

Housing Wealth and Consumption: The Role of Heterogeneous Credit Constraints*

S. Borağan Aruoba

Ronel Elul

University of Maryland and NBER

FRB Philadelphia

Şebnem Kalemli-Özcan

Brown University and NBER

January 2026

Abstract

We quantify the role of heterogeneity in households' financial constraints in explaining the link between the large 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.

KEYWORDS: financial crisis, mortgage, individual-level data, bank health, credit supply, wealth effect

*Correspondence: Aruoba: Department of Economics, University of Maryland, College Park, MD 20742. Email: aruoba@umd.edu, Elul: Research Department, Federal Reserve Bank of Philadelphia, Philadelphia, PA 19106. Email: ronel.elul@phil.frb.org, Kalemli-Özcan: Department of Economics, Brown University, Providence, RI 02912. Email: sebnem_kalemli-ozcan@brown.edu. 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, to 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 these papers 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.

1 Introduction

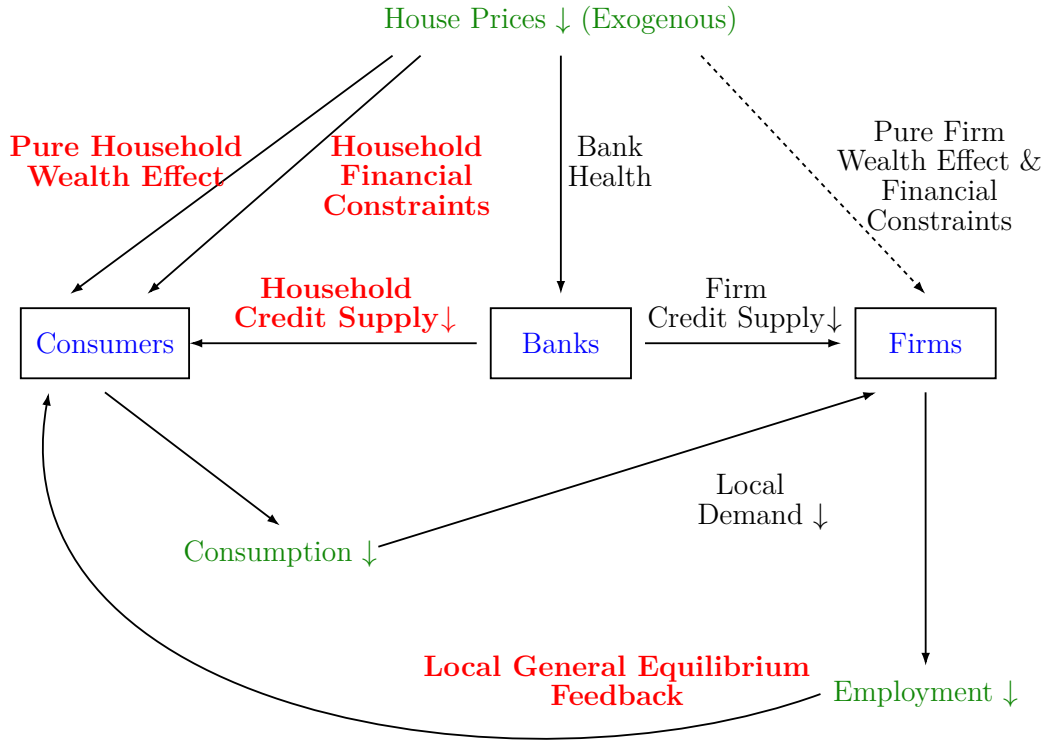
The U.S. economy experienced a large financial crisis together with a large nationwide decline in house prices in 2007–2008. A deep recession with significant declines in consumption, investment, and employment followed. The house price declines across the country showed a large degree of variation: across U.S. 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 datasets 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, deleveraging and consumption decline and thus point to 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 is used, conclusions based on aggregate

¹Some 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.

Figure 1: House Prices and Consumption: Channels



(e.g. ZIP-level) data 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

²Many 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 \(2017\)](#), [Iacoviello \(2005\)](#).

³This 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).

to obtain loans to invest and to produce. [Bahaj et al. \(2020\)](#) provides direct evidence for this channel for the U.K. and [Schmalz et al. \(2017\)](#) show that an increase in one’s house value as a 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 it provides is 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 (IV) 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 (for example 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

⁴[DeFusco \(2018\)](#) is a related paper which uses individual-level data between 1997-2012 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 price 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.

(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 helps 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 detail may be found in [Appendix A](#).

2.1 Individual-Level Data Sources

Our main dataset 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 dataset, 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

⁵The 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. \(2010\)](#).

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 dataset, as it allows us to compute the dollar change in the value of the home of each individual in our dataset 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 in 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 dataset 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 Appendix C.2, are virtually identical those in the body of the paper.

From the 2007 Survey of Consumer Finances, 71% of all homeowners had an active

⁶When 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.

⁷Even 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

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 datasets. 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 Finance Protection Bureau finds that less than 6% of the adult population in the U.S. does not have a credit bureau record.¹⁰

In contrast to other, publicly available datasets such as the PSID, the advantage of our credit bureau data is that it allows 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, 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

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

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

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

¹¹Lenders 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.

create four groups: those with a 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 five years**, and separately those that have a duration **greater than or equal to five 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 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.

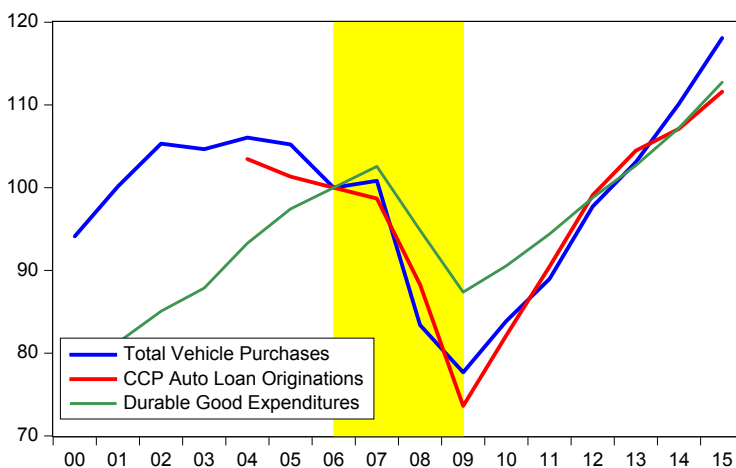
2.2 Construction of the Consumption Proxy

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

¹²A 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).

¹³Panel Study of Income Dynamics, public use dataset. Produced and distributed by the Institute for Social Research, University of Michigan, Ann Arbor, MI (2017). This figure is computed from

Figure 2: BEA vs. CCP: Consumption Expenditures



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.

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 dataset 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

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.

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

While the CRISM data allows 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 does not contain information on their assets. To overcome this problem we use the Survey of Consumer Finances (SCF) to impute non-housing assets of consumers in our data.¹⁴ The details are provided in 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 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

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

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 Instrumental Variable (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 50km 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

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

estate development.

