Sufi (2013), one needs collateralized lending that amplifies the impact of housing wealth on consumption.[3] Our individual-level data and methodology will allow us to separate these effects.

Next, as shown in the figure, there is the effect of house price declines on bank health. If banks are exposed to the real estate market, housing price declines constitute a negative balance sheet shock to banks, which results in banks cutting credit supply both to households and firms. As argued by Justiniano, Primiceri, and Tambalotti (2017) an increase in credit supply is the only force that can match the empirical regularities in the boom period. They argue that looser borrowing constraints cannot account all for the facts since they only shift the demand for credit. In their model these forces interact, a lending constraint on the bank side and a household borrowing constraint are both in play during the boom-bust phase. They argue that in the models without an exogenous credit supply decline, tightening of the household borrowing constraint put upward pressure on interest rates, which has not been observed during the boom phase. Hence, we believe it is important to quantify this effect separately than the previous ones.[4]

Lower credit supply to firms will lead to lower employment and investment. As argued above there is a debate in the empirical literature on the size of this effect. Another possible channel is, as shown by the dotted arrow, a collateral shock to firms' balance sheet if firms' owners use their own housing wealth as collateral to get loans to invest and to produce.[5] We will not be able to study this channel, since we do not have information on firms' or their owners' real estate wealth. In addition, due to low consumption, demand for firms' output will be lower, which will also lead firms to decrease employment, as shown by the "local demand" arrow following the work of Mian and Sufi as cited above. Any firm-level response via lower employment will feed back to lower consumption due to general equilibrium, as shown with the bottom arrow. We will be able to identify these effects collectively using

---

[3]See Berger, Guerrieri, Lorenzoni, and Vavra (2015), Guerrieri and Iacoviello (2017), Iacoviello (2005). More generally (outside housing), see Barro (1976), Stiglitz and Weiss (1981), Hart and Moore (1994), Kiyotaki and Moore (1997), Bernanke, Gertler, and Gilchrist (1999).

[4]Gropp, Krainer, and Laderman (2014), show empirically that renters with low risk scores, compared to homeowners in the same markets, reduced their levels of debt more in counties where house prices fell more. This suggests that the observed reductions in aggregate borrowing were more driven by cutbacks in the provision of credit than by a demand-based response to lower housing wealth.

[5]See Decker (2015) who shows in a model that this channel is important for the decline in start-up activity. See Bahaj, Foulis, and Pinter (2017) for an empirical study of this channel for U.K.

county-level employment.

Not only the literature that studies the Great Recession, but also the broad literature that tries to understand the effect of house prices and housing wealth on consumption takes by and large an aggregate approach. The early literature uses time series data from the U.S. as a whole, and the later literature uses geographic variation across states or counties. In either case, aggregate time-series and cross-sectional correlations make identification hard. For example, expectations about future income, can drive both consumption patterns and house prices. As shown by Attanasio, Blow, Hamilton, and Leicester (2009) and Calomiris, Longhofer, and Miles (2009), the strong aggregate relation between house prices and consumption shown by Case, Quigley, and Shiller (2005), Carroll and Kimball (1996), and Carroll, Otsuka, and Slacalek (2011) goes away once expectations of income and other common factors are controlled. Attanasio, Blow, Hamilton, and Leicester (2009) is an early paper that shows similar responses from renters and home owners, which again indicates the existence of common factors in aggregate data. Demyanyk, Hryshko, Luengo-Prado, and Sørensen (2015) also show that unemployment, income, and debt are important determinants of consumption in the aggregate data.

In the aggregate data, there can also be an omitted variable problem related to compositional changes in the population, such as the effect of age on housing demand. Both Calomiris, Longhofer, and Miles (2012) and Campbell and Cocco (2007) show that age profile is very important for the relation between housing wealth and consumption where older cohorts have larger response.[6] In the context of the Great Recession, two set of authors challenged findings of Mian and Sufi also based on compositional effects. Adelino, Schoar, and Severino (2017) and Albanesi, De Giorgi, and Nosal (2017) argue that credit growth between 2001 and 2007 was concentrated in the prime segment, debt to high risk borrowers was virtually constant for all debt categories during this period, and default among high income prime borrowers were common during the post period.[7] They argue that results of Mian and Sufi confound life-cycle debt demand of borrowers who were young at the start of the boom, with an expansion in credit supply over that period.

Our unique data set will help us to solve this identification problem caused by using aggregate data, and help us to identify the channels outlined above in Figure 1. We use individual-level data from two sources that gives us most detail to-date in terms of individual

---

[6]See also Charles, Hurst, and Notowidigdo (forthcoming).

[7]Albanesi, De Giorgi, and Nosal (2017) also use individual-level data from one of the datasets we use, Federal Reserve Bank of New York/Equifax Consumer Credit Panel, but focus on growth in mortgage debt prior to crisis and subsequent defaults rather than consumption response as we do.

mortgages and Equifax Risk Scores. Our first dataset is the Federal Reserve Bank of New York/Equifax Consumer Credit Panel (CCP), a quarterly database of consumer credit bureau records for a random 5 percent sample of consumers with a credit bureau record. Our second dataset is a match between credit bureau data with more detailed information on residential first mortgages from loan servicing data. This matched dataset is Equifax Credit Risk Insight Servicing (Equifax Credit) and McDash Analytics, LLC, a wholly owned subsidiary of Black Knight Financial Services, LLC. (McDash) known as CRISM. We then restrict attention to those borrowers who can be found in the CCP. As a result we have a random representative sample of borrower-level information on all loans of the borrower, including any auto loans, borrower's Equifax Risk Score, borrower's age and detailed characteristics of the borrower's mortgages, most notably the appraised value of the property, and the type of mortgage.

As a proxy for consumption we use a binary variable at the individual level that represents origination of an auto loan in 2009. This resembles the ZIP code level new car registration data that Mian, Rao, and Sufi (2013) use in their analysis, and it has certain advantages, which we discuss in detail. The most important advantage is that it is at the individual level. Using an individual-level measure of consumption, we are able to see how a decline in housing wealth affects consumption, once we control for various aggregate variables. To further dissect the effects, we are also able to focus on various subgroups in the population based on their borrower characteristics.

In addition to changes in house prices, we have five main controls: first and foremost, the life-cycle age profile is controlled at the individual level by age and age square terms. Then, we include controls for ZIP code level car sales in 2006, change in the county-level unemployment rate between 2006 and 2009, and a measure of county-level bank health. The first variable among these aggregate controls is useful to capture preexisting differences across ZIP codes in consumption (auto purchase) behavior. We obtain this variable by aggregating our individual-level auto loan origination variable. The second variable is a key measure for capturing the general equilibrium effect in Figure 1. Finally, bank health, which we construct using one of Chodorow-Reich (2014)'s bank-level measures, distributed to counties using banks' branch shares in the county, is used to control for a county-wide decline in availability of bank credit. By using the richness of our dataset in terms of information on borrower characteristics, we interact these control variables with a number of categories, which may be as detailed as homeowners with a high Equifax Risk Score, who have a fixed-rate first mortgage, no second mortgage and a loan-to-value (LTV) ratio less than 50%, as an example.

Our results are as follows. Using both datasets, we identify the effect of the combined household wealth and financial constraints channel as accounting for 40-45% of the overall consumption response to house prices. The contribution of the decline in credit supply to households is estimated to be 20-25%. The rest, roughly 35%, as shown in Figure 1, is a general equilibrium effect that combines the feedback through reduced consumption, as well as the direct effect of the decline in credit supply to firms. In order to measure further the contribution of a wealth effect, we focus our analysis on a very specific group of consumers, which we can identify thanks to the detailed information we have in our data. These consumers have high credit Equifax Risk Scores, they own their houses outright or "free and clear" and have not moved between 2006 and 2009. Due to these characteristics, especially the absence of a mortgage, we expect that the only reason these consumers react to a decline in house prices will be due to a wealth effect. We demonstrate that, once other aggregate controls are introduced, these consumers do not react to house prices, indicating that wealth effect is negligible. This leads us to conclude that the 40-45% contribution we referred to above is solely due to households' financial constraints.

We also consider an instrumental variables (IV) strategy to account for the endogeneity of house prices and unemployment, as well as a possible omitted variable bias. We follow Aladangady (2017), Gyourko, Saiz, and Summers (2008), and Saiz (2010) to construct our instruments for house prices. As in those papers, we exploit the variation in lower land availability and tighter land use regulations that create differences in house prices across counties. We also construct a Bartik-type instrument following Keys, Tobacman, and Wang (2017) for employment changes. Our results regarding a negligible wealth effect continues to hold in our IV specification. The contribution of household financial constraints decline slightly to 35%, while the contribution of the decline in bank credit supply to households increase to roughly 50%. This is intuitive since the existence of constrained households and change in house prices in a given locality can be simultaneously determined by other characteristics of the locality, which will be controlled once house prices are instrumented for. Hence the role of exogenous-to-household bank credit supply effect increases.

