

ORIGINAL ARTICLE OPEN ACCESS

Housing Wealth and Consumption: The Role of Heterogeneous Credit Constraints

S. Borağan Aruoba^{1,4} | Ronel Elul² | Şebnem Kalemli-Özcan^{3,4}

¹Department of Economics, University of Maryland, College Park, Maryland, USA | ²Research Department, Federal Reserve Bank of Philadelphia, Philadelphia, Pennsylvania, USA | ³Department of Economics, Brown University, Providence, Rhode Island, USA | ⁴National Bureau of Economic Research, Cambridge, Massachusetts, USA

Correspondence: S. Borağan Aruoba (Aruoba@umd.edu)

Received: 24 June 2026 | **Accepted:** 24 June 2026

Keywords: bank health | consumption | credit constraints | credit supply | Great Recession | household heterogeneity | house prices | housing wealth | individual-level data | mortgage

ABSTRACT

We quantify the role of heterogeneity in households' financial constraints in explaining the link between the decline in aggregate consumption and the decline in house values between 2006 and 2009 using individual-level data. Constraints that are triggered by a decline in house values, and a small fraction of consumers facing particularly severe constraints prior to the decline in house values explain 76% of the aggregate response. Local general equilibrium feedback and a decline in bank credit to consumers make up the remaining 24%. We find no contribution of a pure wealth effect in explaining the consumption decline.

JEL Classification: E21, E32, G21, D14

1 | Introduction

The US economy experienced a large financial crisis together with a large nationwide decline in house prices in 2007–2008. A deep recession with significant declines in consumption, investment, and employment followed. The house price declines across the country showed a large degree of variation: Across US counties the change in house prices between 2006 and 2009 ranged from a decline of around 20% to a slight increase. Moreover, homeowners had different exposures to the changes in house prices, in part based on the kinds of constraints they faced prior to this episode. Although there is an extensive theoretical and empirical literature on the causes and consequences of the crisis, so far, there is little work on the quantitative significance of consumer heterogeneity.

The main challenge in undertaking such an exercise is the availability of individual-level data on house values, wealth, consumption, financial constraints, and individual characteristics.

In this paper, we combine data from two sources based on consumer credit bureau records and mortgage servicing data. These data sets are the FRBNY Consumer Credit Panel/Equifax Data (henceforth CCP) and the Equifax Credit Risk Insight Servicing and ICE, McDash Data (henceforth CRISM).¹

Mian and Sufi (2009), Mian and Sufi (2011), and Mian et al. (2013) have documented that an increase in household leverage predicts the subsequent crisis, de-leveraging, and consumption decline and thus point to the importance of credit constraints. Using our individual-level data, we show how heterogeneous this effect is across consumers, and how this relates to their financial constraints. A small set of papers including Adelino et al. (2016), Adelino et al. (2018), Albanesi et al. (2022), and Foote et al. (2020) uses individual-level data covering the period before the financial crisis and its aftermath, and is very closely related to our analysis. These papers show that, once individual-level data are used, conclusions based on aggregate (e.g., ZIP-level) data

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 Federal Reserve Bank of Philadelphia and The Author(s). *International Economic Review* published by Wiley Periodicals LLC on behalf of The Economics Department of the University of Pennsylvania and the Osaka University Institute of Social and Economic Research Association.




FIGURE 1 | House prices and consumption: Channels.

do not necessarily continue to hold. For example, Albanesi et al. (2022), who also use CRISM, and Adelino et al. (2018) show that although ZIP codes with many subprime borrowers appear to have led to the rise of mortgage defaults during the downturn, it was in fact the non-subprime borrowers **within** these ZIP codes who defaulted at higher rates than historical patterns. However, unlike our paper, these papers do not explore the link between house price declines and consumption.

Figure 1 shows all of the possible channels that we identify as linking a decline in house prices to lower consumption, via three players in the same locality: consumers, banks, and firms. First, on the household side, as a result of declining house prices, there will be both a pure wealth effect, denoted with the arrow “pure household wealth,” and a balance sheet effect (if housing is an important source of collateral for borrowing), denoted with the arrow “household financial constraints.”² It is important to define what we mean by pure wealth effect and how it differs from the term as used by some other papers in the literature. We, as does Aladangady (2017), define the pure wealth effect to capture the change in the consumers’ behavior due to a change in their wealth without any change in the severity of their financial constraints. By contrast, Kaplan et al. (2020) investigate the sources of the decline in house prices and its contribution to decline in consumption using their model, and find that half of the link between the two can be traced to what they term wealth effects. However, they use this term to refer to what is also sometimes called the “endowment effect,” where the initial wealth of the consumer explains the decline in consumption. Thus, while on the surface they refer to similar things, our pure wealth effect and the wealth effect in Kaplan et al. (2020) are very different concepts.

Next, there is the effect of house price declines on bank health. As documented by Rosen (2011), in the period preceding the crisis about 30% of mortgages were originated by local banks. As such, there are at least two ways local banks, defined as depository

institutions with branches in the local area, are exposed to the real estate market and can be adversely affected by a decline in house prices. First, they hold a fraction of the mortgages they originate on their balance sheets and a decline in house prices would shrink the value of these assets. In addition, even for loans that instead end up in securitized pools, they collect fees for originating these mortgages, and a decline in house prices would reduce the housing activity and thus revenues from this part of their business. These effects could then lead local banks to cut credit supply to both households and firms.³

A similar channel can also occur with a direct shock to firm balance sheets instead of bank balance sheets, when firms’ owners use their own housing wealth as collateral to obtain loans to invest and to produce. Bahaj et al. (2020) provides direct evidence for this channel for the United Kingdom and Schmalz et al. (2017) show that an increase in one’s house value as collateral increases one’s probability of becoming an entrepreneur. We are not able to directly identify this channel, since we do not have information on firms’ or their owners’ real estate wealth. If wages are sticky in the short run, this, and other channels, including a decline in “local demand” by consumers and a decline in credit by banks, will lead to a decline in employment by local firms. This will feed back to lower consumption because of a “local general equilibrium feedback effect,” as shown by the bottom arrow.

Using individual-level data, we identify the importance of all these four channels through which consumers change their consumption (shown in bold in Figure 1): pure household wealth effect, household financial constraints, household credit supply, and local general equilibrium feedback. It is important to emphasize that individual data and the heterogeneity the data provide are key for distinguishing the pure wealth effect from financial constraints, since all homeowners in an area would experience a similar pure wealth effect, but they will differ in the severity of their financial constraints.

We follow Mian et al. (2013) and proxy consumption expenditures with information on auto purchases. While they use ZIP code-level new car registrations as one of their measures of consumption, we take advantage of our individual-level credit-bureau data and create a binary variable representing the origination of an auto loan in 2009 by individuals. In addition, as a novel contribution to the literature, by leveraging our individual-level data, we are able to use an estimate of the household-level change in house values in our analysis, instead of simply relying on aggregate measures of changes in local house price indices. Using detailed information on mortgage and borrower characteristics in our data, we also create various measures of credit constraints. We distinguish between two types: ex ante and ex post constraints. Ex ante constraints are those that existed prior to 2006 and likely affected the choices the consumers made prior to 2006, including their mortgage type. We use measures of the creditworthiness of the consumer, the loan-to-value ratio of their mortgage and their type of mortgage as indicators of the ex ante constraints they face prior to 2006. Ex post constraints are those that are triggered by the decline in house values. We focus on a particular measure we term *Bad Mortgage*, in which we identify homeowners who are seriously delinquent in their mortgage payments during this period. We show below that this is triggered by a decline in house prices and it directly affects consumers' ability to originate an auto loan in 2009.

Our main specification is one with auto loan originations in 2009 as the dependent variable and the house value changes between 2006 and 2009 as the key explanatory variable. To account for the fact that house values are not exogenous, we use an instrumental variables (IVs) approach, where we instrument the change in house prices with standard instruments in the literature on housing supply elasticity. We find that 76% of the total response of consumption in 2009 to changes in house values between 2006 and 2009 can be attributed to financial constraints. Of this, the ex post constraint proxy *Bad Mortgage* is responsible for 23%, and ex ante constraints are responsible for the remaining 53%. A small fraction of consumers that have particularly severe ex ante constraints (e.g., those with second mortgages that have high interest rates) are responsible for a large fraction of this 53%. Local general equilibrium (13%) and household credit supply (11%) constitute the remainder of the response. Finally, we turn to the identification of the pure wealth effect. To do so, we focus on consumers that are unlikely to face any credit constraints—creditworthy homeowners with very low loan-to-value ratios—and show that these consumers do not react to changes in their house value.⁴

Turning to the literature, the work of Aladangady (2017) is closest to our paper. To the best of our knowledge, this is the only other paper using individual-level data (restricted-access geographical files from the Consumer Expenditure Survey, in his case) to investigate the consumption response to a change in house prices, although he focuses on the period before the 2007–2009 Great Recession. He finds results similar to ours in terms of importance of household-level financial constraints, emphasizing the importance of collateral effects as opposed to a pure wealth effect. In addition to the time period of analysis, there are two main differences between our paper and his. First, we can account for general equilibrium effects and the effect of credit supply. Second, we have much larger and detailed individual-

level data that help us distinguish both ex ante and ex post borrowing constraints.

Our results are consistent with much of the broader housing wealth and permanent income literature. Many papers generally estimate a small pure wealth effect: five cents out of one dollar in Pistaferri (2016) with aggregate data, and two cents out of one dollar using the PSID in Carroll et al. (2011). Vestman et al. (2026) estimate a pure wealth effect of only 0.12 cents out of one dollar using a quasi-experiment in Sweden. 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 (2010). In the context of a life-cycle model, if homeowners are 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). Garriga and Hedlund (2020) present a rich incomplete-markets macro-housing model where consumers with larger mortgages (and illiquid wealth) respond much more to house price changes, in line with our empirical results. Guren et al. (2021) use historical data and show that responses to changes in housing wealth are consistent with a standard life-cycle model with borrowing constraints, uninsurable income risk, illiquid housing, and long-term mortgages. They also find that housing wealth effects were not particularly large in the 2000s. Accounting for the *heterogeneity* in financial constraints, as we do, seems to be key in explaining the large aggregate response of consumption.