3 From ZIP Code to Individual Level

3.1 Mian, Rao and Sufi’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 U.S. 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 U.S. and then normalizes by the number of households in that ZIP to get a per-capita measure. Let $\Delta HP_z^{2006-2009}$ 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

$$C_{z,2009} - C_{z,2006} = \alpha^{MRS} + \underset{(0.001)}{0.018} \Delta HP_z^{2006-2009} + \varepsilon_z^{MRS} \text{ with } R^2 = 0.153, \text{ and } N = 6,263.$$

This shows a highly significant effect of the change in house prices on change in consumption: an \$18 decline for every \$1,000 drop in house prices, which can be translated to a dollar change of $\overline{\Delta C^{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 \$1,000 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 goal 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. (2013)'s analysis in the previous section, in addition to the obvious one that involves change in the unit of observations. 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

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

¹⁶Appendix 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.

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 is 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 $\{A_l\}$ 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 co-

¹⁷Practically 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.

efficients 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,

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.

which is 13.5%. Following the same approach as in Section 3.1, this translates into a dollar response of $\overline{\Delta C} = -\$1,225$.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 \$1,225, 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:

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

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

¹⁹Based 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 $\overline{\Delta C} = \$12,518 \times \frac{0.0176 \times (-\$0.75109)}{0.1352} = -\$1,225$.

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)	-\$1,225	-\$1,368	-\$1,163	-\$715
ΔU (average: -5.5 p.p.)	-	-\$636	-\$694	-\$756
Credit Supply (-1 s.d.)	-	-	-\$214	-\$254
Bad Mortgage (= 1)	-	-	-	-\$7,674

Notes: 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.

whether or not the consumer has had difficulty in paying their mortgage from 2006-2009. 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” 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

²⁰In 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.

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 \$7,674.

4.2 Decomposing the Total Effect

Using a methodology adapted from the derivations of the omitted variable bias (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 omitted variable bias 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 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 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

²¹The “true” and “misspecified” labels follow the omitted-variable bias 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.

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%

Notes: 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 Appendix B. Numbers in each row may not add to 100% due to rounding.

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 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 allocate the 2008 “Lehman shock” of Chodorow-Reich (2014) to U.S. 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 Appendix C.3. These results confirm the robustness of our main results. Considering the specification with Bank Health,

Table 3: Robustness of Main Results (Marginal Effects, in p.p.)

Main Results		
Δ HV (average: -\$75,109)	-1.32***	-0.77***
Δ U (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		
Δ HV (average: -\$75,109)	-1.32***	-0.81***
Δ U (average: -5.5 p.p.)	-	-0.98**
Bank Health (+1 s.d.)	-	-0.25***
Bad Mortgage (=1)	-	-8.28***
Probit-IV		
Δ HV (average: -\$75,109)	-0.32*	0.09
Δ U (average: -5.5 p.p.)	-	-1.61***
Credit Supply (-1 s.d.)	-	-0.41***
Bad Mortgage (=1)	-	-8.44***

Notes: 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.

the marginal effect for Δ HV is roughly unchanged, while that for Δ U is modestly larger. Probit results show much smaller total effect of Change in House Value (-0.32 versus -1.32), and a much smaller direct effect (an insignificant 0.09 versus -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 two 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 four 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

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

$$+ \alpha_i^j + \alpha_t^j + \alpha_{z(it)}^j + \varepsilon_{it}$$

for individuals with $j(it) = j$, where $j(it)$, $z(it)$, and $c(it)$ denotes 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. 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 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 of any threshold for the relevant statistical tests. Standard errors are clustered at the individual level.

²²We 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.

²³In 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.

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 / 3,868	24,116 / 3,839
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

Notes: 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.

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 p.p. increase in loan originations for house price increases and a 3 to 4 p.p. 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 p.p. (8 p.p.) after experiencing a Bad Mortgage event. Overall these results confirm that our main results remain robust in this larger dataset.

²⁴These results are considerably larger than our baseline result where after all controls loan originations declined 0.77 p.p. 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.

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 p.p. to -0.77 p.p. 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 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 a decline of 5.25 p.p., which is almost 80% of the unconditional probability. Credit

Table 5: House Values, Bad Mortgage and Creditworthiness

<i>Dependent Variable</i>	<i>Bad Mortgage</i>	<i>Non-Prime in 2009</i>	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.)				
Δ HV (average: -\$75,109)	5.25	5.29	2.39	-0.72
Δ U (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

Notes: 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.

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 p.p. 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 p.p. 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 p.p., 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

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

²⁶We 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 distinctly different than 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.

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 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

²⁷One 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

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

Notes: The table shows the marginal effects of an average Change in House Value on originating an auto loan in 2009 in p.p. 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 Appendix C.4.

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 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

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.

categories for Non-Prime homeowners show a significant reaction (at 5% significance) at -1.74 pp and -3.63 p.p., 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 p.p., 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 p.p. 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 p.p.

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 wouldn't 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 dataset, 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 of 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

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 (0.0069)	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)	2,150
LTV0-Prime, Non-Housing Net Worth $\in (25, 50]$ Pct	0.0043 (0.0235)	3,374
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)	1,094
LTV0-Prime, Age $\in [41, 60]$	-0.0039 (0.0111)	19,056
LTV0-Prime, Age > 60	0.0082 (0.0130)	13,987

Notes: 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.

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 p.p. 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

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 F-statistics	9344 / 6207	4870 / 2568
ΔU and Credit Supply	Yes	Yes
Individual FE	Yes	Yes
ZIP FE	Yes	Yes
Year FE	Yes	Yes
N	396,947	262,285

Notes: 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.

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 does 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 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

²⁸We 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 .

change their consumption in response to a change in their house price. This result, once again, shows that 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 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 \$1,225 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.

References

- Adelino, M., A. Schoar, and F. Severino (2016, 03). Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class. *The Review of Financial Studies* 29(7), 1635–1670.
- Adelino, M., A. Schoar, and F. Severino (2018, October). Dynamics of housing debt in the recent boom and great recession. In *NBER Macroeconomics Annual 2017*, pp. 265–311. University of Chicago Press.

- Agarwal, S., G. Amromin, S. Chomsisengphet, T. Landvoigt, T. Piskorski, A. Seru, and V. Yao (2023). Mortgage refinancing, consumer spending, and competition: Evidence from the home affordable refinance program. *The Review of Economic Studies* 90(2), 499–537.
- Aladangady, A. (2017). Housing wealth and consumption: Evidence from geographically-linked microdata. *American Economic Review* 107(11), 3415–46.
- Albanesi, S., G. DeGiorgi, and J. Nosal (2022). Credit growth and the financial crisis: A new narrative. *Journal of Monetary Economics* 132, 118–139.
- Anenberg, E., A. Hizmo, E. Kung, and R. Molloy (2019). Measuring mortgage credit availability: A frontier estimation approach. *Journal of Applied Econometrics* 34(6), 865–882.
- Bahaj, S., A. Foulis, and G. Pinter (2020, July). Home values and firm behavior. *American Economic Review* 110(7), 2225–70.
- Beraja, M., A. Fuster, E. Hurst, and J. Vavra (2019). Regional heterogeneity and the refinancing channel of monetary policy. *Quarterly Journal of Economics* 134(1), 109–183.
- Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2018). House prices and consumer spending. *Review of Economic Studies* 85(3), 1502–1542.
- Buiter, W. H. (2010). Housing wealth isn't wealth. *Economics* 4(1), 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(1), 55–79.
- Chen, X., D. T. Jacho-Chávez, and O. Linton (2016). Averaging of and increasing number of moment condition estimators. *Econometric Theory* 32(1), 30–70.
- 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), 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(6), 2922–2971.
- DeFusco, A. A. (2018). Homeowner borrowing and housing collateral: New evidence from expiring price controls. *Journal of Finance* 73(2), 523–573.