Using information on mortgage characteristics further, we are also able to describe the possible reasons why financial constraints affect consumers. We distinguish between ex-ante and ex-post credit constraints. Ex-ante constraints are those that were in place in 2006, before the house prices declined, while the ex-post constraints arise, as we demonstrate, mostly due to the decline in house prices between 2006 and 2009. We show that segments of the population that are most affected by ex-ante credit constraints, such as those that

do not have high Equifax Risk Scores, have large LTVs, those that have adjustable-rate first mortgages, those that have closed-end second mortgages, or a combination of these characteristics, respond much stronger to change in house prices. These responses areup to an order of magnitude larger than those of much less constrained groups mentioned above. Taking into account both the response of these constrained groups and their population weights, at least 70% of the consumption response due to financial constraints are as a result of ex-ante constraints. Regarding ex-post credit constraints, we show that the decline in house prices is a strong predictor of whether or not the Equifax Risk Score of a consumer falls in 2009, especially for those who were borrowers with a high Equifax Risk Score and a moderate-to-large LTV in 2006. We argue that this is because these borrowers default or fall behind on their mortgages, which reduce their Equifax Risk Scores. This, in turn, means that they have difficulty in getting a loan to purchase a car, which leads to the reduction in their consumption.

The closest paper to our work is by Aladangady (2017). To the best of our knowledge this is the only other paper using individual-level data to investigate the consumption response to change in house prices. He finds results similar to ours in terms of importance of household level financial constraints. There are two main differences between our paper and his. First, we can account for general equilibrium effects and the effect of bank health. Second, we have a much larger and detailed individual-level data that help us identify both ex-ante and ex-post borrowing constraints. His key variables to identify constrained households are refinancing, household leverage and debt service, whereas we have direct data on loan types and individuals's credit risk. His results point to the key role played by financial constraints, whereas our results give an equal role to these constraints and bank health once endogeneity are accounted for.

We proceed as follows. Section 2 discusses the data in detail. Section 3 presents our econometric methodology including the IV analysis. Section 4 presents the results and Section 5 concludes.

2 Data

This section introduces our individual-level data in detail. We will go over the sources of data first and then explain how we construct our variables and show descriptive statistics.

Figure 2: Share in Census vs. Share in CCP (Counties)
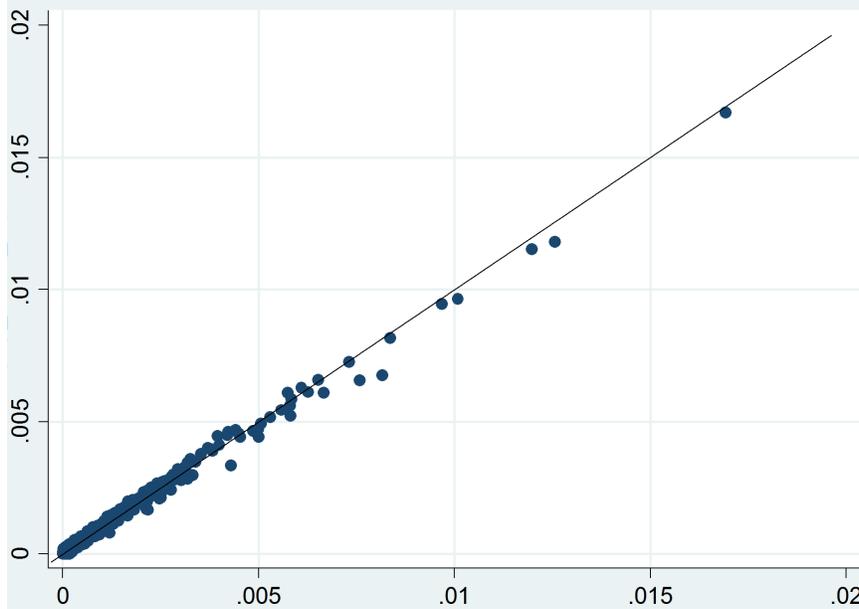
Notes: Each dot is a county. Census share and CCP share are on the x-axis and y-axis, respectively. The line shown is the regression line.

2.1 Data Sources

Our main dataset is the Federal Reserve Bank of New York/Equifax Consumer Credit Panel (CCP), a quarterly database of consumer credit bureau records for a random 5 percent sample of consumers with a bureau record. We restrict attention to primary CCP consumers. Available data fields include total balances and aggregate delinquency status on a variety of consumer credit obligations such as mortgages, auto loans and credit cards, the proprietary Equifax Risk Score, as well as some loan-level information on first and second mortgages. We are also able to calculate the age of the consumers based on the birth year that is provided in CCP. As can be seen from Figure 2, which shows the share of a county within total US census population versus the share of that county within total CCP, this dataset is representative of the broader population.

For a sample of these borrowers we have a match between their credit bureau file and more detailed information on their residential first mortgage. This matched dataset is known as CRISM.[8] This dataset is constructed by taking mortgages originated in the McDash

---

[8]See Elul and Tilson (2015) for more details on the CRISM dataset. The exact details of the matching procedure are proprietary, but it is an anonymous match, using loan amount and other loan characteristics, and is similar to that in Elul, Souleles, Chomsisengphet, Glennon, and Hunt (2010).

dataset and matching them to the primary borrowers Equifax Credit file. The McDash dataset, which forms the starting point for CRISM, captures approximately two thirds of all mortgage originations during this time period. The CRISM database begins in June 2005 and we restrict attention to consumers who had a first lien as of December 2006. The matched data gives more detailed information on the borrower's mortgages, most notably the appraised value of the property (which allows us to calculate a loan-to-value ratio), interest rate, other characteristic such as whether it is fixed or adjustable rate, low documentation, etc., and monthly mortgage performance information. We further restrict attention to those borrowers who appear in CCP (recall that this is a random 5% sample), so that we have a full panel of credit bureau variables for them.

2.2   Defining Groups of Individuals in CCP and in CRISM

Our base CCP dataset consists of 6.5 million consumers, who are in the sample in both 2006Q4 and 2009Q4, and who have an address in the same ZIP code at the start and end of the sample period. This ensures that they are all exposed to the same local aggregate house price shock. In order to correctly decompose the effect of house price changes, we classify consumers according to their homeownership status in CCP, as follows:

1. *Renters* are those who are age 55 or less in 2009, and who had no mortgages in the CCP dataset from 1999 (its inception) through 2009.

2. *Non-mover mortgage-holding homeowners* had a mortgage in both 2006Q4 and 2009Q4, and the same address in both quarters as well.

3. *Free-and-clear homeowners* had no mortgages in 2006Q4 or 2009Q4, but a mortgage at some point prior to 2006Q4, and the same address in both 2006Q4 and 2009Q4.

4. *Moving homeowners* are those with a mortgage in both endpoints but whose address changed in the interim.[9]

5. *Miscellaneous* are those who do not fit in any of the categories above (this includes borrowers with no mortgage, who are too old to be classified as renters, or those who do not have a mortgage in one of the end points.)

---

[9]For the three homeowner categories described so far, we also require the mortgage to remain in good standing between 2002 and 2009.

Table 1: Distribution of Characteristics

(a) CCP

Homeownership Status	Prime	Non-Prime	Total
Renters	5.5%	17.3%	22.8%
Free-and-Clear Homeowners	6.3%	4.2%	10.4%
Non-Mover Homeowners	25.5%	8.8%	34.3%
Moving Homeowners	1.6%	0.8%	2.4%
Miscellaneous	19.3%	10.8%	30.1%
Total	58.2%	41.8%	100.0%

(b) CRISM - 1

LTV Category	Prime	Non-Prime	Total
LTV0	43.1%	11.3%	54.3%
LTV1	22.7%	9.8%	32.5%
LTV2	9.1%	4.0%	13.2%
Total	74.9%	25.1%	100.0%

(c) CRISM - 2

Mortgage Category	Prime			Non-Prime			Total
	LTV0	LTV1	LTV2	LTV0	LTV1	LTV2	
Fixed Rate	23.2%	10.9%	4.0%	6.5%	5.3%	2.0%	51.9%
ARM < 5yr	1.2%	0.9%	0.5%	0.9%	1.2%	0.6%	5.2%
ARM ≥ 5yr	1.4%	1.4%	0.7%	0.3%	0.4%	0.2%	4.3%
CE Second	3.0%	2.2%	0.8%	1.3%	1.3%	0.5%	9.1%
HELOC	14.3%	7.3%	3.1%	2.3%	1.7%	0.7%	29.5%
Total	43.1%	22.7%	9.1%	11.3%	9.8%	4.0%	100.0%

The first panel in Table in 1 shows the share of different types of individuals in the data. Renters make up 23% of the sample, free-and-clear homeowners 10%, non-moving homeowners 34%, moving homeowners 2%, and miscellaneous 30%. Our analysis will focus mostly on the first three groups since we can clearly identify their types and argue that they constitute fairly uniform groups. The other two groups, especially the last one is one with a great deal of heterogeneity that is hard to disentangle.