We proceed as follows. Section 2 discusses the data in detail. Section 3 begins with our replication of the ZIP code-level results in Mian et al. (2013), and then proceeds to our baseline individual-level results. In Section 4, we decompose the total effect of changes in house values on consumption into the various channels. Section 5 digs deeper into the effects of household-level financial frictions and Section 6 focuses on identifying the pure wealth effect for the households. Finally, Section 7 concludes.

2 | Data

This section introduces our data—first the individual-level data followed by aggregate (ZIP-, county-, and MSA-level) data. Further details may be found in Online Appendix A.

2.1 | Individual-Level Data Sources

Our main data set is CCP, a quarterly database of consumer credit bureau records for a random 5% anonymized sample of consumers with a bureau record. From the CCP we can observe total balances and aggregate delinquency status on a variety of consumer credit obligations such as mortgages, auto loans and credit cards, and Equifax Risk Score (henceforth Risk Score), as well as some loan-level information on first and second mortgages. As we explain in Section 2.2, our consumption proxy is computed using auto loan origination information included in credit bureau records. We are also able to calculate the age of consumers based on the birth year that is provided in the CCP.

For the consumers in CCP who have a mortgage in 2006, we also obtain information from a second data set, CRISM, which contains more detailed information on residential first mortgages

from loan servicing data.⁵ CRISM captures approximately two-thirds of all mortgage originations during this time period, and it gives more detailed information on the borrower's mortgages than found in CCP itself, most notably: the appraised value of the property; interest rate; other characteristics such as whether the mortgage is fixed or adjustable rate; and monthly mortgage performance information. Being able to observe the appraised value of homes is a major strength of our data set, as it allows us to compute the dollar change in the value of the home of each individual in our data set between 2006 and 2009. To do so, we begin with the appraised value that is shown in the active mortgage as of December 2006. Then using the percentage changes in the available house price indices we update the value of the home from the date of the appraisal to both December 2006 and December 2009.⁶ For two homeowners in the same ZIP code, the *percentage* change in the value of their house will be the same but the *dollar* change can differ. We restrict attention to CRISM borrowers, who also appear in CCP, are homeowners with a single first-mortgage in December 2006 and December 2009 (though not necessarily the same one, which allows for refinancing), and who have not moved between December 2006 and December 2009; we are left with about 345,000 unique individuals.⁷

We now discuss the representativeness of our sample. The focus of this paper is the impact of changes in the value of one's own home on one's own consumption. So it is natural to focus on the population of homeowners. Aladangady (2017) studies renters (i.e., non-homeowners), and finds a negligible response to local house price changes.

Consumers with a single mortgage represent 79% of all mortgage-holders in the credit bureau data. We center our main analysis on these borrowers in order to be able to construct a clean sample linking house price changes to a borrower. However, we also extended our data set and analysis to the further 14% with two mortgages (thus covering 93% of all mortgage-holders in the data). The results, which can be found in Online Appendix C.2, are virtually identical to those in the body of the paper.

From the 2007 SCF, 71% of all homeowners had an active mortgage in this time period.⁸ These mortgage-holders, who make up our sample, are almost certain to be represented in our credit bureau data. For the remaining homeowners who have no mortgage, if they nevertheless appear in the credit bureau data, our panel results for "free-and-clear homeowners" in Section 6.2 suggest a zero impact of house price changes on this population.

Upper bounds for the fraction of homeowners who would be missed by our analysis because they do not appear in the credit bureau data can be obtained from other data sets. From the PSID, only 7.5% of homeowners who reported purchasing a vehicle in 2008 report no debts of any kind, and thus might not appear in the credit bureau data.⁹ Similarly, the Consumer Financial Protection Bureau finds that less than 6% of the adult population in the United States does not have a credit bureau record.¹⁰

In contrast to other, publicly available data sets such as the PSID, the advantage of our credit bureau data is that the data allow us to explore the role of credit constraints, and in particular distinguish between what we term *ex ante* and *ex post* credit constraints. In order to do so, it is important to control for risk scores, which

are absent from the PSID. In addition, as Gerardi et al. (2018) point out, house values in the PSID (which are self-reported) are not measured particularly accurately, particularly for those with negative equity. Finally, the PSID sample size is considerably smaller than ours, which would make using geographic variation (such as for local economic conditions and credit supply) difficult.

We classify the households in our data in three different ways, all using *ex ante* criteria that are observed as of 2006Q4, the start of our analysis. First, we label households with a Risk Score of 700 or higher as **Prime** and the others as **Non-Prime**.¹¹ Next, we use the updated first-lien LTV ratio as of December 2006 to create four groups: those with an LTV ratio of less than 25%, 25% or higher and below 50%, 50% or higher and below 80%, and greater than or equal to 80%. We refer to these groups as **LTV0**, **LTV1**, **LTV2**, and **LTV3**, respectively. Finally, using the more detailed information about types of mortgages the households have in CRISM, we create five categories: those with a **Fixed-Rate First Mortgage** (and no second lien); those with an **Adjustable-Rate Mortgage (ARM)** that has an initial fixed-rate duration of **less than 5 years**, and separately those that have a duration **greater than or equal to 5 years** (in both cases with no second lien); those with a **Closed-End Second Mortgage** (and any first mortgage); and those with a **Home Equity Line of Credit (HELOC)** (and any first mortgage), all as of December 2006.¹² Table A-2 in the Online Appendix shows the distribution of the households across these categories. These categories will be useful in quantifying the effect of *ex ante* financial constraints.

In order to measure the effect of *ex post* financial constraints, we create a dummy variable called **Bad Mortgage** by identifying households that are seriously delinquent (at least 90 days behind) in any mortgage payment during 2007–2009. Roughly 7% of households had a bad mortgage in this period.

2.2 | Construction of the Consumption Proxy

We proxy for individual-level consumption expenditures by **Auto Loan Originations**. The same measure has been used in Di Maggio et al. (2017), among others. The Auto Loan Tradeline Panel of CCP provides data on auto loans and leases, which include the month of origination. This measure tracks the individual-level incidence of auto loan originations in other sources very well: For example, we find that 10.1% of all consumers originate an auto loan in 2008 in the CCP, whereas from the Panel Study of Income Dynamics (PSID) the origination rate in that year was 10.8%.¹³

The aggregated auto loan originations from CCP also track alternative aggregate vehicle expenditure measures very well, all of which exhibit similar cyclicity. The red line in Figure 2 shows the aggregated Auto Loan Originations and the blue line is Total Vehicle Purchases from the Bureau of Economic Analysis (BEA), both normalized to 100 in 2006. Our measure tracks the BEA measure almost perfectly over time, both in its magnitude, and also its turning points. The figure also shows Durable Good Expenditures from the BEA in green, which contains much more than Vehicle Purchases. For overall durable goods expenditures the peak occurs a year later and the trough is not as deep.




FIGURE 2 | BEA vs. CCP: Consumption expenditures.

As Mian and Sufi (2016) also emphasize, individual-level consumption data, combined with detailed asset and liability information, is hard to come by for the United States. This motivates our use of auto loan originations, which are available in the credit bureau data set we use, as a proxy for individual consumption. By contrast, Mian et al. (2013) use new car registrations at the ZIP code level. Relative to Mian et al. (2013), our measure has some advantages. First and foremost, it is at the individual level, and since it is obtained from credit bureau data, we are also able to exploit other individual-level characteristics, rather than basing our analysis solely on aggregate measures. Second, by focusing on consumer credit records, we are able to isolate auto purchases by consumers, as opposed to businesses. Finally, our measure also captures purchases of used cars, and not just new ones.

2.2.1 | Wealth Imputation Using Survey of Consumer Finances (SCF)

While the CRISM data allow us to capture housing assets and liabilities, households' consumption decisions may also be influenced by their non-housing net worth, defined as assets minus liabilities excluding the value of the house and the loans that are secured by the house. This measure, in turn, may be at least partially correlated with the value of their house. As such, the omission of non-housing net worth may bias our results. Unfortunately, while it has extensive coverage of households' liabilities, credit bureau data do not contain information on their assets. To overcome this problem we use the SCF to impute non-housing assets of consumers in our data.¹⁴ The details are provided in Online Appendix A.2. To prevent measurement and imputation error from affecting our results, we assign households to five equal-sized categories based on their imputed non-housing wealth and use these as controls in our models.

2.3 | Aggregate Data

In addition to individual-level variables, we use selected ZIP code- and county-level variables as controls. Crucially, these variables will also help us identify the portions of the change in consumption that are due to local general equilibrium effect and changes in credit supply. See Online Appendix A for further details.

To proxy for local economic conditions, we use **Change in Unemployment Rate** from December 2006 to December 2009 at the county level, published by the Bureau of Labor Statistics. To capture the effect of changes in banks' credit supply on consumption, we follow the methodology in Gilchrist et al. (2023) and Greenstone et al. (2020) and create a county-level measure of **Credit Supply** shock that cumulates changes in credit supply that are only due to banks' fundamentals. As an alternative to this variable, we create a variable we term **Bank Health**. This is a county-level version of the bank health indicator provided by Chodorow-Reich (2014), who uses, among others, a bank-level measure of the fraction of the syndication portfolio in which Lehman Brothers had a lead role.