- Di Maggio, M., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2017, November). Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review* 107(11), 3550–88.
- Di Maggio, M., A. Kermani, and C. J. Palmer (2020). How quantitative easing works: Evidence on the refinancing channel. *Review of Economic Studies* 87(3), 1498–1528.
- Elul, R., A. Payne, and S. Tilson (2023, September). Owner occupancy fraud and mortgage performance. *Real Estate Economics* 51(5), 1137–1177.
- Elul, R., N. S. Souleles, S. Chomsisengphet, D. Glennon, and R. Hunt (2010, May). What ‘triggers’ mortgage default? *American Economic Review* 100(2), 490–94.
- Foote, C. L., L. Loewenstein, and P. S. Willen (2020, 08). Cross-Sectional Patterns of Mortgage Debt during the Housing Boom: Evidence and Implications. *The Review of Economic Studies* 88(1), 229–259.
- García, D. (2019). Second-home buying and the housing boom and bust. *Real Estate Economics*.
- Garriga, C. and A. Hedlund (2020). Mortgage debt, consumption, and illiquid housing markets in the great recession. *American Economic Review* 110(6), 1603–34.
- Gilchrist, S., M. Siemer, and E. Zakrajsek (2023). Bank health and local economic outcomes. Mimeo.
- Glaeser, E. L., J. Gyourko, and A. Saiz (2008). Housing supply and housing bubbles. *Journal of Urban Economics* 64(2), 198–217.
- Greenstone, M., A. Mas, and H.-L. Nguyen (2020, February). 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(1), 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(2), 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(3), 693–729.
- Hoekstra, M., S. L. Puller, and J. West (2017, July). Cash for corollas: When stimulus reduces spending. *American Economic Journal: Applied Economics* 9(3), 1–35.

- Iacoviello, M. (2005). House prices, borrowing constraints, and monetary policy in the business cycle. *American Economic Review* 95(3), 739–764.
- Imbens, G. W. and J. D. Angrist (1994). Identification and estimation of local average treatment effects. *Econometrica* 62(2), 467–475.
- Justiniano, A., G. E. Primiceri, and A. Tambalotti (2019). Credit supply and the housing boom. *Journal of Political Economy* 127(3), 1317–1350.
- Kaplan, G., K. Mitman, and G. L. Violante (2020). The housing boom and bust: Model meets evidence. *Journal of Political Economy* 128(9), 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(5), 1403–1426.
- Mian, A., K. Rao, and A. Sufi (2013). Household balance sheets, consumption, and the economic slump. *Quarterly Journal of Economics* 128(4), 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(4), 1449–1496.
- Mian, A. and A. Sufi (2010). Household leverage and the recession of 2007–09. *IMF Economic Review* 58(1), 74–117.
- Mian, A. and A. Sufi (2011). House prices, home equity–based borrowing, and the US household leverage crisis. *American Economic Review* 101(5), 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.
- Rosen, R. J. (2011). Competition in mortgage markets: The effect of lender type on loan characteristics. *Federal Reserve Bank of Chicago Economic Perspectives* 35:1.
- Saiz, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics* 125(3), 1253–1296.
- Schmalz, M. C., D. A. Srael, and D. Thesmar (2017). Housing collateral and entrepreneurship. *The Journal of Finance* 72(1), 99–132.
- Sinai, T. and N. S. Souleles (2005). Owner-occupied housing as a hedge against rent risk. *Quarterly Journal of Economics* 120(2), 763–789.

Stock, J. and M. Yogo (2005). Testing for weak instruments in linear iv regression. In *Identification and Inference for Econometric Models*, pp. 80–108. New York: 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.

Internet Appendix (For Online Publication)

A Data

This section introduces our data in detail. We first introduce our individual-level data followed by aggregate (ZIP-, county- and MSA-level) data. Table [A-1](#) shows the descriptive statistics for the key variables we use in the analysis.

A.1 Individual-Level Data Sources

A.1.1 Credit-Bureau Data

Using the information in CRISM, we estimate the first-mortgage loan-to-value (LTV) ratio for the house as of December 2006, by dividing the remaining balance in the first mortgage by the value of the house. In our analyses we drop individuals with an estimated LTV ratio of 125% or higher. Table [A-1](#) shows the descriptive statistics for the key variables we use in the analysis. It shows that the average and median decline in house values in our sample from December 2006 to December 2009 is \$78,100 and \$52,800, respectively, where the average decline is about 20 percent. All but about 5% of the individuals experience a house value decline, with the fifth percentile at a decline of \$229,300.

For computing **Bad Mortgage**, we use both the payment status for the mortgage in CRISM as well as that reported in CCP to identify households that are seriously delinquent (at least 90 days behind) in any mortgage payment during 2007-2009. Note that since we require the presence of a mortgage in December 2009, we generally drop those that completed the foreclosure process by that point. As the majority of defaults in our sample occurred in 2008 and 2009, and the average foreclosure timeline in this period exceeded a year, this is not a significant limitation.

Table A-1: Summary Statistics**(a) Individual-Level Variables**

	Mean	Std. Dev.	5%	Median	95%
Originate Auto Loan in 2009	0.135	0.342	0	0	1
Change in House Value (\$1,000)	-75.1	97.0	-233.0	-45.8	1.0
Bad Mortgage	0.070	0.255	0	0	1
2006 Non-Housing Net Worth (\$1,000)	144.1	399	10.8	92.3	349.6

(b) Aggregate Variables

	Mean	Std. Dev.	5%	Median	95%
Change in Unemployment Rate (county, p.p.)	5.5	1.8	3.0	5.3	8.7
2003 ZIP-Code Auto Sales (per-capita, \$)	3,263.2	855.0	1926.7	3218.9	4745.1
Credit Supply Shocks (county, $\times 100$)	-2.8	8.8	-13.2	-4.8	14.1
Bank Health (county, $\times 100$)	0.64	0.12	0.42	0.65	0.78
Land Unavailable for Development (MSA)	0.29	0.21	0.03	0.251	0.67
WRLURI (MSA)	0.25	0.67	-0.81	0.31	1.60

Notes: Change in House Value and Change in Unemployment Rate are computed between December 2006 and December 2009. See the main text for the definitions of the variables.