For the CRISM dataset, we similarly restrict attention to borrowers who have the same

ZIP code in their address in 2006Q4 and 2009Q4. They must also have a first mortgage in both endpoints as well (although not necessarily the same one.) Our sample size is approximately 650,000 borrowers. For each homeowner in our sample, we compute an estimate of their updated first-lien loan-to-value (LTV) ratio by taking their McDash mortgage balance from December 2006, and updating the appraised value from the time of origination to December 2006. We then categorize CRISM consumers based on this updated first-lien LTV, dropping observations with updated LTV greater than 125%:

1. *LTV0*, less than or equal to 50%

2. *LTV1* above 50% and less than or equal to 80%

3. *LTV2* above 80% and less than or equal to 125%.

These groups roughly represent low, moderate and high levels of LTV. From second panel of Table 1 we see that 54% of consumers fall in the lowest category, 33% in the moderate group, and 13% in the high LTV group.

We further classify borrowers in order to analyze the effect of house price changes on consumption, based on information on the borrowers' mortgages at the end of 2006, thanks to the detailed information coming through CRISM. We define five categories and the third panel of Table 1 reports the shares of these categories in our sample. First, for borrowers who do not have a second mortgage in December 2006, we break them up into three groups, based on information from McDash on their first lien:

1. *fixed-rate first lien* (52% of total sample),

2. *adjustable-rate mortgage (ARM) with a fixed period of less than five years* (5%),

3. *ARM with a fixed period of five years or more* (4%).

Then for borrowers with a second lien, we construct two additional categories, depending on the second mortgage type, dropping 1.8% of our sample who have both types of second liens:

1. *closed-end second* (9%),

2. *home equity line of credit (HELOC)* (30%).

An important piece of data we have, which helps distinguish our work from some of the recent literature, is the Equifax Risk Score of the individual on a quarterly basis. Instead of

Figure 3: Distribution of Fraction of Non-prime Borrowers Across ZIP Codes
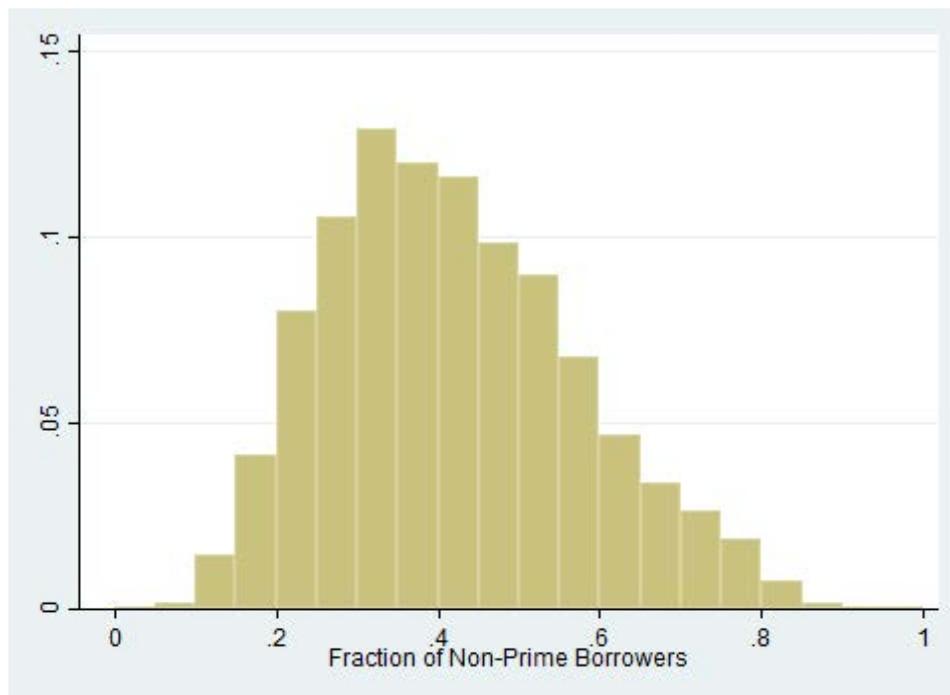
using this score directly in our analysis, we create two groups based on Equifax Risk Scores. We define a consumer as:

1. *non-prime* if he has an Equifax Risk Score of below 700[10];

2. *prime* those with Equifax Risk Scores of 700 or higher are denoted as prime borrowers.

The prime share in the CCP dataset is 58%. It is 75% in CRISM, which higher since homeowners (CRISM by definition is exclusively composed of homeowners) have higher Equifax Risk Scores. Figure 3 shows the distribution of ZIP codes with respect to the fraction of non-prime borrowers. This shows that a vast majority of ZIP codes have a mixture of prime and non-prime borrowers, and thus ZIP code level variables and the individual-level indicator of prime status will contain largely independent information. Table 1 also shows the breakdown of each of the other categories we defined above with respect to prime status. While not central to our analysis, there are some interesting observations such as renters being predominantly non-prime or non-mover homeowners being predominantly prime.

---

[10]This is a relatively high score cutoff for nonprime, and it reflects the fact that our analysis focuses on homeowners, who tend to have higher Equifax Risk Scores.

Figure 4: CEX vs. CCP: Fraction of Consumers with an Auto Loan Origination
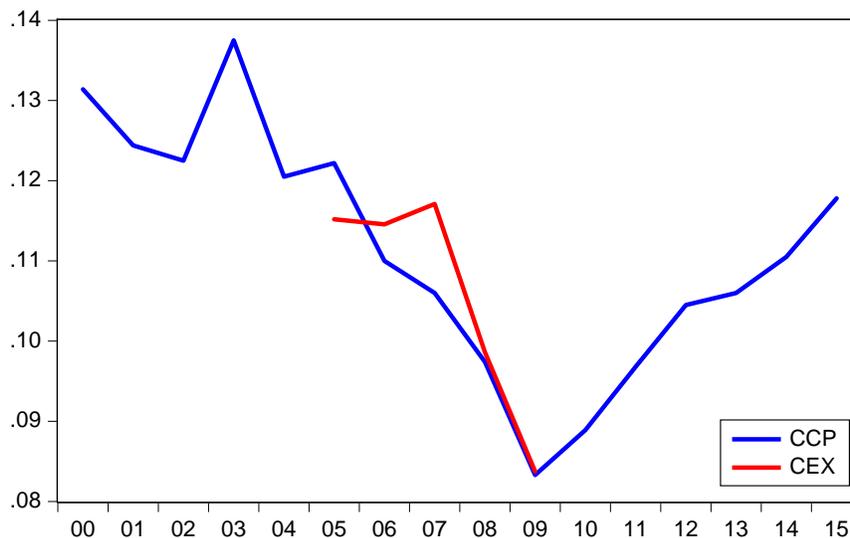
2.3 Construction of Consumption Proxy

We proxy for consumption by computing auto loan originations. As the credit bureau dataset does not give information on individual auto loans, we impute originations by tracking changes in total balances. For a given consumer in a particular quarter in the credit bureau dataset, we identify an auto loan origination by an increase in total auto loan balances of at least $1,000, relative to the previous quarter.[11] This procedure tracks the incidence of auto loan originations in other sources very well: for example, we find that 10.1% of all consumers have an auto loan origination in 2008 in the CCP, whereas from the Panel Study of Income Dynamics (PSID) the origination rate in that year is 10.8%.[12] We also find that it matches the share of auto loans originated in the Consumer Expenditure Survey (CEX) for the period of our analysis (2006-2009) as we show in Figure 4.

By contrast, Mian, Rao, and Sufi (2013) use new auto registrations from Polk, at the ZIP code level.[13] Compared to our measure, the advantage of theirs is that they are able to capture cash purchases that do not involve any financing. However, Johnson, Pence, and Vine (2014) report that about 70 percent of household purchases of new vehicles and 35 percent of household purchases of used vehicles are financed with auto loans. In addition,

---

[11]Our analysis is robust to different definitions of originations.

[12]This is computed from the 2009 wave of the PSID, using the number of respondents with a vehicle that was acquired in 2008, and the share of these which were acquired using a loan or lease.

[13]Kaplan, Mitman, and Violante (2017) replicated all the results of Mian, Rao, and Sufi (2013) using publicly available aggregate data.