Finally, as an additional control, we also use a ZIP code-level measure of **Auto Sales in 2003**, which we compute by aggregating our individual-level auto loan origination variable. This is meant to capture permanent geographical differences in the prevalence of car ownership, holding other things constant (e.g., between Manhattan and Los Angeles). We choose 2003 for this because it is sufficiently far from the 2006–2009 period we consider to represent a baseline.

2.4 | Local Housing Supply Instruments

Most of our analysis is undertaken using an IV approach in order to address the endogeneity of individual-level house values and omitted variables. To this end, following Mian and Sufi (2009), Mian and Sufi (2010), Mian and Sufi (2011), and Mian et al. (2013), we use two MSA-level instruments, which are intended to capture the elasticity of housing supply, and therefore the response of house prices to demand shocks.¹⁵ First is the **Share of Land that is Unavailable for Real Estate Development**, from Saiz (2010), which measures the share of land within a 50-km radius of the MSA centroid that cannot be developed based on geographic features. In addition, we use the MSA-level **Wharton Residential Land Use Regulation Index (WRLURI)** developed using a survey by Gyourko et al. (2008). This is a standardized measure across all municipalities, and lower values can be thought of as reflecting the adoption of more laissez-faire policies toward real estate development.

3 | From ZIP Code to Individual Level

3.1 | Mian et al.'s (2013) ZIP Code-Level Results

We begin by linking our analysis to the ZIP code-level analysis in Mian et al. (2013). Let $R_{z,t}$ denote the number of new car registrations in ZIP z , S_t the aggregate dollar value of new car sales in the United States, and $h_{z,t}$ the number of households in ZIP z , all in year t . We can define $C_{z,t}$, one of the consumption measures in Mian et al. (2013), as $C_{z,t} \equiv S_t \frac{R_{z,t}}{h_{z,t} \sum_{z'} R_{z',t}}$, which simply allocates S_t to each ZIP using the share of new car registrations in that ZIP out of the whole United States and then normalizes by the number of households in that ZIP to get a per-capita measure. Let $\Delta H P_z^{2006-2009}$ denote the average dollar change in house prices in the ZIP between 2006 and 2009.

Given these definitions, one of the headline results at the ZIP code level in Mian et al. (2013), as shown in column 5 of their table V, is

$$C_{z,2009} - C_{z,2006} = \alpha^{\text{MRS}} + \frac{0.018\Delta HP_z^{2006-2009}}{(0.001)} + \varepsilon_z^{\text{MRS}} \text{ with } R^2 = 0.153, \text{ and } N = 6263.$$

This shows a highly significant effect of the change in house prices on change in consumption: an \$18 decline for every \$1000 drop in house prices, which can be translated to a dollar change of $\Delta C^{\text{MRS}} = -\855 using the average change in house prices from 2006 to 2009 they report in their table I. Note that this model is estimated using an IV strategy with one of the two instruments we introduced in Section 2.4, namely Land Unavailable.

3.2 | Individual-Level Results

The number reported in the previous section from the analysis of Mian et al. (2013), an \$18 decline in consumption for each \$1000 decline in house prices at the ZIP code level, can be thought of as the *total effect* of change in house prices on consumption. One of our key goals in this paper is to decompose this total effect into a *direct* effect and a number of *indirect* effects, which we turn to in Section 4. Before doing so, in this section, we provide our own baseline estimate for the *total* effect from our individual-level data. This requires five noteworthy deviations relative to Mian et al.'s (2013) analysis in the previous section, in addition to the obvious one that involves change in the unit of observation. First, our left-hand side variable, y_i is a dummy variable that shows if the individual i has originated an auto loan in 2009. Second, we add some individual- and ZIP code-level controls that were introduced in Section 2 to the specification. Third, instead of the ZIP code-level house *price* changes we are now able to use individual-level dollar change in house *values* as the key variable of interest. This is denoted as ΔHV_i and is measured in \$100,000. Fourth, we use both housing supply elasticity measures as instruments. Fifth, we allow for heterogeneous coefficients for all variables in the specification except for the Change in House Value, because our analysis reported later uncovers considerable heterogeneity in how each of these controls affect the household's consumption decision. We consider heterogeneity in two dimensions: four categories based on the consumer's LTV ratio and two categories based on their Risk Score, both as of 2006.¹⁶ In this section and next, we are interested in the *average* effect of the Change in House Value and as such we do not consider heterogeneity in that dimension. We turn to the heterogeneity in the response to Change in House Value in Section 5.2.

Using i to denote an individual and $z(i)$ the ZIP code of their residence in 2006 and 2009 (since we focus on non-movers), the second stage equation we estimate is given by

$$y_i = \alpha_{j(i)}^0 + \frac{0.0176\Delta HV_i}{(0.0022)} + \alpha_{j(i)}^1 age_i + \alpha_{j(i)}^2 age_i^2 + \alpha_{j(i)l(i)}^3 A_{l(i)} + \alpha_{j(i)k(i)}^4 W_{k(i)} + \alpha_{j(i)}^5 C_{z(i)} + \varepsilon_i. \quad (1)$$

Throughout the paper, for a generic variable X , the notation $\mu_{j(i)}X_i$ is a shorthand for $\sum_{j=1}^J \mu_j D_i^j X_i$ where $\{D_i^j\}$ is a set of dummy variables for J categories, with D_i^j equal to 1 if person i is in group j and 0 otherwise.¹⁷ In addition to an age polynomial, we use two other controls at the individual level. First, we compute the number of auto loans the individual originated in the period 2004–2006. We summarize this information in four dummy variables $\{A_{l_i}\}$ with $l = 1, 2, 3, 4$ where we group the individuals with 0, 1, 2, and 3 or more loan originations, respectively. These dummy variables measure how likely it is for an individual to originate an auto loan, especially given that buying a car is lumpy, and are meant to capture unobserved differences across individuals. Second, we use the four categories of non-housing wealth given by $\{W_k\}$ with $k = 1, 2, 3, 4$ introduced in Section 2.2.1. The only local aggregate control (for now) is C_z , which is the ZIP code-level auto sales in 2003 computed using our loan origination data. The goal in estimating this equation is to control for most individual- and ZIP code-specific factors that affect new loan originations in 2009, and determine how important the change in house value of the consumer between 2006 and 2009 is in explaining the remainder. As such, in contrast to the estimation in Mian et al. (2013), the variation exploited here is across individuals with similar characteristics. In all estimations standard errors are clustered at the ZIP code level.

The first stage naturally has the same controls as the second stage, and uses the two housing supply instruments introduced in Section 2.4. In the first stage, as in the second stage, all variables including the instruments are allowed to have heterogeneous coefficients based on the eight categories of LTV and Risk Score combinations. Crucially, this means that even though the housing supply instruments are measured at the MSA level, due to their interaction with the eight individual-level variables, we have 16 instruments and thus the variation created in the first stage is at the individual level. In the interest of space we do not report the first stage results. Both instruments enter the first stage with negative and highly statistically significant coefficients for all categories except one, which has a statistically insignificant coefficient for one of the instruments. The instruments capture housing supply (in)elasticity: In areas where building regulations are more restrictive or in areas where little land is available to develop, housing supply will be more inelastic; larger values of the two instruments indicate more inelastic supply. This means that in response to a demand shock, we expect a larger price reaction in such areas, since supply cannot respond as much. The period we are considering, between 2006 and 2009, can be thought of as a period with a large negative aggregate housing demand shock. Thus prices should fall by more in areas with inelastic housing supply, which is indeed what the negative coefficients in the first stage indicate.¹⁸ The first-stage F -statistic is 159.7 and is well beyond the critical values provided by Stock and Yogo (2005). Thus, tests of weak identification and, separately, tests of under-identification are all easily rejected. As is the case with most IV-based studies, our analysis is limited to identifying a local average treatment effect (Imbens and Angrist 1994), and therefore may not generalize to other identification strategies.

Turning to the controls' effects on consumption, we find that households that originated more auto loans in the period

2004–2006 or those that have higher non-housing net worth in 2006 are more likely to originate auto loans in 2009. Similarly, households that live in ZIP codes that had a large number of auto loan originations in 2003 tend to have more originations in 2009. Finally, age polynomials show that, in general, auto loan originations fall monotonically with age.

Finally, the estimated coefficient for ΔHV_i in (1) shows that for every \$100,000 decline in house values, the probability of originating an auto loan falls by 1.76 percentage points. Using the average of individual-level house value changes in our data, which is \$75,109, the marginal effect at the mean is a decline of 1.32 percentage points, compared to the fraction of households that originated an auto loan in 2009, which is 13.5%. Following the same approach as in Section 3.1, this translates into a dollar response of $\Delta C = -\$1,225$.¹⁹

To sum up, in their ZIP code-level analysis Mian et al. (2013) find a decline of \$855 in per-capita purchase of autos in response to the average decline in house prices. Our individual-level results show a decline of \$1225, which replicates the large aggregate response of consumption to housing wealth changes that Mian et al. (2013) find. In the remainder of the paper our goal will be to demonstrate that a large fraction of this *total* consumption response is due to heterogeneity in credit constraints across consumers. Some consumers do not react at all to the change in housing wealth, and some have a reaction that is many times larger than the average response; this heterogeneity will be accounted for by various credit constraints.

4 | Main Results

In this section, we first present our full model, which extends the individual-level model in (1) to include additional aggregate controls. This model in conjunction with (1) is used to decompose the total effect we presented in Section 3.2.