A.1.2 Household Classifications

We classify the 345,000 households in our data in three different ways, all using *ex-ante* criteria that are observed as of 2006Q4, the start of our analysis. Table A-2 shows the fraction of households that fall in each group. First, we label households with an Risk Score of 700 or higher as **Prime** and the others as **Non-Prime**. About a quarter of households are in the Non-Prime category. To be sure, 700 is a fairly high cutoff for prime borrowers, reflecting the fact that our analysis focuses on homeowners, who tend to be more creditworthy. Figure A-1 shows the distribution of ZIP codes with respect to the fraction of non-prime borrowers. This shows that a vast majority of ZIP codes have a mixture of Prime and Non-Prime borrowers, and thus ZIP code-level variables and the individual-level indicator of prime status will contain largely independent information.

Table A-2: Distribution of Characteristics

LTV Category	Prime	Non-Prime	Total
LTV0 (LTV ratio less than 25%)	9.9%	1.3%	11.2%
LTV1 (LTV ratio between 25% and 50%)	27.7%	4.9%	32.6%
LTV2 (LTV ratio between 50% and 80%)	31.7%	12.8%	44.6%
LTV3 (LTV ratio greater than 80%)	5.4%	6.3%	11.6%
Total	74.7%	25.3%	100.0%

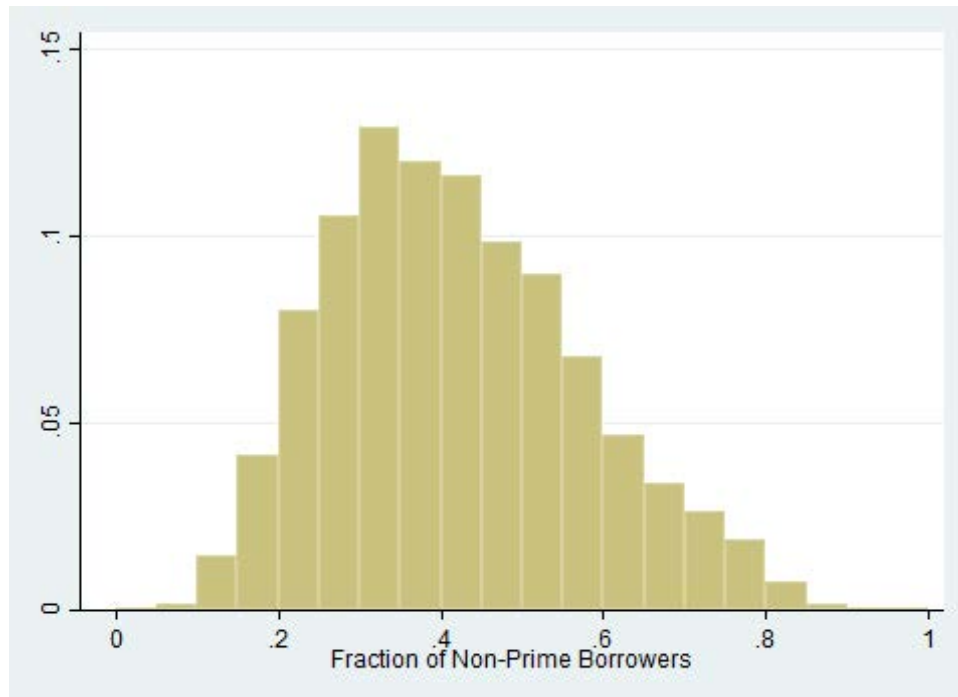
Mortgage	Prime				Non-Prime				Total
	LTV0	LTV1	LTV2	LTV3	LTV0	LTV1	LTV2	LTV3	
Fixed Rate	5.7%	15.4%	16.3%	3.4%	0.7%	2.9%	7.0%	4.4%	55.8%
ARM < 5yr	0.4%	0.5%	1.0%	0.3%	0.2%	0.2%	1.2%	0.9%	4.6%
ARM ≥ 5yr	0.2%	0.8%	1.8%	0.3%	0.1%	0.1%	0.4%	0.1%	3.9%
CE Second	0.4%	1.5%	3.4%	0.5%	0.2%	0.4%	1.8%	0.5%	8.6%
HELOC	3.2%	9.8%	9.2%	0.8%	0.3%	1.2%	2.4%	0.4%	27.0%
Total	9.9%	27.7%	31.7%	5.4%	1.3%	4.9%	12.8%	6.3%	100.0%

Notes: See the main text for the definitions of the categories.

Second, we use the imputed first-lien LTV ratio as of December 2006 to create four groups: those with a 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. Over half of the households have a LTV ratio below 50%. About a third of households have a LTV ratio between 50% and 80% and about 13% of households have a LTV ratio of 80% or higher. Not surprisingly, LTV ratio and Prime Status (or Equifax Risk Score) are somewhat negatively correlated: while the ratio of prime to non-prime is 3 to 1 in the general population, it is over 5 to 1 for LTV0 and about 2 to 1 for LTV3.

Finally, CRISM contains more detailed information about the type of the mortgages the households have. Using this information, 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 five**

Figure A-1: Distribution of Fraction of Non-prime Borrowers Across ZIP Codes



years or greater than or equal to five years (and 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. Over 55% of households have only a fixed-rate first mortgage and no second mortgage and a large majority of these households are prime. Only 8.5% of households have an ARM and about 9% of households have a closed-end second mortgage with about two prime households for every non-prime household. Finally 27% of households have a HELOC with an overwhelming majority being prime. We drop individuals (about 1% of our sample) that have both types of second mortgages.

A.1.3 Consumption Proxy and Cash for Clunkers

Astute readers may recall that a government rebate program designed to stimulate new car sales called Cash for Clunkers (CfC) was in effect in July-August 2009 and

one may be worried that this could influence the usefulness of our consumption proxy. Based on the results of [Hoekstra et al. \(2017\)](#), we conclude that CfC may have only slightly increased our measure by moving a small amount of sales (around 3%) that would have occurred in early 2010 to 2009. Moreover, we find that at the state level, the state’s share of all CfC registrations is uncorrelated with change in house prices between 2006 and 2009. Thus we conclude that CfC is unlikely to influence our results in any meaningful way.

A.2 Wealth Imputation using Survey of Consumer Finances

We use the Survey of Consumer Finances (SCF) to impute non-housing assets of households. To do this, using the 2007 wave for SCF, we regress non-housing assets on variables that are common in both the SCF and CCP/CRISM: age (+), income (+), auto loan or student loan balances (+), number of auto loans (+), an indicator for having a second mortgage (+), balance of second mortgage (−), having a credit card (+), aggregate credit credit balance (+), aggregate credit card limit (+), credit utilization (−), first mortgage balance (+), value of primary residence (+), and LTV ratio (−), where signs in parentheses show the signs we obtain in this regression.²⁹ The SCF regression has 8,765 observations and an R^2 of about 0.40. Using the estimated equation and the information we have for these right-hand side variables in CCP/CRISM, we compute the implied non-housing assets for each of our households. We then deduct the non-housing liabilities from CCP to get a measure of non-housing net worth. To account for the measurement error that may arise due to the imputation of assets, we create four equally-sized bins. Our computations yield mean and median non-housing net worth of \$144,196 and \$92,372, respectively. The 5 to 95 percentile range is \$10,753 to \$349,579.