Johnson, Pence, and Vine (2014) find some additional cyclicality in loan originations, as compared to auto sales, where in bad times consumers substitute away from new cars into used cars and they are more likely to use a loan to do so. Using loan originations brings with it the following advantages. First, we are able to capture both new as well as used car sales, whereas Mian, Rao, and Sufi (2013) only have new car registrations available to them. This is an important feature of our analysis, as new vehicles make up only 38% of all consumer auto purchases.[14] In addition, we are able to focus our analysis on household purchases, whereas the measure used by Mian, Rao, and Sufi (2013) also include business purchases. Finally, the most important advantage of our data is that it is at the individual level and sdince it is obtained from credit bureau data, we are also able to exploit other individual-level characteristics found in our datasets, rather than basing our analysis solely on aggregate measures.[15]

2.4   Descriptive Statistics

In Table 2 we show summary statistics of our consumption proxy, as well as four aggregate control variables we use in our analysis. For 2009, the probability that an individual originated an auto loan is 8.56% for the CCP overall, and 13.90% for the CRISM sample. For each consumer we compute the percent change in the local house prices, ΔHP index from December 2006-December 2009, using the CoreLogic Solutions single family combined house price index (ZIP code level if available, and otherwise county-level); we do this regardless of the consumers housing status. The average change is a drop of 18.4% for CCP and a drop of 20.4% for CRISM borrowers. We also compute the change in county unemployment rates from the Bureau of Labor Statistics from December 2006 to December 2009, ΔU and this averages a 5.5 percentage point increase. Both of these variables display a large degree of dispersion – in CCP the 5-95 percentile range is from -47.2% to 2.2% for ΔHP and from 2.9% to 8.7% for ΔU. We also use a ZIP code level measure of auto sales for 2006, ZIP Control, which we compute by aggregating the individual-level auto loan origination variable. This is meant to capture permanent geographical differences in auto sales, holding other things constant, for example those between Manhattan and Los Angeles which have similar characteristics in many dimensions accept for prevalence of car ownership. The units of this

---

[14]See Federal Reserve Board (2016). Furthermore, the new car share is pro-cyclical, which would tend to heighten the cyclical behavior of their measure.

[15]Another shortcoming of our measure is that it is a binary variable and we do not know the value of the auto purchased. However this is also similar to the variable used by Mian, Rao, and Sufi (2013) who start by car registrations, which do not have any value attached to them, and then aggregate to a dollar value using an aggregate auto sales measure produced by the Census Bureau. See Section 2.5 for more details.

Table 2: Summary Statistics

(a) CCP

	Mean	Std. Dev.	Min	5%	Median	95%	Max
Originate × 100	8.56	27.98	0	0	0	100	100
ΔHP	-18.4	15.1	-64.3	-47.2	-15.7	2.2	25.9
ΔU	5.50	1.86	0.0	2.9	5.3	8.7	14.0
ZIP Control	7.15	4.49	0	2.59	6.33	13.82	95.76
Bank Health	0.62	0.14	0	0.38	0.64	0.78	1.22

(b) CRISM

	Mean	Std. Dev.	Min	5%	Median	95%	Max
Originate × 100	13.90	34.60	0	0	0	100	100
ΔHP	-20.4	15.0	-64.3	-48.4	-17.7	0.8	25.9
ΔU	5.54	1.83	-0.4	3.0	5.4	8.7	15.0
ZIP Control	7.95	5.34	0	3.22	6.80	15.88	95.76
Bank Health	0.64	0.12	0	0.41	0.65	0.78	1.22

Notes: Originate is a variable that can take the values 0 or 1. The statistics in this table are for Originate × 100 for more detail. ΔHP and ΔU are the changes in house prices and unemployment rate, respectively, expressed in percentage points. Bank Health shows the fraction of the syndication portfolio of the banks in a county, in which Lehman Brothers had a lead role for the banks in a county and it is in percentage points. ZIP control shows the per-capita sales of cars in 2006 based on the origination variable, expressed in thousands of dollars.

variable is in $1,000 per capita and it averages $7,150 in CCP and $7,950 in CRISM.

Our final aggregate variable, Bank Health, is a county-level version of a key indicator of bank health as provided by Chodorow-Reich (2014). We start by the bank-level measure of fraction of the syndication portfolio where Lehman Brothers had a lead role. Next, we collect information on how many branches/affiliates each bank is located in each of the U.S. counties.[16] The final step is to distribute the national value of the banks to counties by using the share of branches each bank has in the county.[17] The resulting variable will be such that a higher values indicates worse bank health and will capture a decline in availability of credit in the county.

In some of our analysis, we address the endogeneity of house prices using two instruments.

---

[16]Lists of branches and their addresses are from the Federal Financial Institutions Examination Council's (FFIEC) and the banks' websites. ZIP codes of banks' addresses are then matched with the county names using the FIPS county-code sheet from US Census. When a ZIP code is shared by two or more counties, we manually look up that branch's address in Google Maps to determine which county it belongs to.

[17]For example, if there are two banks in a county with national bank health values of X_1 and X_2 and B_1 and B_2 branches in a county, then the county's bank health will be $(B_1 X_1 + B_2 X_2)/(B_1 + B_2)$.

Table 3: ZIP Code-Level Analysis

Dependent variable: Change in Consumption (2006 to 2009)

	Mian Rao and Sufi (2013)	Our Measure
Change in HP	0.018	0.004
(2006 to 2009)	(0.0010)	(0.0004)
R^2	0.153	0.025
Observations	6,263	6,220

Notes: The dependent variables in the regressions are computed using auto registrations for Mian, Rao, and Sufi (2013) and car loan originations for our measure. Observations are at the ZIP code level, aggregated using the method in MRS. Regressions include a constant which is not reported. All regressions are weighted by the number of households in the ZIP code. Robust standard errors are reported in parenthesis.

Both of these instruments capture the elasticity of housing supply, and therefore the response of house prices to demand shocks. First is the share of land in the borrower's MSA that is unavailable for real estate development, from Saiz (2010), which reflects physical constraints governing land development. In addition, we use the MSA-level Wharton Residential Land Use Regulation Index (WRLURI) from Gyourko, Saiz, and Summers (2008). The WRLURI is a summary measure of the stringency of the local regulatory environment in each MSA, based both on local and state-level factors, with higher levels reflecting greater stringency. We also construct a Bartik-style instrument for the change in county-level unemployment rates from 2006-2009, along the lines of Keys, Tobacman, and Wang (2017), by using the interaction of the 2003 industry mix of employment in that local labor market and the national change in industry employment (exclusive of the given county) from 2006-2009. These measures are constructed using the Quarterly Census of Employment and Wages (QCEW) at the county level.

2.5 ZIP Code-Level Analysis

To obtain a ZIP code level dataset analogous to that of Mian, Rao, and Sufi (2013), we take our CCP-based individual-level auto loan origination variable, and aggregate to the ZIP code level. This gives us 6,224 observations that simply count the number of auto loan originations in a ZIP code.[18] Along the lines of Mian, Rao, and Sufi (2013), we then allocate

---

[18]We have 43 less ZIP codes relative to Mian, Rao, and Sufi (2013). The difference may be due to the fact that we do not need to restrict to ZIP codes represented in the Polk data and we also use a more recent release of the CoreLogic Solutions house price index which affects the availability of the house price index

annual national retail auto sales (from the Census Bureau) across ZIP codes in proportion to their share of auto loan originations in our data; for example, if a ZIP code in our dataset accounted for 5% of all auto loan originations for that year, it would be allocated 5% of national retail auto sales. In contrast to Mian, Rao, and Sufi (2013), we use total auto sales, both new and used, since our loan origination data does not distinguish between the two (and, as we have argued above, this is more appropriate when studying consumer spending). We then divide by the number of households in the ZIP code, which we obtain by applying the national population growth rate to the ZIP code populations in the 2000 Census. Note that the ZIP Code control variable we referred to in Section 2.1 is the 2006 version of this variable.

Table 3 shows our replication of the results reported in column (5) in Table V of Mian, Rao, and Sufi (2013). This is a simple OLS regression with change in consumption between 2006 and 2009 as the dependent variable and change in house prices in the same period as the independent variable. Their estimate shows an $18 decline in auto consumption for every $1,000 decline in house values. It is highly significant at 1% level. Our results shows a smaller elasticity, $4 for every $1,000, which is also highly significant. This is reasonable due to the exclusion of used car purchases in the measure used by Mian, Rao, and Sufi (2013). When a consumer chooses to buy a used car instead of a new car in 2009, this purchase shows up in our dataset (and thus consumption in 2009 does not fall as much) while it does not show up in the measure used by Mian, Rao, and Sufi (2013).

3   Empirical Strategy: Individual-Level Analysis

As we explained in the previous section, our key dependent variable, auto loan originations for an individual in 2009, is a binary variable. We conduct our analysis by estimating various linear probability models using ordinary least squares (OLS) or instrumental variables (IV). Results are very similar if we use a probit model instead of a linear probability model.