4.1 | Full Model

We begin our analysis by expanding (1) with three additional controls:

$$y_i = \alpha_{j(i)}^0 + \beta^1 \Delta HV_i + \beta_{j(i)}^2 \Delta U_{c(i)} + \beta_{j(i)}^3 CS_{c(i)} + \beta_{j(i)}^4 BM_i + \alpha_{j(i)}^1 age_i + \alpha_{j(i)}^2 age_i^2 + \alpha_{j(i)k(i)}^3 A_{k(i)} + \alpha_{j(i)k(i)}^4 W_{k(i)} + \alpha_{j(i)}^5 C_{z(i)} + \varepsilon_i. \quad (2)$$

Here $c(i)$ denotes the county of residence for the individual i in 2006 and 2009, and ΔU and CS are county-level measures of the change in unemployment rate and the change in credit supply. BM is the binary Bad Mortgage measure, which captures whether or not the consumer has had difficulty in paying their mortgage from 2006 to 2009. We defer the discussion of why this variable impacts consumption to Section 5.1, taking it for granted for now.

The estimates of (2) are presented in the first panel of Table 1. We present the results in four columns where the first column replicates the results from the estimation of (1) in the previous section, and each subsequent column adds one additional variable, arriving at the full specification in Column 4. The value

shown for ΔHV is the estimate of β^1 , while for the other three variables we combine the estimates of $\{\beta_{j(i)}^2, \beta_{j(i)}^3, \beta_{j(i)}^4\}$ using the sample weights for each of the j groups to get the marginal effect for each variable shown on the table, adjusting the standard errors accordingly. The first and most important result to highlight is that once the additional variables are included, the importance of ΔHV falls by 40%. This indicates that a sizable fraction of the *total* effect of Change in House Value on consumption was in fact due to its *indirect* effect via other variables, indicating the importance of these other channels. We pick up this “omitted variable bias” (OVB) logic more formally below and provide a decomposition. The second result to highlight from the first panel is that the remaining variables are mostly significant and the coefficients are quite similar as we go across columns, indicating the absence of much correlation among these three variables. ΔU has a negative sign—an increase in unemployment in the county reduces consumption—and Credit Supply has a positive sign—a decline in credit supply in the county reduces consumption. Consumers who experience a Bad Mortgage also reduce their consumption.²⁰

The second panel evaluates the marginal effects in the first panel at the means of the respective variables for ΔHV and ΔU , which are a house value decline of \$75,109 and an unemployment rate change of -5.5 percentage points, respectively. For Credit Supply, since it is a flow variable in levels and its mean is roughly zero, we consider the marginal effect of moving one standard deviation below the mean. For Bad Mortgage we simply report the decline in probability of originating an auto loan when comparing an individual with no Bad Mortgage to one with Bad Mortgage; since it is a dummy variable this is simply the coefficient in the first panel. Finally in the last panel we convert these marginal effects to dollar values via the same method used in Section 3.1.

The important conclusion from the third panel is that the other three variables are also important contributors in their own right to the decline in consumption: Considering Column 4, the average increase in unemployment rate reduces consumption by \$756, which is about the same as the *direct* effect of the Change in House Value in Column 4. A one standard deviation decline in Credit Supply reduces consumption by \$254, and a borrower with a Bad Mortgage has their consumption reduced by \$7674.

4.2 | Decomposing the Total Effect

Using a methodology adapted from the derivations of the OVB, we decompose the *total* effect of the Change in House Value on consumption into its various channels, as identified in Figure 1. For the purposes of this decomposition, we consider the model in (2) as the “true” model and the model in (1) as the misspecified model, as it omits three variables that may be relevant.²¹ Our goal is to interpret the coefficient on ΔHV in (1) as the *total* effect of the Change in House Values on consumption and decompose it into *indirect* effects that go via the three omitted variables and the remaining *direct* effect. For this, we use the well-known derivations for the OVB with three modifications. First, these derivations are typically used for OLS and we adapt them to our IV approach. Second, since we allow for heterogeneous effects for each of the three “omitted” variables, the derivations need to

TABLE 1 | Main results.

Marginal effects				
	(1)	(2)	(3)	(4)
Change in House Value	0.0176*** (0.0022)	0.0197*** (0.0034)	0.0167*** (0.0039)	0.0103*** (0.0039)
Change in Unemployment Rate	—	−0.0013* (0.0007)	−0.0014* (0.0007)	−0.0015** (0.0007)
Credit Supply	—	—	0.0263** (0.0106)	0.0312*** (0.0105)
Bad Mortgage	—	—	—	−0.0829*** (0.0031)
Marginal effects (in p.p.)				
ΔHV (average: −\$75,109)	−1.32	−1.48	−1.26	−0.77
ΔU (average: 5.5 p.p.)	—	−0.69	−0.75	−0.82
Credit Supply (−1 s.d.)	—	—	−0.23	−0.27
Bad Mortgage (= 1)	—	—	—	−8.29
Marginal effects (in dollars)				
ΔHV (average: −\$75,109)	−\$1225	−\$1368	−\$1163	−\$715
ΔU (average: 5.5 p.p.)	—	−\$636	−\$694	−\$756
Credit Supply (−1 s.d.)	—	—	−\$214	−\$254
Bad Mortgage (= 1)	—	—	—	−\$7674

Note: All equations are estimated via instrumental variables using the two housing supply instruments interacted by an eight-dimensional categorical variable and have a sample size of $N = 345,033$. The mean of the dependent variable (Auto Loan Origination in 2009) is 0.135. First stage F -statistics are 160, 101, 90, and 88, respectively. The first panel reports marginal effects for respective variables. For ΔHV this is the estimated coefficient. For the other variables it is computed as the weighted average of the coefficients of the variable interacted with the eight-dimensional categorical variables for Risk Score and LTV, where the standard errors are appropriately computed. The second panel converts these marginal effects to percentage point units, by multiplying with the average of ΔHV and ΔU . For Credit Supply we look at the effect of a one standard deviation decline. For Bad Mortgage we show the effect of having the binary variable being equal to unity. The third panel converts the numbers in the second panel to dollar values by using \$926 for each 1% decline. This is obtained by combining the average probability of originating an auto loan in 2009 and the average price of a car in 2009, which are 0.135 and \$12,518, respectively.

TABLE 2 | Decomposition of the total effect into channels.

	Local general equilibrium	Household credit supply	Pure wealth and other constraints	Bad mortgage
Baseline	13%	11%	53%	23%
Use Bank Health	14%	9%	55%	23%
Probit-IV	9%	9%	57%	25%

Note: This table reports the decomposition of the total effect of Change in House Value on auto loan originations reported in Column 1 of Table 1 into four channels using the methodology presented in Section 4 and in Online Appendix B. Numbers in each row may not add to 100% due to rounding.

be generalized for the total marginal effect. Third, we have some additional control variables that are present in both (1) and (2). The methodology is explained in detail in Online Appendix B.

Table 2 reports the results for the decomposition. This decomposition takes the total effect and decomposes it into a direct effect and three indirect effects, one for each of the additional variables: Change in the Unemployment Rate, Credit Supply, and Bad Mortgage. In doing so, it is useful to link these back the channels in Figure 1. We consider the share of the total

effect that is due to the Change in the Unemployment Rate a measure of the Local General Equilibrium channel, and the share due to Credit Supply a measure of the Household Credit Supply channel. Table 2 shows that in our baseline specification these channels receive 13% and 11% shares, respectively, leaving 76% for the remaining two channels in Figure 1: Pure Household Wealth Effect and Household Financial Constraints. Bad Mortgage is a measure of a key financial constraint and it receives a share of 23%. Thus, the remaining 53% out of the 76% is then due to the Pure Wealth Effect and other financial constraints. We turn to the

TABLE 3 | Robustness of main results (marginal effects, in percentage points).

Main results		
ΔHV (average: $-\$75,109$)	-1.32^{***}	-0.77^{***}
ΔU (average: 5.5 p.p.)	—	-0.82^{**}
Credit Supply (-1 s.d.)	—	-0.27^{***}
Bad Mortgage (= 1)	—	-8.29^{***}
Bank health instead of credit supply		
ΔHV (average: $-\$75,109$)	-1.32^{***}	-0.81^{***}
ΔU (average: 5.5 p.p.)	—	-0.98^{**}
Bank Health (+1 s.d.)	—	-0.25^{**}
Bad Mortgage (= 1)	—	-8.28^{***}
Probit–IV		
ΔHV (average: $-\$75,109$)	-0.31^*	0.09
ΔU (average: 5.5 p.p.)	—	-1.61^{***}
Credit Supply (-1 s.d.)	—	-0.41^{***}
Bad Mortgage (= 1)	—	-8.44^{***}

Note: The first panel repeats the results in Columns 1 and 4 in the second panel of Table 1. The other panels report the summary results analogous to the first panel with the change described in the title. The statistical significance signs follow from the point estimates. See Table A-5. See the notes to Table 1.

detailed analysis of the Financial Constraints channel in Section 5 and of the Pure Wealth channel in Section 6.

4.3 | Robustness of the Decomposition of the Aggregate Effect

We consider two variations to investigate the robustness of our results presented thus far. First, one may be concerned about the identification of credit supply shocks and their exogeneity. To address this, we replace Credit Supply with Bank Health, which allocates the 2008 “Lehman shock” of Chodorow-Reich (2014) to US counties based on the number of branches of affected banks in each county. Second, we use a Probit instead of a linear probability model in the second stage of our IV. In Table 3, we present the marginal effects in percentage points for the baseline results and these two variations. Full results are relegated to Online Appendix C.3. These results confirm the robustness of our main results. Considering the specification with Bank Health, the marginal effect for ΔHV is roughly unchanged, while that for ΔU is modestly larger. Probit results show much smaller total effect of Change in House Value (-0.31 vs. -1.32), and a much smaller direct effect (an insignificant 0.09 vs. -0.77). The marginal effects of other variables are about twice as large.

The decomposition results with these two versions in Table 2 are very similar to the baseline results as well. Using Bank Health instead of Credit Supply yields decomposition results within 2 percentage points of the baseline. With Probit, the effect of the Local General Equilibrium and Household Credit Supply

channels get a lower combined share, with Pure Wealth and Other Constraints showing a 4 percentage point increase.