²⁹When analyzing the two-mortgage borrowers in [Appendix C.2](#), we also include SCF households with two first mortgages, and control for this second home’s mortgage.

A.3 Aggregate Data

Table A-1 shows that unemployment rate increased by about 5.5 percentage points for the average county during this time period, with a 5-95 percentile range from 3% to 8.7%.

We also use a ZIP code-level measure of **Auto Sales in 2003**, which we compute by aggregating the individual-level auto loan origination variable. We start with the ZIP code-level sum of loan originations. Along the lines of Mian et al. (2013), we then allocate annual national retail auto sales (from the Census Bureau) across ZIP codes in proportion to their share of auto loan originations in our data; for example, if a ZIP code in our dataset accounted for 5% of all auto loan originations for that year, it would be allocated 5% of national retail auto sales. We then divide by the number of households in the ZIP code, which we obtain by applying the national population growth rate to the ZIP code populations in the 2000 census.

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). In particular we use Home Mortgage Disclosure Act (HMDA) and bank balance sheet data from call reports to identify the part of credit growth in a county that can be exclusively attributable to changes in credit supply. More specifically, we follow the approach in Gilchrist et al. (2023) and first regress the change in mortgage lending in a county, by a bank and in a year on a county-time fixed effect (to capture demand) and on a bank-time fixed effect (to capture supply). Next we project the bank-time fixed effect on bank balance sheet variables that capture bank health. This step ensures that we keep only the changes in bank credit supply that are related to banks' fundamentals. Finally, this bank-time variable is distributed to counties using the market share of each bank in each county. We obtain the credit supply shocks in 2006-2007, 2007-2008 and 2008-2009 and sum these to get the appropriate credit supply shock that corresponds to the period from 2006 to 2009. Table A-1 shows that while the mean and median of **Credit Supply** are negative at -2.8% and -4.8%, respectively, the

5-95 percentile range is very wide at -13.2% to 14.1%, indicating very different credit supply shocks across counties in this period.

Finally, as an alternative to the Credit Supply variable, we create a variable we call **Bank Health**. This is a county-level version of bank health indicator provided by [Chodorow-Reich \(2014\)](#), who uses, among others, a bank-level measure of the fraction of the syndication portfolio where Lehman Brothers had a lead role. Using information about the number of branches / affiliates each bank has in each of the U.S. counties, we distribute this measure to counties. The resulting variable shows each county's exposure to Lehman Brothers and, since this exposure is determined before 2008, it can also serve as an exogenous measure of bank health in each county relative to auto loan originations in 2009. Table [A-1](#) shows that this measure has a 5-95 percentile range of 0.42% to 0.78% with an average of 0.64%.

A.4 Local Housing Supply Instruments

We use two instruments which are meant to capture the elasticity of housing supply and therefore the response of house prices to demand shocks. First is the **Share of Land that is Unavailable for Real Estate Development**, from [Saiz \(2010\)](#), which measures the share of land within a 50km radius of the MSA centroid that cannot be developed based on geographic features. It ranges from 0.004 to 0.86 in our sample, with higher values corresponding to more unavailable land. The second instrument is 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.

Following [Mian and Sufi \(2009\)](#), [Mian and Sufi \(2010\)](#), [Mian and Sufi \(2011\)](#) and [Mian et al. \(2013\)](#), we use these instruments as a measure of housing supply elasticity. High values of both instruments indicate more inelastic housing supply. The adverse housing demand shock in the period we study represents a reversal of an earlier boom. Thus locations with more inelastic housing supply display larger declines in

house prices, since the previous boom lasted longer, and prices rose further, in these areas (see Glaeser et al. (2008)).


A.5 Panel IV Strategy

In Aladangady (2017), this national variable is the 10-year real interest rate, which is negatively correlated with the national house price changes in his sample of 1985-2008. This correlation is in fact stronger before the 2000s and in the 2000s it turns strongly positive. This suggests that it is likely not a good housing-demand shifter in the period of our analysis. Instead we use an annual measure of mortgage credit availability that we create along the lines of Anenberg et al. (2019). This measure is created using data on first-lien mortgage originations from ICE, McDash and CoreLogic to compute the maximal debt-to-income ratio (DTI) available to mortgage borrowers with a FICO score of 620 or less in that year.³⁰ A decline in this measure would indicate that borrowers with same risk and income are now able to borrow less than before. This measure is steady around 55% in the early 2000s and falls drastically between 2007 and 2010, reaching 25% by the end of the sample. This pattern closely matches that of annual national house price growth, with a correlation of 0.84.

Figure A-2 shows this measure in red (right scale), along with the aggregate house price growth in this period. The comovement between the two series is very clear (their correlation is 0.84) and while steady around 55% in the early 2000s, the maximum debt-to-income ratio falls drastically between 2007 and 2010, reaching 25% by the end of the sample.

³⁰FICO is at origination, from McDash.

Figure A-2: House Price Growth and Maximum DTI for Origination FICO ≤ 620



Notes: The two variables are plotted at an annual frequency.

B Details of the Decomposition of the Total Effect

Consider the **true** model (model A), written in matrix form as

$$y = x_1 b_1^A + X_2 b_2^A + u^A \quad (\text{A-1})$$

where y and x_1 are $N \times 1$ (N is the sample size), and X_2 matrix is $N \times n_2$. Here x_1 is the Change in House Value. In this demonstration $n_2 = 3$ and X_2 contains Change in Unemployment, Credit Supply and Bad Mortgage. We assume, without loss of generality, all variables are demeaned. We further assume that x_1 and X_2 are related via

$$X_2 = x_1 \gamma' + w \quad (\text{A-2})$$

where γ is a $n_2 \times 1$ vector and w is a $N \times n_2$ variable. Finally, we assume x_1 and X_2 are both exogenous satisfying $\mathbb{E}(x_1' u^A) = 0$ and $\mathbb{E}(X_2' u^A) = \mathbf{0}$, as well as $cov(w, u^A) = \mathbf{0}$, $cov(w, x_1) = 0$ and that $cov(w)$ is a diagonal matrix.

The **misspecified** model, one that drops X_2 (model B), is given by

$$y = x_1 b_1^B + u^B \tag{A-3}$$

B.1 OLS Case

The OLS estimate for b_1^B is given by

$$\hat{b}_1^B = (x_1' x_1)^{-1} x_1' y \tag{A-4}$$

Using the definition of y in model A and using \hat{u}^A to denote the residuals from the OLS estimation of (A-1), this can be written as

$$\hat{b}_1^B = \hat{b}_1^A + \hat{\gamma}' \hat{b}_2^A \tag{A-5}$$

where we use the definition $\hat{\gamma} = (x_1' x_1)^{-1} x_1' X_2$ and the result $x_1' \hat{u}^A = 0$, which follows from the properties of OLS estimation.