The generic equation we estimate in either of our datasets is

$$y_{izc} = \alpha + \lambda_1 age_i + \lambda_2 age_i^2 + \sum_{k=1}^{K} \sum_{j=1}^{4} \beta_{jk} C_{izc}^k X_{cz}^j + \varepsilon_{izc} \tag{1}$$

where the subscripts i, z and c refer to an individual, a ZIP code and a county. The dependent variable shows whether or not the individual originated an auto loan in 2009. We control

---

used in the analysis.

for any life-cycle effects by a quadratic polynomial in age.

In our full model, we have four other controls, each of which are interacted with a full set of individual-level categorical variables. These controls, denoted by X_{cz}^j, are ZIP code-level house price change, ΔHP, county-level change in the unemployment rate, ΔU, the 2006 ZIP code auto sales control ZIP Control and the county-level bank health variable Bank Health in (1). We keep the age polynomial, ΔHP and ZIP Control in all regressions and in addition to the full model consider specifications that exclude one or both of the remaining two control variables. These regressions help us identify the key channels of the effect of house prices as we explain shortly.

All four of our control variables are interacted by a full set of dummy variables obtained from up to three individual-level categorical variables, which are denoted by C_{icz}^k for $k = 1, ..., K$. In regressions using CCP we categorize individuals in two dimensions: whether or not they are prime (two values) and their homeownership status (five values). Considering all combinations, and dropping as necessary to avoid multicollinearity, we get $K = 9$ interaction variables per control variable. In CRISM, on the other hand, we can have up to a three-way interaction that includes prime status (two values), mortgage type (five values) and LTV (three values), leading to $K = 29$ interaction variables per control variable.

In all our estimations we cluster standard errors at the ZIP code level. Since our estimations result in tens of coefficient estimates, we focus on one key number, the average marginal impact of a unit change in ΔHP, and report it either in aggregate or for various subcategories j. In OLS, naturally this amounts to the sum of the appropriate combinations of β_{jk}. We compute the standard error of these marginal impacts using the delta method.

We use three regressions in each of our datasets in order to identify the importance of the three channels for explaining the effect of changes in house prices on consumption. We start by using only ΔHP and ZIP Control as controls. We record the marginal impact for ΔHP and normalize this to 100. Next we add ΔU to the model and compute the marginal impact for ΔHP in this model. Controlling for ΔU typically reduces the marginal impact for ΔHP and this decline relative to the marginal impact we obtained in the first regression is our measure of the general equilibrium effect. Next we add Bank Health in to the regression and compute the marginal impact for ΔHP – this is our full model. The difference between this and the one we computed from the second regression is our measure of the effect of the decline in bank credit supply to households. Finally, after using all the controls, what remains in terms of the marginal impact of ΔHP is the combination of the household wealth and the household credit constraints.

We identify the magnitude of the wealth effect using three segments of the population, two using CCP and one using CRISM. These are (a) prime homeowners who did not move between 2006 and 2009 and own their houses without a mortgage or "free and clear"; (b) prime homeowners who did not move between 2006 and 2009 and hold a mortgage; (c) prime homeowners who have a fixed-rate mortgage, no second mortgage and an LTV that is less than 50%. Recall that for us to label it wealth effect, a consumer should react to a change in house prices only because it reduces his wealth, and not because some constraints the consumer faces either today or in the future become more binding, or because the change in house prices are correlated with other aggregate things (such as unemployment risk) he cares about. All three segments of the population we use for this purpose fit this broad definition. First, because they are all prime, they are less sensitive to aggregate conditions we may not be controlling for. Second, the free-and-clear group does not hold a mortgage and thus they have no financial constraints that is directly related to house prices. Similarly, the second and third groups are least likely to have binding financial constraints. The third group is especially relevant since they are not worried about changing terms of their mortgage when house prices change. Moreover with a low LTV, they are immune to adverse effects of large changes in house prices – for example it would take a decline in house prices over 50% to wipe out their equity in their house, which happened for only a small fraction of homeowners in our sample.

The change in house prices and employment are endogenous. Note that, it is not plausible to have individual level auto loan origination to effect ZIP code level house prices and county level employment, and hence we do not worry about reverse causality, nevertheless an omitted factor, both at the ZIP code and/or county level may drive our dependent and independent variables simultaneously. This is why we instrument both house prices changes and employment changes. We follow the literature to instrument house prices changes based on elasticity of housing supply and for employment changes we construct a Bartik-type instrument.

4   Results

This section presents our decomposition results, first with CCP and then with CRISM. We then provide IV results .

Table 4: CCP Decomposition

	Only ΔHP	ΔHP and ΔU	Full
Overall HP Effect	0.0353 (***)	0.0230 (***)	0.0141 (***)
% of Only Δ HP	100%	65%	40%
Categories			
Prime	0.0275 (***)	0.0156 (***)	0.0034
Non-Prime	0.0527 (***)	0.0390 (***)	0.0348 (***)
Renters	0.0267 (***)	0.0151 (***)	0.0073 (**)
Free-and-Clear Homeowners	0.0544 (***)	0.0351 (***)	0.0225 (***)
Non-Mover Homeowners with Mortgage	0.0220 (***)	0.0167 (***)	0.0102 (***)
Moving Homeowners with Mortgage	0.0894 (***)	0.0694 (***)	0.0583 (***)
Miscellaneous	0.0295 (***)	0.0217 (***)	0.0155 (***)
Prime Renters	0.0272 (***)	0.0145 (***)	0.0028
Prime Free-and-Clear Homeowners	0.0397 (***)	0.0212 (***)	0.0034
Prime Non-Mover Homeowners with Mortgage	0.0123 (***)	0.0096 (***)	0.0008
Number obs.	6,553,884	6,553,884	6,553,884

Notes: All regressions include age, age^2 and as well as 2006 ZIP-code control interacted with a full set of dummies. (***), (**) and (*) denote significance at 1%, 5% and 10% levels, respectively.

4.1 Decomposition of Channels using CCP

We start by estimating the model using CCP. As we discussed earlier, CCP is representative of the U.S. population and as we now demonstrate it contains a large degree of heterogeneity. Table 4 reports our results.[19] Each column shows the estimated model, starting from the one with only ΔHP and ZIP Control, then adding ΔU and Bank Health in the second and third columns.

The first row shows the overall marginal impact of the change in house prices on consumption. Before controlling for key aggregate variables, the marginal impact on consumption is 0.0353 and it is highly significant. To put this number in perspective, since the average change in house prices is −18.4, our results show that the probability of originating a car loan goes down by about 0.65 percentage points. Considering that the unconditional probability of originating an auto loan is 8.56%, this is a sizable response.

Controlling for ΔU reduces the effect of house prices by 35% and controlling for Bank Health reduces it by a further 25%. These results constitute our first key result. Out of

---

[19]In all tables that follow we use (***), (**) and (*) to denote significance at 1%, 5% and 10% levels, respectively. Moreover unless otherwise specified, these tables will report the marginal impact of a unit change in ΔHP in the aggregate or for some subgroups of individuals.

the overall response of consumption to house prices measured without controls (except for ZIP-code sales in 2006), 35% of it is explained by general equilibrium effects and 25% of it is explained by the decline in credit due to deteriorating bank health. The remaining 40% is the direct effect of house prices on consumption – the leftmost arrow in Figure 1. Much of the rest of the paper will be devoted to understanding and further decomposing this 40% response.

In the rest of Table 4, we show how various subgroups in our sample are affected by the change in house prices, and how this varies across different specifications. In CCP we have two main categories: prime / nonprime and the five homeownership categories we defined in Section 2. The third and fourth rows of Table 4 show the results for the prime and non-prime groups and the next five rows show the results for each homeownership category. Finally the next three rows show the results for prime individuals who belong to one of the three key homeownership categories. Looking at the first column reveals the extent of heterogeneity: non-prime consumers react almost twice as much as prime consumers; homeowners with a mortgage who has moved respond four times as much as those that did not move. Controlling for ΔU reduces the responses across the board but all groups show highly statistically significant responses. When we also control for Bank Health in the last column, some very interesting results emerge. First, prime consumers' reaction to changes in house prices become insignificant. Second, looking deeper, the reactions of prime consumers in all three important homeownership categories become insignificant. This shows that the strong responses for these groups that we found in the first column, were not actually due to the decline in house prices but due to other aggregate developments (such as the increase in unemployment or decline in bank health), which are related to but distinct from the decline in house prices. Third, despite the decline in the overall response, there is still considerable heterogeneity in responses: homeowners with a mortgage who has moved now respond over five times as much as those that did not.

The results for the two homeowner groups allow us to disentangle the household wealth effect and the effect of household financial constraints. Consumers that are prime and either own their home without a mortgage or have a mortgage and have not moved, would be most immune from any other effect of house price changes but the wealth effect. Thus the results for these groups allow us to conclude that the wealth effect is negligible and all of the 40% that we attributed to the wealth effect and the effect of household financial constraints is indeed due to the latter.