4.4 | Panel Results With House Prices

Unfortunately the house value data in CRISM only allow us to conduct an analysis between two dates. As such, we must use controls to capture individual-specific heterogeneity, and are not able to include individual fixed effects in our analysis. In this section, we use panel data from CCP covering the period 2002–2010, which allows us to capture unobserved individual heterogeneity. This comes at the cost of no longer being able to use house value changes at the individual level. We must instead rely on house price changes at the ZIP code level; we also cannot control for non-housing net worth. We focus our analysis on consumers with a mortgage and who do not move from the previous year into the current year, in order to investigate the robustness of our results thus far. In Section 6.2, we turn to homeowners without a mortgage in order to measure the pure wealth effect.

We conduct our analysis, splitting the sample into “Prime” and “Non-Prime” borrowers using the Risk Score of the consumers in our analysis. Similar to our analysis above, we define a consumer as **Prime** in year t if their Risk Score in Q4 of year $t - 1$ is greater than 700 and **Non-Prime** otherwise. We estimate the following linear probability model with an IV strategy separately for each of the two Risk Score categories j :

$$y_{it} = \beta_j^{1+} \Delta H P_{z(it)t}^+ + \beta_j^{1-} \Delta H P_{z(it)t}^- + \beta_j^2 \Delta U_{c(it)t} + \beta_j^3 CS_{c(it)t} + \beta_j^4 BM_{it} + \alpha_i^j + \alpha_t^j + \alpha_{z(it)}^j + \varepsilon_{it} \quad (3)$$

for individuals with $j(it) = j$, where $j(it)$, $z(it)$, and $c(it)$ denote the category, ZIP code, and county the individual i belongs to in year t . The dependent variable is whether or not the consumer originates an auto loan in the current year. We use two separate house price variables, one for increases ($\Delta H P_{zt}^+$) and one for decreases ($\Delta H P_{zt}^-$), both measured as the percentage change in the house price of the ZIP code the individual lives in. These variables are meant to capture the possible asymmetric effect of house price changes on consumption.²² In addition to the two county-level controls, change in the unemployment rate (ΔU) and credit supply (CS), we also include individual, time, and ZIP code fixed effects. We also add an indicator of having a Bad Mortgage in the previous year (BM) in some specifications. Our IV strategy also adapts to the panel structure. Following Aladangady (2017), we interact the two housing supply instruments used throughout the paper with a national variable that captures shifts in housing demand, and use these to instrument for the two house price growth variables. Online Appendix A.5 provides further detail.

Ideally, one would want to conduct the same decomposition into channels as in our main analysis. However with our panel specification, which includes ZIP code fixed effects, the ΔU and CS variables will capture only a small fraction of what they capture in the main specification. Similarly, some of what we capture in the Bad Mortgage variable in the earlier analysis will now be absorbed by the individual fixed effect. Therefore, instead

TABLE 4 | Panel results for homeowners with a mortgage.

	Prime		Non-Prime	
House Price Growth (Positive)	0.5412*** (0.1094)	0.5257*** (0.1090)	0.7463*** (0.1377)	0.6818*** (0.1377)
House Price Growth (Negative)	0.7870*** (0.1160)	0.7835*** (0.1159)	0.9054*** (0.2799)	1.0674*** (0.2798)
Bad Mortgage	—	−0.0803*** (0.0044)	—	−0.0684*** (0.0011)
Marginal Effects for HP+ (in p.p.)	0.021	0.020	0.031	0.028
Marginal Effects for HP− (in p.p.)	−0.031	−0.030	−0.034	−0.040
First-Stage F-statistics	53,664 / 29,639	53,609 / 29,469	24,305 / 3868	24,116 / 3839
ΔU and Credit Supply	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes
ZIP FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
N	1,828,786	1,828,786	758,800	758,800

Note: Each column shows the results of an IV estimation. The dependent variable is a binary variable showing if the consumer originated an auto loan in a particular year. The mean of the dependent variables is 0.139 for prime borrowers and 0.164 for subprime borrowers. See the text for details.

of providing a full decomposition, we have two objectives in this section. First, to show that even with all the controls, there is a large impact of changes in house prices on auto loan originations. This would be an indication that ex ante financial constraints and the pure wealth effect combined are still very important. Second, even with a specification that includes individual fixed effects, Bad Mortgage continues to be highly significant, indicating that ex-post constraints are important.

The first two columns of Table 4 show the results for Prime homeowners while the latter two columns show the results for Non-Prime homeowners. For each group, we show the results with and without Bad Mortgage.²³ The table also shows first-stage *F*-statistics that are in the thousands, well above any threshold for the relevant statistical tests. Standard errors are clustered at the individual level.

We highlight a few results. First, both positive and negative house price changes lead to large responses in loan originations. Looking at the second panel for the marginal effects of an average house price change, we see that there is a 2 to 3 percentage points increase in loan originations for house price increases and a 3 to 4 percentage points decrease in loan originations for house price decreases.²⁴ Second, homeowners react more to negative house price changes relative to positive changes and Non-Prime homeowners respond more relative to Prime borrowers, both of which indicate that financial constraints are possibly at work. Finally, when we include Bad Mortgage it has a very large and significant coefficient—a Non-Prime (Prime) homeowner reduces their probability of originating an auto loan by 6.8 percentage points (8 percentage points) after experiencing a Bad Mortgage event. Overall these results confirm that our main results remain robust in this larger data set.

5 | Household Financial Constraints and House Values

In this section, we study the relevance of the ex ante and ex post constraints we defined in the Introduction, in turn.

5.1 | Ex Post Financial Constraints

We already introduced our measure of an ex post financial constraint, Bad Mortgage, which is a binary variable that captures whether the consumer has been seriously delinquent (at least 90 days behind) on any mortgage payment or has experienced a foreclosure at any point during 2007–2009. Comparing Columns 3 and 4 of Table 1 it is clear that the introduction of *BM* has a significant impact on the coefficient of Change in House Value: the marginal effect at the mean changes from −1.26 percentage points to −0.77 percentage points. This shows that a sizable fraction of the average effect of Change in House Value was operating through Bad Mortgage. At the end, the results in Table 2 show that Bad Mortgage alone is responsible for 23% of the total effect.

In this section we demonstrate that: (a) house value declines have a very strong effect on Bad Mortgage, (b) Bad Mortgage increases the likelihood of a major decline in the borrower's Risk Score, or what we call being Non-Prime, and (c) Bad Mortgage mainly affects consumption through its negative impact on the consumer's credit history, which makes it more difficult to qualify for an auto loan.

Table 5 shows the results, where, similar to Table 1 we show both the estimated marginal effects in levels and also in percentage points in the second panel. The first column uses Bad Mortgage

TABLE 5 | House values, bad mortgage, and creditworthiness.

Dependent variable	Bad mortgage	Non-prime in 2009		Originate 2009
Change in House Value	−0.0699*** (0.0033)	−0.0705*** (0.0045)	−0.0318*** (0.0037)	0.0096** (0.0039)
Change in Unemployment Rate	−0.0011* (0.0006)	−0.0023*** (0.0008)	−0.0016** (0.0007)	−0.0015** (0.0007)
Credit Supply	0.0522*** (0.0073)	−0.0183* (0.0109)	−0.0419*** (0.0096)	0.0303*** (0.0105)
Bad Mortgage	—	—	0.7125*** (0.0040)	−0.0718*** (0.0036)
ZIP Code and Individual Controls	Yes	Yes	Yes	Yes + 2009 Credit Status
Marginal effects (in p.p.)				
ΔHV (average: −\$75,109)	5.25	5.29	2.39	−0.72
ΔU (average: 5.5 p.p.)	−0.62	−1.23	−0.85	−0.84
Credit Supply (−1 s.d.)	−0.46	0.16	0.37	−0.27
Bad Mortgage (= 1)	—	—	71.25	−7.18

Note: In the first column the dependent variable is having a Bad Mortgage, whose unconditional probability is 7%. In the second and third columns the dependent variable is an indicator for being Non-Prime (Risk Score less than 700) in 2009 and its unconditional probability is 25.3%. The last column repeats the baseline estimation with Loan Origination in 2009 as the dependent variable, mimicking Column 4 of Table 1 and adds controls for the 2009 credit status. The first panel shows the marginal effects of each variable. The second panel shows these marginal effects converted to percentage points. See the notes to Table 1 for details.

as the dependent variable, and the same set of regressors and the estimation strategy as in the main results. The unconditional probability of Bad Mortgage is 7%. The results show that the Change in House Value between 2006 and 2009 has a very strong effect on Bad Mortgage: At the mean it amounts to an increase of 5.25 percentage points, which is about 75% of the unconditional probability. Credit Supply also has a meaningful effect: In counties where credit supply to households was one standard deviation below the mean, the probability of Bad Mortgage was higher by 0.46 percentage points. Local labor market conditions only have a marginally significant effect on Bad Mortgage, indicating that a channel that goes through job loss is not very significant.