In order to complete the decomposition, we divide (A-5) by \hat{b}_1^B and expand the term $\hat{\gamma}' \hat{b}_2^A$ to obtain

$$1 = \frac{\hat{b}_1^A}{\hat{b}_1^B} + \frac{\hat{\gamma}'_1 \hat{b}_{2,1}^A}{\hat{b}_1^B} + \frac{\hat{\gamma}'_2 \hat{b}_{2,2}^A}{\hat{b}_1^B} + \frac{\hat{\gamma}'_3 \hat{b}_{2,3}^A}{\hat{b}_1^B} \tag{A-6}$$

where the first term on the right-hand side shows the share of the *total* effect of Change in House Value that is due to the *direct* effect and the remaining three terms show the share that is due to the *indirect* effects that is coming via each of the three additional variables.

In the OLS case we presented above, there were no other controls that is common between Model A and Model B for simplicity. Our full specification includes further controls such as the 2003 ZIP measure for auto sales or other individual controls like

age. Thus, for the derivations here, y is the *residual* from a regression that contains all these controls on the right hand side and the individual origination variable on the left hand side.

B.2 IV Case

Here we first continue with the IV case with one instrument and finally the general case we use in the paper, which is IV with k instruments. Consider the same **true** model (model A) as in (A-1) but now we have $\mathbb{E}(x_1' u^A) \neq 0$, violating the key condition for OLS to be valid. Through (A-2), we see that $\mathbb{E}(X_2' u^A) \neq \mathbf{0}$ also must hold, but we assume $\mathbb{E}(w' u^A) = \mathbf{0}$. We have an instrument Z , which is collected in a $N \times k$ matrix that satisfies $\mathbb{E}(Z' u^A) = \mathbf{0}$. Note that (A-2) can no longer be estimated consistently via OLS since x_1 may be correlated with ω , or in other words $\mathbb{E}(x_1' w) \neq 0$.

The second stage of Model B is still given by (A-3). In estimating this model, we ignore X_2 but we still instrument using Z . This means in the first stage we only use Z , and X_2 is omitted. Define $P_Z \equiv Z(Z'Z)^{-1}Z'$ and the IV estimate for b_1^B is given by

$$\hat{b}_1^B = (x_1' P_Z x_1)^{-1} x_1' P_Z y \quad (\text{A-7})$$

B.2.1 IV - Single Instrument

Even though our general case multiple instruments, we now focus on the case where $k = 1$, that is we have a single instrument. This will prove to be useful. The IV estimate for b_1^B can be further simplified

$$\hat{b}_1^B = (x_1' Z(Z'Z)^{-1}Z'x_1)^{-1} x_1' Z(Z'Z)^{-1}Z'y \quad (\text{A-8})$$

$$= (Z'x_1)^{-1}(Z'Z)(x_1'Z)^{-1}x_1'Z(Z'Z)^{-1}Z'y \quad (\text{A-9})$$

$$= (Z'x_1)^{-1}Z'y \quad (\text{A-10})$$

which we can do since $Z'Z$, $Z'x_1$ and x_1Z' are all square matrices of the same size. Note that, the IV estimator, written this way, solves for the b_1^B that satisfies $Z'\hat{u}^B = 0$.³¹

Using the definition of y in model A, this can be written as

$$\hat{b}_1^B = (Z'x_1)^{-1}Z'(x_1\hat{b}_1^A + X_2\hat{b}_2^A + \hat{u}^A) \quad (\text{A-11})$$

$$= (Z'x_1)^{-1}Z'x_1\hat{b}_1^A + (Z'x_1)^{-1}Z'X_2\hat{b}_2^A + (Z'x_1)^{-1}Z'\hat{u}^A \quad (\text{A-12})$$

$$= \hat{b}_1^A + \hat{\gamma}'b_2^A \quad (\text{A-13})$$

where the last term in the second line drops out because $Z'\hat{u}^A = 0$ and $\hat{\gamma}'$ is the IV estimate of γ' in (A-2) with x_1 instrumented by Z , $\hat{\gamma}' = (Z'x_1)^{-1}Z'X_2$. The bias between the estimates from the two models in this case is given by $\hat{\gamma}'b_2^A$, which is the same expression as in the OLS case except, of course, now $\hat{\gamma}$ is computed using IV.

B.2.2 IV - Multiple Instruments

If $k > 1$ then the system is over-identified and the simplifications in (A-10) will not hold. Thus the generalized version of (A-11) is given by

$$\hat{b}_1^B = (x_1'P_Zx_1)^{-1}x_1'P_Z(x_1\hat{b}_1^A + X_2\hat{b}_2^A + \hat{u}^A) \quad (\text{A-14})$$

$$= (x_1'P_Zx_1)^{-1}(x_1'P_Zx_1)\hat{b}_1^A + (x_1'P_Zx_1)^{-1}x_1'P_ZX_2\hat{b}_2^A + (x_1'P_Zx_1)^{-1}x_1'P_Z\hat{u}^A \quad (\text{A-15})$$

$$= \hat{b}_1^A + \hat{\gamma}'\hat{b}_2^A + \hat{\delta} \quad (\text{A-16})$$

where once again we use $\hat{\gamma}'$ to represent the IV estimate of γ' in (A-2) as $\hat{\gamma}' = (x_1'P_Zx_1)^{-1}x_1'P_ZX_2$ and define $\hat{\delta} \equiv (x_1'P_Zx_1)^{-1}x_1'P_Z\hat{u}^A$. It is easy to see that $\hat{\delta}$ refers to the IV estimate of regressing \hat{u}^A on x_1 with instruments Z . While asymptotically $\mathbb{E}(Z'u^A) = 0$ would hold and $\hat{\delta} \rightarrow 0$, in finite samples, $\hat{\delta}$ will not drop out from this expression since $Z'\hat{u}^A \neq 0$.

³¹Similarly, though we do not explicitly use, the IV estimation of Model A sets $Z'\hat{u}^A = 0$ and $X_2'\hat{u}^A = 0$.

In order to do the decomposition, we proceed as follows. Rewrite (A-16) as

$$\hat{b}_1^B - \hat{\delta} = \hat{b}_1^A + \hat{\gamma}'\hat{b}_2^A \quad (\text{A-17})$$

where we consider the left-hand side of the equation to be the total effect of x_1 on y and the two terms on the right-hand side as the direct effect of house value changes on consumption and the indirect effect of X_2 that comes via house value changes, respectively. It is convenient to report these as shares and we use $\hat{b}_1^A/(\hat{b}_1^B - \hat{\delta})$ as the share of the total effect that's direct and $\hat{\gamma}'\hat{b}_2^A/(\hat{b}_1^B - \hat{\delta})$ as the share of the total effect that's indirect and due to X_2 .³²

One final practical note is about the interaction of the all controls with categorical dummy variables. As we explain in Section 3.2 we allow all heterogenous effects in the first and second stage (except for the ΔHV coefficient in the second stage) with respect to the eight household categories. This means that each of the three main controls we have, Change in Unemployment Rate, Credit Supply and Bad Mortgage, are interacted with eight dummy variables. This means β_2^A actually isn't a $n_2 = 3$ dimensional vector but it has $n_2 = 24$ elements, eight for each of the controls, corresponding to one of the categories. In order to compute $\hat{\gamma}'\hat{b}_2^A$, then, for each of the controls we create eight versions, each interacted with a specific dummy variable and run an IV estimation of this variable on ΔHV . Then the part of $\hat{\gamma}'\hat{b}_2^A$ that is due to a particular control is the weighted average of the relevant eight terms in this multiplication, using the sample weights.