Table 5: CRISM OLS Results

(a) Decomposition

	Only ΔHP	Δ HP and ΔU	Full
Overall HP Effect	0.0682 (***)	0.0436 (***)	0.0298 (***)
% of Only Δ HP	100%	64%	44%
Categories			
Prime	0.0566 (***)	0.0351 (***)	0.0186 (***)
Non-Prime	0.1020 (***)	0.0699 (***)	0.0647 (***)
LTV0	0.0540 (***)	0.0240 (***)	0.0097 (*)
LTV1	0.0755 (***)	0.0574 (***)	0.0457 (***)
LTV2	0.1071 (***)	0.0920 (***)	0.0764 (***)
Prime LTV0	0.0435 (***)	0.0143 (**)	-0.0037
Number obs.	677,918	677,918	677,918

Notes: (***), (**) and (*) denote significance at 1%, 5% and 10% levels, respectively.

(b) Implied Change in Probability of Origination Based on Average Value in Variable (in p.p.)

	Only ΔHP	ΔHP and ΔU	Full
House Price	-1.39	-0.81	-0.61
Unemployment	-	-1.72	-1.83
Bank Health	-	-	-3.64

4.2 Decomposition of Channels using CRISM

In this section we confirm that the results regarding the decomposition of the house price response also hold using CRISM. We do so using OLS like we did with CCP in the previous section, as well as an IV specification that accounts of the endogeneity in house price and unemployment changes.

The first panel of Table 5 shows the results from CRISM analogous to Table 4. To keep things simple we only report results for the prime / nonprime groups, the three LTV groups and the prime LTV0 group. The first two rows show that the overall results are larger than those for CCP, but remarkably similar in therm of the percentage changes across

columns reported in the second row. Controlling for the change in unemployment reduces the house price response by 36% and further controlling for bank health reduces the response by another 20%.

Looking at the rest of the first panel of Table 5, we see that the largest effect of controlling for the change in unemployment and bank health was on prime consumers, whose response goes down by almost 70%. This causes a difference that is 3.5 times between prime and non-prime consumers. As we further show in the next section, the non-prime response itself is very heterogenous. LTV is also an important determinant of their consumption response: those with low LTV have a negligible response while those with higher LTVs show sizable responses; those with LTV greater than 80% respond over 2.5 times the average response.

Finally, in CRISM, we argue we can identify the wealth effect by prime consumers whose LTV is less than 50%. This group does not show a statistically significant response once we properly include the controls. Thus we again conclude that the wealth effect is negligible.

Our discussion so far has focused on the effect of the house price changes on consumption and how controlling for the change in the unemployment rate or bank health is crucial for properly measuring this. However, it is important to emphasize that these two variables actually do more than just absorbing some of the effect of house prices on consumption. The second panel of Table 5 show how each of the three variables contribute to explaining the probability of originating an auto loan. To ease interpretation, we report the change in origination probability due to each variable, which multiplies the marginal effect of each variable with its average of value. For reference, recall that the unconditional probability of originating an auto loan in the CRISM sample is 13.9%. Without other controls, ΔHP would create a decline of 1.39 percentage points, which, at 10% of the unconditional probability, is really large. Once other controls are introduced, this falls down to 0.61 percentage points. In contrast, the average change in unemployment reduces the origination probability by 1.83 percentage points and the average bank health reduces it by 3.64 percentage points. This shows that the two aggregate controls have very important independent effects on consumption.

We also conduct an IV estimation using CRISM in order to take into account endogeneity and omitted variable problems. We introduced the three instruments we use in Section 2. Since our full specification includes interactions of what we consider to be endogenous variables in this IV (ΔHP and ΔU) and individual level dummy variables, instead of estimating the model using the full sample, we estimate separate IV models for each subsample. This is equivalent to, and simpler than, running a single IV regression. We do this using six

Table 6: CRISM IV Results

(a) Sample First Stage (For Full Model, LTV0, Prime)

	ΔHP	ΔU
WRLURI	-0.0089 (***)	-0.2887 (***)
Unavailable	-0.2345 (***)	3.1699 (***)
Bartik	1.3103 (***)	-23.6776 (***)
ZIP Code Control	-0.0048 (***)	0.0233 (***)
Bank Health	-22.8514 (***)	162.3166 (***)
N	251,169	251,169
R^2	0.21	0.19
F-stat	252.05	302.2

Notes: A constant and estimates for age and age^2 are omitted from the table. (***), (**) and (*) denote significance at 1%, 5% and 10% levels, respectively.

(b) Marginal effect of ΔHP in various IV specifications.

Category	HP Only	HP and U	Full
Prime LTV0	0.0924 (***)	0.0465	-0.0349
Prime LTV1	0.1148 (***)	0.1262 (***)	0.0834 (**)
Prime LTV2	0.1556 (***)	0.2242 (***)	0.1838 (***)
Non-Prime LTV0	0.0914 (***)	0.0283 (***)	0.0022
Non-Prime LTV1	0.1193 (***)	0.1266 (***)	0.1227 (**)
Non-Prime LTV2	0.1616 (***)	0.1421 (*)	0.1028
Overall	0.1087 (***)	0.0909 (***)	0.0379 (***)
% of HP Only	100%	84%	35%

Notes: (***), (**) and (*) denote significance at 1%, 5% and 10% levels, respectively.

subsamples where we group consumers based on their prime status and LTV. Our results are presented in Table 6. The first panel shows the first stages in one of the subgroups as an example. All other subgroups have very similar first stage estimates both in terms of signs and magnitudes. All first stages pass weak and under-identification tests.

The second panel shows the marginal effects for each subgroup for the three specifications in terms of which controls are included. There are some significant differences relative to the OLS results – IV results are typically larger. Focusing on the breakdown at the last row, which is computed as the weighted average of the six subsample results, the importance of ΔU is smallerthan the OLS results at 16%, and the importance of the bank credit supply

Table 7: Decomposition of Channels

	CCP	CRISM - OLS	CRISM - IV
Household Wealth	0%	0%	0%
Household Financial Constraints	40%	44%	35%
Bank Credit Supply to Households	25%	20%	49%
General Equilibrium	35%	36%	16%

to households is larger at nearly 50%. The wealth effect is still negligible as shown by the response of prime LTV0 borrowers.

Table 7 summarizes our results in terms of the importance of each channel across different specifications and datasets.

4.3 Household Financial Constraints

Our results so far show that between 35% and 44% of the overall consumption response to house price changes is driven by household financial constraints. In this section we investigate further and attempt to identify which constraints are responsible for this large response.

We think of financial constraints in two broad categories: ex-ante and ex-post. By ex-ante financial constraints we mean those that affected consumers in 2006 or earlier, before house prices declined. Ex-post constraints are those that affect the consumers in 2009 and they are likely tightened at least in part due to the decline in house prices. Our detailed individual-level data allows us to cut the data various ways to identify these constraints.

4.3.1 Ex-Ante Constraints

We provide three ways of observing ex-ante financial constraints at work. First two are shown in Table 8. Here, using CRISM we show how the interaction of LTV and prime status affects the consumption response to ΔHP. This is just a more detailed breakdown of the results in Table 5. There are three clear conclusions. First, being non-prime in 2006 significantly increases the consumption response in 2009. Second, having high LTV in 2006 also significantly increases the response. In fact, while the prime and LTV0 groups each show no response to ΔHP, either being non-prime or having higher LTV matters a lot. Third, not surprisingly, the interaction of the two creates a significant consumption response. Once again after appropriate rescaling, the 0.1100 marginal effect we show correspond to a 1.5

Table 8: Ex-Ante Financial Constraints 1 - Prime Status and LTV

	Prime	Non-Prime	Overall
LTV0	-0.0037	0.0495 (***)	0.0097
LTV1	0.0371 (***)	0.0717 (***)	0.0457 (***)
LTV2	0.0651 (***)	0.1100 (***)	0.0764 (***)
Overall	0.0186 (***)	0.0647 (***)	0.0298 (***)

Notes: (***), (**) and (*) denote significance at 1%, 5% and 10% levels, respectively.

percentage point decline in the probability of auto loan origination.

We view both of these characteristics as being a sign of having constraints in 2006. Being nonprime shows the presence of some adverse credit activity and this further influences the type of credit the consumer gets access to. Moreover, non-prime status very persistent. Our data shows that there is a 72% probability that a person who is non-prime in 2006 remains non-prime in 2009. Being non-prime in 2009 has obvious adverse effects on access to credit in 2009 and this limits how much consumption the consumer can have, especially using our measure of auto loan originations. LTV in 2006 directly reflect the severity of one of the most important financial constraints, the implicit collateral constraint of a mortgage. The higher the LTV, the more constrained the consumer is and thus the more vulnerable he is to house price changes. To sum up, both of these characteristics have implications about how easy it is for the consumers to refinance their mortgage, how likely it is for them to default and more generally how constrained they are.