The second and third columns show how having a low Risk Score in 2009, in other words being Non-Prime, is influenced by Bad Mortgage. In our sample of consumers, 25.3% are Non-Prime in 2009. Column 2 omits Bad Mortgage and otherwise includes all the controls we have in our main results. The results show a very strong effect of Change in House Value: At the mean it leads to a 5.29 percentage points increase in the probability of being Non-Prime in 2009. Once Bad Mortgage is included in the estimation in Column 3, this large effect of Change in House Value is reduced by 55%, though it remains statistically significant. Moreover, Bad Mortgage itself increases the probability of being Non-Prime by about 71 percentage points, and it is by far the most important determinant of Non-Prime status in 2009. The last column repeats the baseline estimation in Column 4 of Table 1, but adds 2009 credit status as an additional control. Comparing the results in this column with those in Table 1, while the marginal effects of all other variables are unchanged, the effect of Bad Mortgage is reduced by 14% (from 8.29% to 7.18%). This shows that a portion of the effect of Bad Mortgage on consumption is due to the former's effect on the creditworthiness of consumers.²⁵

The results in this section, along with our baseline results that show how Bad Mortgage affects consumption, can be interpreted as follows. Consumers whose house values decline are more likely to fall behind on their mortgage payments. This triggers a significant decline in their Risk Score, increasing the probability that they will be in the Non-Prime category in 2009. Once the consumers' Risk Score falls, their ability to borrow declines and they will be less likely to be able to originate an auto loan in 2009, which is our measure of consumption. Given that Bad Mortgage explains 23% of the effect of change in house values on consumption, this ex post credit constraint channel is compelling and empirically relevant.²⁶

5.2 | Ex Ante Financial Constraints: Uncovering the Heterogeneity

We now turn to the impact of ex ante constraints. Our detailed individual-level data allows us to cut the data in various ways to identify ex ante financial constraints. We consider three ways of observing these constraints at work: Risk Score, LTV ratio, and the type of mortgage, all measured in 2006, *before* the decline in house values. Being Non-Prime indicates the presence of some prior adverse credit activity, which can directly limit future credit access. It can also reflect other (unobserved) financial constraints that may make future credit access more difficult. LTV ratio in 2006 directly reflects the severity of one of the most important financial constraints, the collateral constraint of a mortgage at the time of origination. The higher the LTV ratio, the more constrained the consumer, and thus the more vulnerable they are to house value changes.

To understand why mortgage type reflects ex ante constraints, it is important to keep in mind that borrowers are not allocated

TABLE 6 | Ex ante financial constraints.

(a) Risk Score and LTV								
	Prime				Non-Prime			
LTV0 (First-mortgage LTV < 25%)	−0.16				−3.65*			
LTV1 (First-mortgage LTV between 25% and 50%)	−0.23				0.06			
LTV2 (First-mortgage LTV between 50% and 80%)	−0.93*				−1.74**			
LTV3 (First-mortgage LTV ≥ 80%)	−2.05				−3.48**			
(b) Mortgage Type								
	Prime				Non-Prime			
	LTV0	LTV1	LTV2	LTV3	LTV0	LTV1	LTV2	LTV3
Fixed Rate	−1.05	−0.63	−1.56*	−0.14	−5.27*	−0.91	−3.60***	−4.32*
ARM < 5yr	−0.24	3.90	−1.28	5.38	−8.03*	1.73	2.27	−6.04*
ARM ≥ 5yr	0.97	−3.05	−2.51	−9.28	−2.30	−2.55	2.28	7.10
CE Second	8.27	−0.45	−2.21	−20.53***	13.05	−2.54	0.51	−10.02*
HELOC	0.28	0.01	0.05	−3.12	−0.89	3.18	−0.85	1.90

Note: The table shows the marginal effects of an average Change in House Value on originating an auto loan in 2009 in percentage points. The unconditional probability of the dependent variable is 13.5% and it varies from 9.4% for the Non-Prime / LTV0 group to 16.4% for the Non-Prime / LTV3 group. Prime status, LTV category and mortgage type are all measured as of 2006. Each number is obtained from a separate IV estimation with all standard controls including Bad Mortgage. These are shown in Online Appendix C.4.

randomly to different mortgage types, but rather select the mortgage that best suits their situation, including the financial constraints they face. For example, borrowers with Closed-End Second Mortgages typically originate these because they lack the resources to make the standard 20% down payment on their first mortgage. Further analyzing the distribution of consumers in Table A-2, we see a few more interesting patterns that suggest that they also capture choices by consumers. For example, short-maturity ARMs seem to be chosen by Prime low-to-moderate-LTV borrowers (perhaps because they intend to pay off their loan in a short period of time) or Non-Prime moderate-to-high-LTV borrowers (perhaps because this was the only product they qualified for and they hope to refinance before the ARM resets). HELOCs seem to be favored by Prime borrowers with low-to-moderate LTV ratios. It is plausible that these consumers use the extra liquidity from their HELOCs to finance some consumption expenditures.²⁷ Thus, a decline in house values would make their constraints bind, since banks can (and did) reduce HELOC limits for consumers with increased LTV ratios.

To sum up, all three of these characteristics have implications for the ease of consumers in refinancing their mortgage, how likely it is for them to default and more generally how much their consumption would be affected by changes in house values. Panel (a) of Table 6 shows how consumers in each of the eight categories of LTV and Risk Score react to House Value Change once all controls including Bad Mortgage are included. Each coefficient is obtained from a separate IV estimation, which are reported in Online Appendix C.4, and the table reports them as marginal effects at the average of House Value Change in percentage points. The results show that Prime homeowners do not react to changes

in house values, regardless of LTV ratio, except for a marginally significant response for the LTV2 group. The two highest LTV categories for Non-Prime homeowners show a significant reaction (at 5% significance) at −1.74 percentage points and −3.48 percentage points, which are two to five times the average response.

Panel (b) shows a deeper cut of the results, where we also condition on the type of mortgage the consumer held in 2006. It is useful to interpret the results alongside the distribution of characteristics reported in Table A-2. We find the following results noteworthy. Consumers who are Prime that only have a Fixed Rate first mortgage are about 41% of the population, and they show no reaction to Change in House Value (regardless of LTV), again, except for a marginally significant response for the LTV2 group. In fact, with the exception of a noteworthy group, which we turn to, Prime consumers, 74.7% of the population, do not react to Change in House Value, which is consistent with the results in Panel (a). The Prime group that has a significant reaction is those that have a CE Second mortgage and high LTV. This group, which is only 0.5% of the population, reduces their probability of originating an auto loan by 20.5 percentage points, which is over 26 times the average response. Various groups within Non-Prime consumers show marginally significant responses. The biggest response comes from those with a fixed-rate mortgage and moderately high LTV with a 3.6 percentage points reduction. It is also noteworthy to point out that the Non-Prime consumers with a CE Second and high LTV also show a very large response at 10 percentage points.

We interpret these results as indicating a significant degree of heterogeneity in the consumption response that depends on the financial constraints the consumer had prior to the decline in

house prices. Only a negligible part of the total response comes from Prime consumers with a Fixed Rate, arguably those that are least likely to have ex ante constraints. The three most important groups are all those that have significant ex ante constraints: They are Non-Prime or they have a Closed-End Second mortgage.

In closing this section, we acknowledge that the three sets of measures we use for identifying ex ante constraints—creditworthiness, LTV ratio, and the type of mortgage—may not fully identify all possible ex ante constraints. If this is the case, and if the missing ex ante constraints are observable to banks, they may show up as binding ex post constraints, if banks in fact reduce credit as house prices fell based on the ex ante constraints they observe. This means that some of what we pick up by our ex post constraint measure Bad Mortgage may be banks' reaction to some unobserved (to us) ex ante constraints. Nevertheless, this would not change the conclusion that a majority of the reaction of consumption to house prices is due to credit constraints.

6 | Identifying the Pure Wealth Effect

To take stock of the results so far, we have shown that there is a large response of consumption to Change in House Value and that this can be decomposed into various channels. Local General Equilibrium and Credit Supply channels jointly capture 24% of the total effect. The particular ex post constraints proxied by Bad Mortgage are responsible for another 23%, which leaves 53% for ex ante constraints, other unmodeled ex post constraints and the pure wealth effect. In the previous section, we demonstrated the importance of ex ante constraints. In this section, we show that the pure wealth effect is, in fact, negligible.

We do this in two ways. In Section 6.1, we repeat our baseline estimation for various subsets of consumers for which we would expect that credit constraints should not be important. Thus, if these consumers display a reaction to house value changes, it would likely be due only to the pure wealth effect. In Section 6.2, we use the panel structure of the CCP data set, which allows us to identify homeowners that do not hold a mortgage. A priori, the expectation would be that consumers who own a house without a mortgage would be less likely to be affected by credit constraints, and any consumption response would reflect the pure wealth effect. Our results will show that in all of these cases there is no consumption response to house price changes.

6.1 | Cross-Section Subsample Results

Our first approach to identifying the pure wealth effect relies on the assumption that a consumer who was Prime and had an LTV ratio less than 25% in 2006 would be unlikely to be affected by ex ante credit constraints. The house value of the consumer would have had to decline by more than 75% between 2006 and 2009 for the loan value to exceed the value of the house, something that did not occur in this period. Under this assumption, once we estimate our baseline specification (2), the response to Change in House Value should only reflect the pure wealth effect.

Table 7 reports the coefficient for Change in House Value in a series of subsamples with all the relevant controls (coefficients for

these are not reported). The first row shows the key subsample for our argument, which is Prime consumers with an LTV ratio less than 25%, LTV0. The estimate of 0.0021 has a p -value of 0.80 and it is clearly insignificant. This is our key evidence that the pure wealth effect is negligible. The next two rows show the results when we replace Credit Supply with Bank Health and when we use Probit instead of a linear probability model. In both cases the estimates are insignificant.

There are two possible concerns with this identification, both of which arise from the fact that Prime and low-LTV statuses are not random. First, someone who brought their first-mortgage LTV ratio to a low level may in fact have enough liquid wealth to buy a car without an auto loan. For these individuals we may incorrectly conclude that they did not consume (purchase a car) even though they may have done so using cash. Returning to our main equation in (2), the estimation includes non-housing net worth categories as well as the age polynomial as controls, both of which can be thought of as wealth proxies. The coefficients for non-housing net worth categories show that the wealthier households have about 2 percentage points less probability of originating an auto loan than the poorest category. Considering the unconditional probability is 13%, this is rather a small difference. Similarly the fitted value of the age polynomial is fairly flat indicating that the probability of originating an auto loan does not vary by age. These results address this concern, because we show that the likelihood of origination for the Prime-LTV0 group does not vary significantly by wealth. Thus, it is unlikely that we would be missing the auto purchases of these individuals any more than we would miss them for a random consumer.