³²An alternative is to follow the approach in [Chen et al. \(2016\)](#) and compute the decomposition twice, each with only one of the instruments and take the average. Doing so does not alter the results in a meaningful way.

C Detailed Results

C.1 Results without Interactions

All the results in the paper allow for heterogeneous effects of all controls except for the Change in House Value, which means all controls are interacted with a set of eight dummy variables. In this Appendix we remove these interactions to report how our main results change. Table A-3 is the counterpart of Table 1. The main coefficients of interest, those for Change in House Value are about 0.007 higher, which corresponds to about 0.53 percentage points in the probability of originating an auto loan. This suggests that not allowing for the heterogeneity of the effects of other controls introduces a significant bias to the coefficient of Change in House Value. Inspecting the other coefficients, the effect of Credit Supply is essentially unchanged and the coefficient of Bad Mortgage is higher by 0.4, indicating a 0.4 percentage point increase in the effect of this variable. The coefficient of Change in Unemployment Rate is smaller (in absolute value) by 0.0009, which correspond to a roughly 0.51 percentage point reduction in the effect of this variable. Using these estimates we repeat the decomposition reported in Table 2. We get the following decomposition (with results in Table 2 in parenthesis for convenience) : Pure Wealth and Other Constraints : 66% (53%), Local General Equilibrium : 4% (13%), Household Credit Supply : 10% (11%) and Bad Mortgage : 20% (23%). Consistent with how the estimated coefficients changed, the biggest change is in the Local General Equilibrium channel (reduced significantly) and the Pure Wealth and Other Constraints channel (increased by 13 p.p).

Given that the model in column (4) in Table A-3 is nested in the model reported in column (4) in Table 1, we can run a simple Wald test to test if relaxing the restrictions in the model with no interactions is warranted. The Wald test statistic is 13.12 which is distributed as $F(70, 83)$ (83 parameters in the unrestricted model and 70 parameters restrictions.) The critical value for a one-sided test at 0.1% significance would be 2.03. The p-value of the test statistic is 0 and as such the restrictions are

rejected at any level of significance. A likelihood ratio test yields a test statistic of 974.28, which is distributed as $\chi^2_{(70)}$ and it also has a p-value of 0. Therefore we conclude that the results we report in the main text, including the decomposition results are the appropriate ones to look a

Table A-3: Main Results with No Interactions

Marginal Effects				
	(1)	(2)	(3)	(4)
Change in House Value	0.0243*** (0.0021)	0.0260*** (0.0032)	0.0226*** (0.0036)	0.0173*** (0.0036)
Change in Unemployment Rate	-	-0.0005 (0.0007)	-0.0006 (0.0007)	-0.0006 (0.0007)
Credit Supply	-	-	0.0316*** (0.0104)	0.0320*** (0.0103)
Bad Mortgage	-	-	-	-0.0869*** (0.0018)
Marginal Effects (in p.p.)				
Δ HV (average: -\$75,109)	-1.82	-1.96	-1.70	-1.30
Δ U (average: -5.5 p.p.)	-	-0.28	-0.33	-0.31
Credit Supply (-1 s.d.)	-	-	-0.28	-0.28
Bad Mortgage (= 1)	-	-	-	-8.69
Marginal Effects (in Dollars)				
Δ HV (average: -\$75,109)	-\$1,689	-\$1,812	-\$1,570	-\$1,202
Δ U (average: -5.5 p.p.)	-	-\$256	-\$310	-\$284
Credit Supply (-1 s.d.)	-	-	-\$257	-\$261
Bad Mortgage (= 1)	-	-	-	-\$8,044

Notes: See the notes for Table 1.

C.2 Results for Consumers with Two Mortgages

In this section we consider consumers with more than one first lien. In order to avoid including large investors, whose behavior would be very different, we limit our analysis to those with two such mortgages. This restriction also allows for a cleaner wealth imputation for these borrowers. Together with the single-mortgage borrowers, we cover 93% of all mortgage-holders in the dataset. The following tables present key summary statistics for these two-mortgage borrowers.

(a) Individual-Level Variables: Two-Mortgage Borrowers

	Mean
Originate Auto Loan in 2009	0.153
Change in House Value (\$1,000)	-97.2
Bad Mortgage	0.099
2006 Non-Housing Net Worth (\$1,000)	243.3

(b) Distribution of Characteristics: Two-Mortgage Borrowers

LTV Category	Prime	Non-Prime	Total
LTV0 (LTV ratio less than 25%)	8.7%	1.0%	9.7%
LTV1 (LTV ratio between 25% and 50%)	26.6%	4.0%	30.6%
LTV2 (LTV ratio between 50% and 80%)	38.2%	12.3%	50.4%
LTV3 (LTV ratio greater than 80%)	6.2%	3.1%	9.3%
Total	79.7%	20.3%	100.0%

As expected, two-mortgage borrowers are of better ex-ante credit quality than the single-mortgage borrowers in our main analysis, wealthier, and more likely to originate auto loans. But they also experience larger declines in house value. In addition, they are likelier to default on their mortgages, which is consistent with [Elul et al. \(2023\)](#).

For these borrowers, we augment the main data from CRISM by merging their other mortgage's balance from CCP. We (i) use this information to control for the fact that they have another mortgage in the results below; and also (ii) use this other mortgage balance in our wealth imputation.

Results are reported in Table A-4. In panel (a) we show results for consumers that have one or two mortgages and in panel (b) we show results for consumers with only two mortgages. Results in panel (a) are very close to those in Table 1.

Table A-4: Main Results - By Number of Mortgages

(a) One or Two Mortgages				
Marginal Effects				
	(1)	(2)	(3)	(4)
Change in House Value	0.0171*** (0.0021)	0.0187*** (0.0032)	0.0159*** (0.0036)	0.0091** (0.0036)
Change in Unemployment Rate	-	-0.0013* (0.0007)	-0.0014** (0.0007)	-0.0016** (0.0007)
Credit Supply	-	-	0.0250** (0.0103)	0.0317*** (0.0103)
Bad Mortgage	-	-	-	-0.0832*** (0.0029)
Marginal Effects (in p.p.)				
Δ HV (average: -\$75,109)	-1.28	-1.40	-1.19	-0.68
Δ U (average: -5.5 p.p.)	-	-0.71	-0.77	-0.85
Credit Supply (-1 s.d.)	-	-	-0.22	-0.27
Bad Mortgage (= 1)	-	-	-	-8.32
Marginal Effects (in Dollars)				
Δ HV (average: -\$75,109)	-\$1,186	-\$1,298	-\$1,105	-\$630
Δ U (average: -5.5 p.p.)	-	-\$657	-\$715	-\$792
Credit Supply (-1 s.d.)	-	-	-\$204	-\$246
Bad Mortgage (= 1)	-	-	-	-\$7,706

Notes: See the notes for Table 1. All equations have a sample size of $N = 365,498$. First stage F-statistics are 167, 108, 97 and 95, respectively.