The third way of identifying ex-ante constraints in our data is presented in Table 9. To produce this table, we repeat our benchmark CRISM estimation but in addition to prime status and LTV, we include a third layer of interaction with the mortgage type variable defined in Section 2.1. To see how mortgage type is a sign of ex-ante constraints, note that borrowers are not allocated randomly to different mortgage types, but they select the mortgage that best suits their situation, including financial constraints they face. For example, borrowers with closed-end second mortgages typically get these mortgages because they lack the resources to make a 20% downpayment, which is the standard amount in most mortgages. Further analyzing the distribution of consumers in Table 1 we see a few more interesting patterns that suggest choices by consumers. For example short-maturity ARMs seems to be chosen by prime low LTV borrowers (perhaps because they intend to pay off their loan in a short period of time) or non-prime moderate-LTV borrowers (perhaps becuase this was the only product they qualified for and they hope to refinance before the ARM re-

Table 9: Ex-Ante Financial Constraints 2 - Mortgage Type

(a) Marginal effects by category

	Prime			**Non-Prime**		
	LTV0	**LTV1**	**LTV2**	**LTV0**	**LTV1**	**LTV2**
Fixed Rate	0.0044	0.0432 (***)	0.0613 (***)	0.0523 (***)	0.0565 (***)	0.0826 (***)
ARM < 5yr	-0.0381	0.0492	0.108 (**)	0.0544	0.0781 (**)	0.0593
ARM ≥ 5yr	-0.077 (***)	-0.0558 (*)	0.007	-0.0272	0.0899	0.0941
CE Second	0.0485 (**)	0.0787 (***)	0.1063 (**)	0.092 (***)	0.0793 (**)	0.219 (***)
HELOC	-0.0177	0.0238	0.047 (**)	0.0196	0.0873 (***)	0.099 (**)

Notes: (***), (**) and (*) denote significance at 1%, 5% and 10% levels, respectively. Color coding in cells show the weight of each cell in the overall CRISM population using the distribution reported in Table 1. Green represents a group with more than 10% weight, yellow shows a weight between 5% and 10% and purple shows a group that has a weight between 1% and 5%.

(b) Share of each category in overall effect

	Prime			**Non-Prime**			
	LTV0	**LTV1**	**LTV2**	**LTV0**	**LTV1**	**LTV2**	**Total**
Fixed Rate	4%	17%	9%	12%	11%	6%	58%
ARM < 5yr	-2%	2%	2%	2%	3%	1%	8%
ARM ≥ 5yr	-4%	-3%	0%	0%	1%	1%	-5%
CE Second	5%	6%	3%	4%	4%	4%	27%
HELOC	-9%	6%	5%	2%	5%	3%	12%
Total	-6%	28%	19%	20%	24%	15%	100%

Notes: This table takes the marginal effect of a cell in panel (a), multiplies with the share of this cell in the population as reported in Table 1 and divides by the overall effect. Red color denotes cells whose contribution is greater than 10%.

sets). HELOCs seem to be favored by prime borrowers with low-to-moderate LTVs. It is plausible that these consumers use the extra liquidity from their HELOCs to finance some consumption expenditures. Thus a decline in house prices would make their constraints bind since banks can (and did) reduce HELOC limits of consumers with increased LTVs.[20]

The first panel of Table 9 shows the marginal impact of house price changes, broken down into these three categories. To focus on the important results, we use colors to show the weight of each subgroup in the whole population using the distribution in Table 1. Green represents a group with more than 10% weight, yellow shows a weight between 5% and

---

[20]One may be tempted to think consumers can use cash they get from their HELOCs to finance an auto purchase completely without the need for an auto loan. If this was the case then it is not clear how we can identify our results based on auto loan originations for people with HELOCs. Results reported by McCully and Vine (2015) show that very few consumers purchase cars outright using HELOCs or cash-out refinancing using data from three nationally representative surveys.

10% and purple shows a group that has a weight between 1% and 5%. We find that prime borrowers with a fixed-rate mortgage and a low LTV, a group that represents over 23% of the population, do not respond to changes in house prices. This is yet another indication that wealth effect is not important as these consumers would have no financial constraints. We find significant responses for the remainder of the fixed-rate group who have higher LTV and/or are non-prime. There are negligible responses from consumers with short-maturity ARMs. Prime consumers with long-maturity ARMs, on the other hand, respond strongly and negatively to house price changes. This means they benefited from the decline in house prices. Consumers with closed-end second mortgages seem to be the most responsive group where those that are non-prime and have high LTV shows a response that is over 7 times the average response. Finally consumers with HELOCs that also have moderate-to-high LTVs show significant responses.

While the results in the preceding paragraph are interesting, they do not fully answer our main goal of finding out what constraints are responsible for the large response of consumption to house prices. To do so, we compute the contribution of each cell in the first panel of Table 9 to the overall response. This amounts to taking the marginal effect of a particular group, multiplying by the weight in population in Table 1 and dividing by the overall response. Results are reported in the second panel of Table 9. To ease interpretation, we highlight cells that show a contribution grater than 10%.

We find that almost 60% of the consumption response to changes in house prices come from consumers with fixed-rate mortgages, especially those that are non-prime (29%) or are prime and have moderate-to-high LTV (26%). The contribution of consumers with ARMs is negligible. Consumers with second mortgages constitute about 40% of the response, with those with closed-end second at 27%.

Our reading of these results are as follows. We think all of the non-prime response, amounting to 59%, represents ex-ante constraints. Similarly, any mortgage choice other than a fixed rate mortgage also shows the presence of ex-ante constraints, which is an additional 11%. Thus our results indicate that at least 70% of the house price response stripped of any general equilibrium and bank credit effects (recall that this itself is 44% of the full response) is due to ex-ante credit constraints.

4.3.2 Ex-Post Constraints

To demonstrate the importance of ex-post financial constraints, those that become more binding due to the decline in house prices, we consider a simple analysis. We use the same

Table 10: Ex-Post Financial Constraints - 2009 Prime Status

	2006 Prime	2006 Non-Prime	Overall
LTV0	2.18 (***)	0.29	1.69 (***)
LTV1	5.63 (***)	3.02 (***)	4.86 (***)
LTV2	8.78 (***)	5.24 (***)	7.82 (***)
Overall	4.24 (***)	1.82 (***)	3.53 (***)

Notes: The table reports the marginal effect of each characteristic given in a cell to the probability of becoming non-prime in 2009, expressed in percentage points. Unconditional probability of being non-prime in 2009 is 26%. (***), (**) and (*) denote significance at 1%, 5% and 10% levels, respectively.

regression model we used in CRISM, where all controls are interacted by all the combinations of 2006 prime status and LTV, to predict consumers' 2009 prime status. The goal here is to demonstrate the importance of ΔHP in converting borrowers to non-prime in 2009, which is precisely how our concept of ex-post constraints work. Table 10 shows the results. Here to ease interpretation all marginal effects are converted to change in the probability of becoming non-prime. To put things in perspective, the unconditional probability of being non-prime in 2009 is 26%. It is also helpful to note that even after controlling for the effects of the control variables, being non-prime in 2006 increases the probability of being non-prime in 2009 by 15 percentage points.

Table 10 shows that the average decline in house prices lead to a 3.53 percentage point increase in the probability of an average person to become non-prime in 2009. This is already very sizable. When we look at those that were prime in 2006 and especially those that had moderate-to-high LTV, the increase in probability is 1.5 to 2.5 times larger – as high as 8.78 percentage points for prime borrowers who had a high LTV in 2006. Borrowers that were non-prime in 2006, on the other hand, are much less affected by the house price decline – the average effect is about half of the overall effect. This is because the large persistence in non-prime status we mentioned above.

Our results regarding prime borrowers in 2006 explain why in Table 9 about 40% of the response to house prices come from prime borrowers. While some of them certainly could have financial constraints that influence them in 2006, at least for some their reaction to consumption is due to the simple reason that they have become non-prime due to the decline in house prices. We think that this is likely due to the consumer falling behind (either voluntarily or involuntarily) his mortgage payments. This, in turn, reduced the creditworthiness of the consumer to the point where we label him non-prime. Crucially, it also means that he is less likely to qualify for an auto loan and thus we observe that he

reduces his consumption.

It is important to emphasize that this result seems to also be relevant for the current debate in the literature in terms of whether the boom period mortgage borrowing was driven by prime and sub-prime borrowers. Albanesi, De Giorgi, and Nosal (2017) show that borrowing in sub-prime ZIP codes is driven by prime borrowers in those ZIP codes, whereas borrowing by subprime individuals was constant in the boom period. Our results of prime borrowers turning into non-prime due to a decline in house prices is consistent with these findings given that prime borrowers did most of the borrowing during boom period, which then fell behind in their payments.