The second concern is that an individual who has a low LTV ratio may also have high non-housing wealth and thus the decline in housing wealth may constitute a small share of their total wealth. To address this concern, the rest of Table 7 shows the same marginal effects for two sets of subsamples, first one broken down by non-housing net worth, and next by age. None of the marginal effects are significant, which indicates that the response to changes in house value does not vary with wealth or age and remains insignificant. Thus we conclude that there is no evidence of a significant pure wealth effect.

6.2 | Panel Results—Free-and-Clear Homeowners

In this section, we take a different approach to identifying the pure wealth effect. Our analysis thus far has focused on homeowners with a mortgage, in large part in order to utilize the detailed data we have in CRISM, including individual-level House Value Change. However, this meant we had to leave out an important group of homeowners, one that can uniquely help in identifying the magnitude of the pure wealth effect: homeowners without a mortgage, whom we term free-and-clear homeowners. In this section, we use the panel data introduced earlier, which allows us to capture this group of homeowners.

Unfortunately, the credit bureau data do not contain any direct information on the homeownership status of consumers. However, we are able to use additional information in the bureau records to identify these free-and-clear homeowners.²⁸

TABLE 7 | Identifying the pure wealth effect.

Sample	Δ HV coefficient	Number of obs
LTV0-Prime (Benchmark)	0.0021 (0.0082)	34,137
LTV0-Prime (with Bank Health)	−0.0032 (0.0102)	34,137
LTV0-Prime (using Probit)	−0.0014 (0.0083)	34,137
LTV0-Prime, Non-Housing Net Worth $\leq 25^{th}$ Pct	0.0079 (0.0281)	2150
LTV0-Prime, Non-Housing Net Worth $\in (25, 50]$ Pct	0.0043 (0.0235)	3374
LTV0-Prime, Non-Housing Net Worth $\in (50, 75]$ Pct	0.0176 (0.0164)	10,164
LTV0-Prime, Non-Housing Net Worth $> 75^{th}$ Pct	−0.0043 (0.0115)	18,449
LTV0-Prime, Age < 41	0.0398 (0.0451)	1094
LTV0-Prime, Age $\in [41, 60]$	−0.0039 (0.0111)	19,056
LTV0-Prime, Age > 60	0.0082 (0.0130)	13,987

Note: The table shows the estimated β^1 coefficient from (2) for the subset of consumers as shown in the first column. See notes to Table 1.

TABLE 8 | Panel results.

	Prime	Non-Prime
House Price Growth (Positive)	0.1755 (0.2427)	−0.0382 (0.2574)
House Price Growth (Negative)	0.3323 (0.2130)	−0.4006 (0.2437)
First-Stage <i>F</i> -statistics	9344 / 6207	4870 / 2568
Δ U and Credit Supply	Yes	Yes
Individual FE	Yes	Yes
ZIP FE	Yes	Yes
Year FE	Yes	Yes
N	396,947	262,285

Note: Each column shows the results of an IV estimation. The dependent variable is a binary variable showing if the consumer originated an auto loan in a particular year. The mean of the dependent variables across columns are: 0.077 and 0.098. See the text for details.

Table 8 shows the panel estimation results where we show the estimates β^{1+} and β^{1-} for Prime and Non-Prime consumers. Neither of the estimates for β^{1+} or β^{1-} are statistically significant, indicating that Free-and-Clear Homeowners do not change their consumption in response to a change in their house price. This result, once again, shows that the pure wealth effect is not

important in shaping the reaction of consumption to changes in house prices.

7 | Conclusion

We investigate the role of household heterogeneity—in terms of wealth and financial constraints—on the response of aggregate consumption to the decline in house prices during the 2007–2009 crisis. Unlike most other studies that focus on the link between house values and consumption, we use individual-level data drawn from consumer credit bureau records linked to mortgage data, proxying consumption by auto loan originations. This allows us to use not only several key characteristics of consumers such as their age, their creditworthiness, and the type of mortgage they have, but also the change in their houses' value, all at the individual level.

The average decline in house values between 2006 and 2009 leads to a decline of about \$1225 in consumption in 2009. We decompose this effect into four channels: 13% for the Local General Equilibrium (declining employment reducing consumption), 11% for the Household Credit Supply (declining bank wealth), 23% for the effect of ex post financial constraints, and 53% for the effect of ex ante household credit constraints. We show that the pure wealth effect of house value changes on consumption is negligible. These results can help design theories and policies in the future.

Acknowledgments

The authors thank participants at various conferences and seminars for helpful comments, Aditya Aladangady, Stefania Albanesi, Saleem Bahaj, John Chao, Jessica Goldberg, Adam Guren, Guido Kuersteiner, Felipe Saffie, and Frank Schorfheide for useful discussions; Charles Nathanson and Kurt Mitman for their discussions; and Omar Ahmad, Aaron Payne, and Di Wang for excellent research assistance. We also thank Jesús Fernández-Villaverde, the editor, and two anonymous referees for their constructive and useful comments. The views expressed in this paper are solely those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Endnotes

¹ Some of the other papers that also use CRISM data are Beraja et al. (2019) to investigate the response of auto consumption to monetary policy; Agarwal et al. (2023) and Di Maggio et al. (2020) to investigate refinancing; and García (2019) to investigate the secondary housing market.

² Many argued that to be able to match the large responses of consumption to house price changes found in the data, one needs collateralized lending that amplifies the impact of housing wealth on consumption. See Berger et al. (2018), Guerrieri and Iacoviello (2017), and Iacoviello (2005).

³ This importance of the credit supply channel has found support in the work of Justiniano et al. (2019); they show that an increase in credit supply is necessary to match the empirical regularities in the boom period (where the increase in house prices served as a positive shock to bank balance sheets).

⁴ DeFusco (2018) is a related paper which uses individual-level data between 1997 and 2012 from Montgomery County, Maryland. He uses the expiration of price controls, which allows homeowners to use the market value of their house as collateral, as a source of exogenous variation. Since market prices of houses are unchanged, he is able to disentangle the pure wealth effect and the effect of collateral constraints. His primary variable of interest is the change in collateralized borrowing and he shows a large response of this to the increase in collateral. As such, his results are in line with ours, where a change in collateral without a change in wealth results in a sizable response by consumers.

⁵ 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 et al. (2010).

⁶ When available, we use house price indices at the ZIP code level. If this is not available we use the next highest level (county or state). All house price indices are from CoreLogic.

⁷ Even though both CRISM and CCP are available as panels, in the majority of our analysis we draw information from individual years and conduct a cross-sectional analysis, so as to more crisply identify the change in house value. In Sections 4.4 and 6.2, we utilize the panel structure of CCP. We postpone the introduction of this analysis and the data used to Section 4.4.

⁸ See <https://www.federalreserve.gov/pubs/bulletin/2009/pdf/scf09.pdf>.

⁹ This is an upper bound, as any other debt would likely appear in their credit bureau data, even if recently paid off.

¹⁰ https://files.consumerfinance.gov/f/documents/cfbp_update-credit-invisibles-estimate_2025-06.pdf.

¹¹ Lenders use a variety of credit scores, often different ones depending on the context (e.g., mortgages versus auto loans) and finer categories when they make lending decisions. Our categorization serves as a rough proxy for the broad creditworthiness of the households.

¹² A closed-end second mortgage is one that is junior to the first mortgage, and also does not allow any further draws following the origination date (in contrast to a Home Equity Line of Credit).

¹³ PSID, public use data set. Produced and distributed by the Institute for Social Research, University of Michigan, Ann Arbor, MI (2017). This figure 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.

¹⁴ Coibion et al. (2020) do a similar imputation for the income of consumers in CCP.

¹⁵ Aladangady (2017) relies on the same pair of instruments interacted with an aggregate demand shifter in a panel structure.

¹⁶ Online Appendix C.1 reports results from a restricted model where we do not allow for this heterogeneity. There is overwhelming statistical evidence in favor of allowing for this heterogeneity; a test with a null of the restricted model is rejected at any level of significance.

¹⁷ Practically speaking, we estimate the equation by interacting the two-dimensional Risk Score categories and the four-dimensional LTV categories with all regressors in the first and second stage, except for ΔHV_i . Doing so yields eight estimated parameters per regressor.

¹⁸ Mian et al. (2013) use one of the same instruments, Land Unavailable, and obtain the same sign for the period 2006 to 2009. Mian and Sufi (2009) also show that between 2002 and 2006, when there was a strong increase in housing demand across the country due to cheaper credit, regions with inelastic housing supply showed larger house price increases.

¹⁹ Based on Bureau of Transportation Statistics data, the average price of a car (new or used) in 2009 was \$12,518. Combining these we find that the change in consumption in 2009 at the mean house value change is $\Delta C = \$12,518 \times (0.0176 \times (-\$0.75109)) / 0.1352 = -\$1225$.

²⁰ In Online Appendix C.2 we redo the estimation in Table 1 using a sample of consumers that have one or two mortgages and, separately, those that have exactly two mortgages. Both sets of results are very similar to those in Table 1.

²¹ The “true” and “misspecified” labels follow the OVB literature. These labels are somewhat inappropriate for the general IV case since both models satisfy the assumptions of IV. In this case the smaller model captures the combined effect of all channels and the larger model “over controls” and identifies the individual channels. We thank an anonymous referee for emphasizing this nuance. Moreover, even though in Table 1 we introduced the additional variables in a particular order, the decomposition only requires the results of Column 1 as the misspecified model and Column 4 as the true model, and as such the order in which the variables were introduced in Table 1 is not relevant.