(b) Only Two Mortgages (continued)

Marginal Effects				
	(1)	(2)	(3)	(4)
Change in House Value	0.0143*** (0.0056)	0.0157** (0.0080)	0.0153* (0.0091)	0.0080 (0.0092)
Change in Unemployment Rate	-	-0.0001 (0.0024)	-0.0001 (0.0023)	-0.0001 (0.0023)
Credit Supply	-	-	0.0012 (0.0445)	0.0072 (0.0442)
Bad Mortgage	-	-	-	-0.0872*** (0.0029)
Marginal Effects (in p.p.)				
Δ HV (average: -\$75,109)	-1.07	-1.18	-1.15	-0.60
Δ U (average: -5.5 p.p.)	-	-0.04	-0.08	-0.03
Credit Supply (-1 s.d.)	-	-	-0.01	-0.06
Bad Mortgage (= 1)	-	-	-	-8.72
Marginal Effects (in Dollars)				
Δ HV (average: -\$75,109)	-\$995	-\$1,092	-\$1,064	-\$559
Δ U (average: -5.5 p.p.)	-	-\$35	-\$76	\$28
Credit Supply (-1 s.d.)	-	-	-\$10	-\$59
Bad Mortgage (= 1)	-	-	-	-\$8,080

Notes: All equations have a sample size of $N = 20,465$. First stage F-statistics are 110, 77, 66 and 64, respectively.

C.3 Details of Robustness Results

Table A-5 presents the corresponding coefficient estimates for the marginal effects in percentage points reported in Table 3.

Table A-5: Coefficient Estimates for Robustness Results

	Marginal Effects				
	Baseline		Bank Health		Probit
	(1)	(2)	(3)	(4)	(5)
Change in House Value	0.0176*** (0.0022)	0.103*** (0.0039)	0.108*** (0.0042)	0.0042* (0.0023)	-0.0011 (0.0041)
Change in Unemployment Rate	-	-0.0015** (0.0007)	-0.0018** (0.0008)	-	-0.0030*** (0.0007)
Credit Supply	-	0.0312*** (0.0105)	-	-	0.0467*** (0.0102)
Bank Health	-	-	-2.126** (0.8534)	-	-
Bad Mortgage	-	-0.0829*** (0.0031)	-0.0828*** (0.0031)	-	-0.0844*** (0.0029)

Notes: See the notes to Tables 1 and 3.

C.4 Ex-Ante Constraints - Type of Mortgage

Tables A-6 reports the detailed estimation results (omitting all controls) for the marginal effects in percentage points presented in Table 6.

Table A-6: Ex-Ante Constraints - Type of Mortgage

(a) Fixed First Mortgage, No Second

LTV	Non-Prime				Prime			
	0	1	2	3	0	1	2	3
Δ HV	0.0702* (0.0387)	0.0121 (0.0200)	0.0479*** (0.0176)	0.0576* (0.0327)	0.0140 (0.0115)	0.0083 (0.0096)	0.0207* (0.0118)	0.0019 (0.0329)
p-value	0.070	0.548	0.006	0.078	0.223	0.384	0.078	0.954
First Stage Signs	Negative	Negative	Negative	Negative	Negative	Negative	Negative	Negative
First Stage F-stat	23.01	214.50	471.36	371.66	84.14	296.94	438.39	305.59
N	2,390	10,079	24,061	15,221	19,582	53,131	56,133	11,820

(b) ARM <5 Years, No Second

LTV	Non-Prime				Prime			
	0	1	2	3	0	1	2	3
Δ HV	0.1070* (0.0612)	-0.0231 (0.0587)	-0.0302 (0.0349)	0.0804* (0.0431)	0.0032 (0.0274)	-0.0520 (0.0411)	0.0171 (0.0310)	-0.0716 (0.0650)
p-value	0.081	0.694	0.388	0.062	0.263	0.206	0.581	0.270
First Stage Signs	Neg + Insig	Negative	Negative	Negative	Neg + Insig	Neg + Insig	Negative	Neg + Insig
First Stage F-stat	8.57	21.52	131.93	153.29	25.51	20.31	94.25	58.82
N	523	842	4,146	2,944	1,287	1,565	3,298	1,194

(c) ARM \geq 5 Years, No Second

LTV	Non-Prime				Prime			
	0	1	2	3	0	1	2	3
Δ HV	0.0306 (0.0973)	0.0339 (0.0991)	-0.0303 (0.0545)	-0.0945 (0.1303)	-0.0130 (0.0571)	0.0405 (0.0249)	0.0334 (0.0254)	0.1235 (0.0932)
p-value	0.753	0.732	0.578	0.469	0.820	0.104	0.189	0.185
First Stage Signs	Neg + Insig	Neg + Insig	Negative	Neg + Insig	Marginal	Neg + Insig	Negative	Negative
First Stage F-stat	3.34	4.84	61.26	17.92	1.48	25.24	80.96	32.43
N	94	375	1,367	479	819	2,878	6,359	1,059

Table A-6: Ex-Ante Constraints - Type of Mortgage (continued)

(d) Closed-End Second

LTV	Non-Prime				Prime			
	0	1	2	3	0	1	2	3
Δ HV	-0.1737 (0.1966)	0.0338 (0.0741)	-0.0068 (0.0406)	0.1334* (0.0762)	-0.1101 (0.1163)	0.0060 (0.0414)	0.0295 (0.0319)	0.2734*** (0.1020)
p-value	0.377	0.649	0.867	0.080	0.344	0.886	0.355	0.007
First Stage Signs	Neg + Insig	Negative	Negative	Negative	Insig	Negative	Negative	Neg + Marginal
First Stage F-stat	4.54	55.83	240.65	79.15	6.78	87.34	232.78	41.71
N	347	1,429	6,342	1,654	1,299	5,186	11,806	1,690

(e) HELOC

LTV	Non-Prime				Prime			
	0	1	2	3	0	1	2	3
Δ HV	0.0118 (0.0460)	-0.0423 (0.0278)	0.0114 (0.0226)	-0.0253 (0.0641)	-0.0037 (0.0125)	-0.0001 (0.0091)	-0.0007 (0.0105)	0.0415 (0.0517)
p-value	0.798	0.129	0.615	0.693	0.765	0.994	0.947	0.422
First Stage Signs	Negative	Negative	Negative	Negative	Negative	Negative	Negative	Neg + Insig
First Stage F-stat	25.12	138.31	298.77	84.91	60.72	207.48	326.43	66.59
N	1,028	4,263	8,410	1,389	11,150	32,852	31,829	2,713

Notes: See notes to Table 1. This table only reports the coefficients for Δ HV, the sign of instruments and the F-stat in the first stage and the number of observations in each estimation.