5   Conclusion

We use individual level data to decompose the response of consumption to declining house prices during 2006–2009. We find that wealth effect is not important for this response, whereas financially constrained households and lower credit supply from banks who got hit by the crisis explain most of the response. Our decomposition exercise is based on accounting for the role of employment changes and bank health in a given county and identifying the groups of individuals carefully so that persons exposed to wealth effects and a possible tightening of financial constraints can be investigated separately.

In terms of our estimates, tightening household-level financial constraints can explain 40-45 percent of the response of consumption to declining house prices. Deteriorating bank health leads to reduced credit supply to households which explains 20-25 percent of the consumption response. The remaining 35 percent is a general equilibrium effect that works via a decline in employment as a result of either lower credit supply to firms or the feedback from lower consumer demand.

Using elasticity for housing supply and prior national sectoral employment growth as instruments for changes in house prices and unemployment, we run IV regressions. Our estimate of a negligible wealth effect is robust to accounting for the endogeneity of house prices and unemployment. The contribution of tightening household financial constraints goes down to 35 percent, whereas declining bank credit supply to households captures about half of the overall consumption response, once we account for endogeneity.

References

ADELINO, M., A. SCHOAR, AND F. SEVERINO (2017): "Dynamics of Housing Debt in the Recent Boom and Great Recession," Discussion paper, National Bureau of Economic Research.

ALADANGADY, A. (2017): "Housing Wealth and Consumption: Evidence from Geographically-linked Microdata," *American Economic Review*, 107(11), 3415–46.

ALBANESI, S., G. DE GIORGI, AND J. NOSAL (2017): "Credit Growth and the Financial Crisis: A New Narrative," Discussion paper, National Bureau of Economic Research.

ATTANASIO, O. P., L. BLOW, R. HAMILTON, AND A. LEICESTER (2009): "Booms and Busts: Consumption, House Prices and Expectations," *Economica*, 76(301), 20–50.

BAHAJ, S. A., A. FOULIS, AND G. PINTER (2017): "Home Values and Firm Behaviour," .

BARRO, R. J. (1976): "The Loan Market, Collateral, and Rates of Interest," *Journal of Money, Credit and Banking*, 8(4), 439–456.

BERGER, D., V. GUERRIERI, G. LORENZONI, AND J. VAVRA (2015): "House prices and consumer spending," *The Review of Economic Studies*.

BERNANKE, B. S., M. GERTLER, AND S. GILCHRIST (1999): "The Financial Accelerator in a Quantitative Business Cycle Framework," *Handbook of macroeconomics*, 1, 1341–1393.

BUITER, W. H. (2008): "Housing Wealth Isn't Wealth," Discussion paper, National Bureau of Economic Research.

CALOMIRIS, C., S. D. LONGHOFER, AND W. MILES (2009): "The (Mythical?) Housing Wealth Effect," Discussion paper, National Bureau of Economic Research.

CALOMIRIS, C. W., S. D. LONGHOFER, AND W. MILES (2012): "The Housing Wealth Effect: The Crucial Roles of Demographics, Wealth Distribution and Wealth Shares," Discussion paper, National Bureau of Economic Research.

CAMPBELL, J. Y., AND J. F. COCCO (2007): "How Do House Prices Affect Consumption? Evidence from Micro Data," *Journal of Monetary Economics*, 54(3), 591–621.

CARROLL, C. D., AND M. S. KIMBALL (1996): "On the Concavity of the Consumption Function," *Econometrica*, pp. 981–992.

CARROLL, C. D., M. OTSUKA, AND J. SLACALEK (2011): "How Large are Housing and Financial Wealth Effects? A New Approach," *Journal of Money, Credit and Banking*, 43(1), 55–79.

CASE, K. E., J. M. QUIGLEY, AND R. J. SHILLER (2005): "Comparing Wealth Effects: the Stock Market versus the Housing Market," *Advances in Macroeconomics*, 5(1).

CHARLES, K. K., E. HURST, AND M. J. NOTOWIDIGDO (forthcoming): "Housing Booms and Busts, Labor Market Opportunities, and College Attendance," .

CHEN, B. S., S. G. HANSON, AND J. C. STEIN (2017): "The Decline of Big-bank Lending to Small Business: Dynamic Impacts on Local Credit and Labor Markets," Discussion paper, National Bureau of Economic Research.

CHODOROW-REICH, G. (2014): "The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008–9 Financial Crisis," *The Quarterly Journal of Economics*, 129(1), 1–59.

DECKER, R. A. (2015): "Collateral Damage: Housing, Entrepreneurship, and Job Creation," Discussion paper, Federal Reserve Board.

DEMYANYK, Y. S., D. HRYSHKO, M. J. LUENGO-PRADO, AND B. E. SØRENSEN (2015): "The Rise and Fall of Consumption in the '00s," Discussion paper, Federal Reserve Bank of Boston.

DUYGAN-BUMP, B., A. LEVKOV, AND J. MONTORIOL-GARRIGA (2015): "Financing Constraints and Unemployment: Evidence from the Great Recession," *Journal of Monetary Economics*, 75, 89–105.

ELUL, R., N. S. SOULELES, S. CHOMSISENGPHET, D. GLENNON, AND R. HUNT (2010): "What "Triggers" Mortgage Default?," *American Economic Review*, 100(2), 490–94.

ELUL, R., AND S. TILSON (2015): "Owner Occupancy Fraud and Mortgage Performance," Discussion paper, Federal Reserve Bank of Philadelphia.

FEDERAL RESERVE BOARD (2016): "Report on the Economic Well-Being of U.S. Households in 2015," Discussion paper.

GILCHRIST, S., M. SIEMER, AND E. ZAKRAJEK (2017): "Real Effects of Credit Booms and Busts: A County-Level Analysis," Discussion paper, Mimeo.

GREENSTONE, M., A. MAS, AND H.-L. NGUYEN (2015): "Do Credit Market Shocks affect the Real Economy? Quasi-Experimental Evidence from the Great Recession and Normal Economic Times," .

GROPP, R., J. KRAINER, AND E. LADERMAN (2014): "Did Consumers Want Less Debt? Consumer Credit Demand versus Supply in the Wake of the 2008-2009 Financial Crisis," .

GUERRIERI, L., AND M. IACOVIELLO (2017): "Collateral Constraints and Macroeconomic Asymmetries," *Journal of Monetary Economics*, 90, 28–49.

GYOURKO, J., A. SAIZ, AND A. SUMMERS (2008): "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index," *Urban Studies*, 45(3), 693–729.

HART, O., AND J. MOORE (1994): "A Theory of Debt Based on the Inalienability of Human Capital," *The Quarterly Journal of Economics*, 109(4), 841–879.

IACOVIELLO, M. (2005): "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, 95(3), 739–764.

JOHNSON, K., K. PENCE, AND D. VINE (2014): "Auto Sales and Credit Supply," Discussion paper, Federal Reserve Board, Finance and Economics Discussion Series, 2014-82.

JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2017): "Credit Supply and the Housing Boom," Discussion paper.

KAPLAN, G., K. MITMAN, AND G. L. VIOLANTE (2016): "Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data," Discussion paper, National Bureau of Economic Research.

——— (2017): "The Housing Boom and Bust: Model Meets Evidence," Discussion paper, National Bureau of Economic Research.

KEYS, B. J., J. TOBACMAN, AND J. WANG (2017): "Rainy Day Credit? Unsecured Credit and Local Employment Shocks," Discussion paper, University of Pennsylvania.

KIYOTAKI, N., AND J. MOORE (1997): "Credit Cycles," *Journal of Political Economy*, 105(2), 211–248.

McCULLY, BRETT A., P. K. M., AND D. J. VINE (2015): "How Much Are Car Purchases Driven by Home Equity Withdrawal? Evidence from Household Surveys," Discussion paper, Federal Reserve Board, Finance and Economics Discussion Series, 2015-106.

MIAN, A., K. RAO, AND A. SUFI (2013): "Household Balance Sheets, Consumption, and the Economic Slump," *The Quarterly Journal of Economics*, pp. 1687–1726.

MIAN, A., AND A. SUFI (2009): "The Consequences of Mortgage Credit Expansion: Evidence from the US Mortgage Default Crisis," *The Quarterly Journal of Economics*, 124(4), 1449–1496.

——— (2011): "House Prices, Home Equity–based Borrowing, and the US Household Leverage Crisis," *The American Economic Review*, 101(5), 2132–2156.

PHILIPPON, T., AND V. MIDRIGAN (2016): "Household Leverage and the Recession," .

SAIZ, A. (2010): "The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics*, 125(3), 1253–1296.

SINAI, T., AND N. S. SOULELES (2005): "Owner-occupied Housing as a Hedge against Rent Risk," *The Quarterly Journal of Economics*, 120(2), 763–789.

STIGLITZ, J. E., AND A. WEISS (1981): "Credit Rationing in Markets with Imperfect Information," *The American Economic Review*, 71(3), 393–410.