²² We estimated the version with only a single variable as well and a Wald test testing the restriction has a p -value of 0 indicating strong rejection of this restriction at any level of significance.

²³ In both first stages for Positive House Price Growth the instruments have a positive sign, and for Negative House Price Growth they have a negative sign. This is consistent with the earlier results from a linear model where, when faced with a positive demand shock (that increases prices nationwide), areas with more inelastic housing supply see large price increases and this is reversed when the demand shock turns negative.

²⁴ These results are considerably larger than our baseline result where after all controls loan originations declined 0.77 percentage points in 2009. We attribute these differences to using house price changes at the ZIP code level versus using homeowner-specific house value changes, not being able to use a key control, the LTV ratio, and using 3-year changes versus 1-year changes.

²⁵ As 2009 credit status is a binary indicator, this likely does not capture the full impact of Bad Mortgage on creditworthiness.

²⁶ We conclude this section with two additional points. Detailed results regarding these are available upon request. First, one might be concerned that our Bad Mortgage variable is picking up the effect of other shocks that the consumer experienced over this time period. To address this, we augment the baseline specification in (2) as shown in the fourth column of Table 1 with the Utilization Rate of consumers by the end of 2009. This shows the fraction of available credit consumers are using and captures the cumulative adverse liquidity shocks (such as being unemployed or having health issues) consumers experienced in the preceding years. In this specification Bad Mortgage remains highly significant with a coefficient that is only slightly smaller than our baseline specification. This shows that Bad Mortgage measures something distinct from general liquidity shocks. Second, we augment the baseline with a variable we call Bad Card, which is the counterpart of Bad Mortgage but computed for credit cards instead of mortgages. In this specification the coefficient of Bad Mortgage falls from -0.083 to -0.058 and remains significant, while Bad Card has a coefficient of -0.061 . This shows that, once again, Bad Mortgage measures something distinct, triggered by house price declines, and not merely reflecting general credit problems.

²⁷ One may be concerned that consumers can use cash they obtain 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 could identify our results using auto loan originations for people with HELOCs. Results reported by McCully et al. (2019), however, show—using data from three nationally representative surveys—that very few consumers purchase cars outright using HELOCs or cash-out refinancing.

²⁸ We use the following algorithm. If the consumer was not a Free-and-Clear Homeowner in year $t - 1$, then if they have a mortgage on their record in Q4 of year $t - 1$ and no mortgage in Q4 of year t , they do not have a mortgage foreclosure in year $t + 1$ (so the lack of a mortgage reflects paying it off), and their address is the same in Q4 of year $t - 1$ and Q4 of year t , then they are labeled as a free-and-clear consumer. If the consumer was identified as a Free-and-Clear Homeowner in year $t - 1$, then as long as they continue to have no mortgage in Q4 of year t , they do not have a mortgage foreclosure in year $t + 1$ and their address is the same in Q4 of year $t - 1$ and Q4 of year t , then they are again labeled as a Free-and-Clear Homeowner in year t .

References

- Adelino, M., A. Schoar, and F. Severino. 2016. "Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class." *Review of Financial Studies* 29: 1635–1670.
- Adelino, M., A. Schoar, and F. Severino. 2018. "Dynamics of Housing Debt in the Recent Boom and Great Recession." In *NBER Macroeconomics Annual 2017*, 265–311. University of Chicago Press.
- Agarwal, S., G. Amromin, S. Chomsisengphet, et al. 2023. "Mortgage Refinancing, Consumer Spending, and Competition: Evidence From the Home Affordable Refinance Program." *Review of Economic Studies* 90: 499–537.
- Aladangady, A. 2017. "Housing Wealth and Consumption: Evidence From Geographically-Linked Microdata." *American Economic Review* 107: 3415–3446.
- Albanesi, S., G. DeGiorgi, and J. Nosal. 2022. "Credit Growth and the Financial Crisis: A New Narrative." *Journal of Monetary Economics* 132: 118–139.
- Bahaj, S., A. Foulis, and G. Pinter. 2020. "Home Values and Firm Behavior." *American Economic Review* 110: 2225–2270.
- Beraja, M., A. Fuster, E. Hurst, and J. Vavra. 2019. "Regional Heterogeneity and the Refinancing Channel of Monetary Policy." *Quarterly Journal of Economics* 134: 109–183.
- Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra. 2018. "House Prices and Consumer Spending." *Review of Economic Studies* 85: 1502–1542.

- Buiter, W. H. 2010. "Housing Wealth Isn't Wealth." *Economics* 4: 20100022.
- 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: 55–79.
- Chodorow-Reich, G. 2014. "The Employment Effects of Credit Market Disruptions: Firm-Level Evidence From the 2008–9 Financial Crisis." *Quarterly Journal of Economics* 129: 1–59.
- Coibion, O., Y. Gorodnichenko, M. Kudlyak, and J. Mondragon. 2020. "Greater Inequality and Household Borrowing: New Evidence From Household Data." *Journal of the European Economic Association* 18: 2922–2971.
- DeFusco, A. A. 2018. "Homeowner Borrowing and Housing Collateral: New Evidence From Expiring Price Controls." *Journal of Finance* 73: 523–573.
- Di Maggio, M., A. Kermani, B. J. Keys, et al. 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review* 107: 3550–3588.
- Di Maggio, M., A. Kermani, and C. J. Palmer. 2020. "How Quantitative Easing Works: Evidence on the Refinancing Channel." *Review of Economic Studies* 87: 1498–1528.
- Elul, R., N. S. Souleles, S. Chomsisengphet, D. Glennon, and R. Hunt. 2010. "What 'Triggers' Mortgage Default?" *American Economic Review* 100: 490–494.
- Foote, C. L., L. Loewenstein, and P. S. Willen. 2020. "Cross-Sectional Patterns of Mortgage Debt During the Housing Boom: Evidence and Implications." *Review of Economic Studies* 88: 229–259.
- García, D. 2019. "Second-Home Buying and the Housing Boom and Bust." *Real Estate Economics* 50, no. 1: 33–58.
- Garriga, C., and A. Hedlund. 2020. "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession." *American Economic Review* 110: 1603–1634.
- Gerardi, K., K. F. Herkenhoff, L. E. Ohanian, and P. S. Willen. 2018. "Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default." *Review of Financial Studies* 31: 1098–1131.
- Gilchrist, S., M. Siemer, and E. Zakrajsek. 2023. "Bank Health and Local Economic Outcomes." Mimeo. New York University.
- Greenstone, M., A. Mas, and H.-L. Nguyen. 2020. "Do Credit Market Shocks Affect the Real Economy? Quasi-experimental Evidence From the Great Recession and 'Normal' Economic Times." *American Economic Journal: Economic Policy* 12: 200–225.
- Guerrieri, L., and M. Iacoviello. 2017. "Collateral Constraints and Macroeconomic Asymmetries." *Journal of Monetary Economics* 90: 28–49.
- Guren, A. M., A. McKay, E. Nakamura, and J. Steinsson. 2021. "Housing Wealth Effects: The Long View." *Review of Economic Studies* 88: 669–707.
- 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: 693–729.
- Iacoviello, M. 2005. "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle." *American Economic Review* 95: 739–764.
- Imbens, G. W., and J. D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62: 467–475.
- Justiniano, A., G. E. Primiceri, and A. Tambalotti. 2019. "Credit Supply and the Housing Boom." *Journal of Political Economy* 127: 1317–1350.
- Kaplan, G., K. Mitman, and G. L. Violante. 2020. "The Housing Boom and Bust: Model Meets Evidence." *Journal of Political Economy* 128: 3285–3345.
- McCully, B. A., K. M. Pence, and D. J. Vine. 2019. "How Much Are Car Purchases Driven by Home Equity Withdrawal?" *Journal of Money, Credit and Banking* 51: 1403–1426.

- Mian, A., K. Rao, and A. Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics* 128: 1687–1726.
- Mian, A., and A. Sufi. 2009. "The Consequences of Mortgage Credit Expansion: Evidence From the US Mortgage Default Crisis." *Quarterly Journal of Economics* 124: 1449–1496.
- Mian, A., and A. Sufi. 2010. "Household Leverage and the Recession of 2007–09." *IMF Economic Review* 58: 74–117.
- Mian, A., and A. Sufi. 2011. "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis." *American Economic Review* 101: 2132–2156.
- Mian, A., and A. Sufi. 2016. "Who Bears the Cost of Recessions? The Role of House Prices and Household Debt." *Handbook of Macroeconomics* 2: 255–296.
- Pistaferri, L. 2016. "Why Has Consumption Remained Moderate After the Great Recession?" Mimeo. Stanford University.
- Rosen, R. J. 2011. "Competition in Mortgage Markets: The Effect of Lender Type on Loan Characteristics." *Federal Reserve Bank of Chicago Economic Perspectives* 35, no. 1: 2–21.
- Saiz, A. 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125: 1253–1296.
- Schmalz, M. C., D. A. Sraer, and D. Thesmar. 2017. "Housing Collateral and Entrepreneurship." *Journal of Finance* 72: 99–132.
- Sinai, T., and N. S. Souleles. 2005. "Owner-Occupied Housing as a Hedge Against Rent Risk." *Quarterly Journal of Economics* 120: 763–789.
- Stock, J., and M. Yogo. 2005. "Testing for Weak Instruments in Linear IV Regression." In *Identification and Inference for Econometric Models*, 80–108. Cambridge University Press.
- Vestman, R., J. Bojeryd, B. Tyrefors, and D. Kessel. 2026. "The Housing Wealth Effect: Quasi-Experimental Evidence." *Journal of Monetary Economics* forthcoming 158: 103892.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.

Data S1